

Classification of genus-two surfaces in S^3

Filippo Baroni

Abstract

We describe an algorithm to decide whether two genus-two surfaces embedded in the 3-sphere are isotopic or not. The algorithm employs well-known techniques in 3-manifolds topology, as well as a new algorithmic solution to a problem on free groups.

Contents

1	Introduction	2
1.1	The classification problem	2
1.2	Notation	4
1.3	Classical algorithms on 3-manifolds	5
1.4	Examples	10
2	Homeomorphisms of 3-manifolds	11
2.1	Rationale	11
2.2	JSJ decomposition	13
2.3	Seifert pieces	15
2.4	I -bundle pieces	20
2.5	Simple pieces	22
2.6	Computing the JSJ decomposition	27
2.7	Piecing things together	29
3	An algorithmic problem on free groups	33
3.1	Introducing band systems	33
3.2	Band systems and free groups	38
3.3	Topological applications	46
4	The classification algorithm	48
4.1	Outline	48
4.2	General strategy for canonical compression discs	49
4.3	Compression discs for the boundary	51
4.4	One non-separating compression disc	53
4.5	One separating compression disc	55
4.6	Small I -bundles in the JSJ decomposition	59
4.7	Product I -bundle over punctured torus	60
4.8	Twisted I -bundle over punctured Klein bottle	63
	References	64

1 Introduction

1.1 The classification problem

In mathematics, by “classification” we usually mean a list, finite or infinite, of all objects of a given type, up to a given equivalence relation. For instance, the list

$$\left\{ \mathbb{Z}/4\mathbb{Z}, \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z} \right\}$$

is a classification of groups of order 4 up to isomorphism, and the list

$$\left\{ \underbrace{\text{[diagram of } g \text{ holes]}}_{g \text{ holes}} : g \in \mathbb{Z}_{\geq 0} \right\}$$

is a classification of closed orientable surfaces up to homeomorphism. Naturally, we want the description of such a list to be somewhat explicit: the reader will surely agree that

$$\{\text{finitely generated Abelian groups up to isomorphism}\}$$

is not a classification of finitely generated Abelian groups up to isomorphism, but

$$\left\{ \mathbb{Z}^k \times \mathbb{Z}/a_1\mathbb{Z} \times \cdots \times \mathbb{Z}/a_n\mathbb{Z} : k \in \mathbb{Z}_{\geq 0}, n \in \mathbb{Z}_{\geq 0}, a_1, \dots, a_n \in \mathbb{Z}_{\geq 2}, a_1 \mid \cdots \mid a_n \right\}$$

is. This example raises the question of what degree of explicitness is required for something to be considered a classification.

For instance, let us consider the set of knots in the 3-sphere: can we classify them up to isotopy?

- A task that is surely within our grasp is deciding whether two given knots in S^3 are isotopic. In fact, two knots being isotopic is equivalent to their two complements being homeomorphic via a homeomorphism which is meridian-preserving; this can be decided algorithmically thanks to Theorem 1.3.
- Moreover, it is not hard to devise an algorithm which produces an infinite list \mathcal{L}_0 of knot diagrams, so that every knot is isotopic to one element (and possibly more) of our list; this follows from the fact that there are only finitely many knot diagrams with n crossings up to isotopy for every $n \geq 0$, and they can be enumerated algorithmically.

We can combine these two ingredients to produce:

- 1 a computable list \mathcal{L} of knots containing exactly one representative for each isotopy class, by taking knots in \mathcal{L}_0 which are not isotopic to elements appearing earlier in the list;
- 2 an algorithm which takes a knot as input and returns the unique knot in \mathcal{L} which is isotopic to it.

Given how complex the world of knots is, we cannot expect to find a classification which is as neat as the ones given above for surfaces or Abelian groups. This algorithmic answer is probably the best we can hope for, and we believe it is explicit enough to be considered a classification.

We hope that, in light of the above discussion, the reader will accept the following definition as a sensible one. A *classification* of a set \mathcal{X} up to an equivalence relation \sim is the datum of

- 1 an algorithm producing a list \mathcal{L} containing exactly one representative for each equivalence class of \mathcal{X}/\sim , and
- 2 an algorithm which takes an element of \mathcal{X} as input and returns the unique element of \mathcal{L} equivalent to it.

Like in the knots example, in order to have a classification of \mathcal{X} up to \sim , it is enough to provide a computable list \mathcal{L}_0 containing *at least* one representative for each equivalence class of \mathcal{X}/\sim , and an algorithm to decide equivalence; since the second ingredient is usually the hardest to find, we call it the *classification algorithm*.

Going back to the case of knots, a routine topological argument shows that the knot classification problem is equivalent to the classification of tori embedded in S^3 up to isotopy: the one-to-one correspondence is given by associating each knot to the boundary of its regular neighbourhood. It is then only natural to try and generalise this question to higher genus surfaces¹. In fact, the following appears as Problem 3.11 in Kirby's problem list.

Problem ([13, Problem 3.11]). *Classify embeddings of orientable surfaces in S^3 .*

This is precisely the aim of this article, which solves a special case of the above problem by providing a classification of genus-two surfaces in S^3 . Producing a redundant list \mathcal{L}_0 is very easy, by simply enumerating all simplicial genus-two surfaces in all subdivisions of S^3 (see Section 1.3). Therefore, the only goal of this article can be summarised as follows.

Theorem. *There is an algorithm to decide whether two genus-two surfaces in S^3 are isotopic.*

The article is structured as follows. In the rest of Section 1, we introduce the notation and terminology we will use in our exposition, and we briefly discuss some foundational results in low-dimensional algorithmic topology. We also give a few examples of phenomena which make the genus-two classification problem intrinsically harder than the knot classification problem. Section 2 is devoted to the computation of mapping class groups of 3-manifolds. The exposition in this section closely follows the work of Johannson in [12], revisited to provide constructive and effective proofs. The arguments here are somewhat technical and involved, and may be skipped on a first reading; the main result of this section, that is Theorem 2.16, can be assumed as a black box without compromising the understanding of the rest of the article. In Section 3, we present a solution to an algorithmic problem on free groups, as well as an application to a topological decision problem. Finally, in Section 4, we focus on the genus-two classification problem. A careful case analysis and elementary topological arguments, combined with Theorem 2.16 and Corollary 3.3, allow us to prove the theorem stated above; a more detailed outline of the inner workings of the classification algorithm is given at the end of Section 4.1.

¹It should be noted that the classification of genus-zero surfaces in S^3 is trivial, since by Alexander's theorem there is only one isotopy class of embedded 2-spheres (see [1]).

1.2 Notation

✿ Manifolds.

- ⦿ Throughout this article, we will always be working in the PL category: all manifolds will have a PL structure, and functions between manifolds will be assumed to be PL.
- ⦿ All 3-manifolds will be compact, connected and oriented unless otherwise stated. Codimension-zero submanifolds and boundaries of 3-manifolds will be implicitly oriented accordingly.
- ⦿ All surfaces (2-manifolds) will be compact, but we make no assumptions about connectedness or orientability.
- ⦿ For integers $g \geq 0$ and $k \geq 0$, we denote by $\Sigma_{g,k}$ the connected orientable surface of genus g with k holes (also referred to as “punctures”). If $g \geq 1$, we denote by $U_{g,k}$ the surface obtained by taking the connected sum of g projective planes and then removing k open discs. In both cases, the integer k may be omitted when it is equal to 0. We will occasionally refer to the torus Σ_2 by the symbol T^2 .
- ⦿ If X is a manifold and $Z \subseteq X$ is a subspace, we denote by $\text{clos}(Z)$ and $\text{int}(Z)$ its closure and interior in X respectively; the ambient space X will always be clear from the context.

✿ Fibre bundles.

- ⦿ For a fibre bundle $p: M \rightarrow B$ with fibre F , we denote by $\partial_h M$ its horizontal boundary, that is the induced ∂F -bundle on B ; we will refer to the vertical boundary $p^{-1}(\partial B)$ as $\partial_v M$.
- ⦿ We denote by I the 1-manifold $[0, 1]$. For every surface B , there is exactly one (orientable) I -bundle over B ; we denote it by $B \times I$ (the “product I -bundle”) if B is orientable, and by $B \tilde{\times} I$ (the “twisted I -bundle”) if it is not.
- ⦿ If X is a manifold and $Z \subseteq X$ is a properly embedded submanifold, we denote by $\mathcal{N}(Z)$ and $\overset{\circ}{\mathcal{N}}(Z)$ the closed and open regular neighbourhoods of Z in X respectively. The closed neighbourhood $\mathcal{N}(Z)$ is naturally endowed with a bundle structure over Z with fibre the $(\text{codim } Z)$ -dimensional disc.

✿ Surfaces in 3-manifolds.

- ⦿ If X is a manifold and $Z \subseteq X$ is a properly embedded codimension-one submanifold, we define $X \setminus\setminus Z = X \setminus \overset{\circ}{\mathcal{N}}(Z)$ to be the result of cutting X along Z ; if Z is separating, the manifold $X \setminus\setminus Z$ will be disconnected.
- ⦿ As far as the notions of irreducible, boundary irreducible, incompressible, boundary incompressible, and sufficiently large are concerned, we adopt the same definitions as [20]. Unless otherwise stated, by “(boundary) compression disc” we will always mean “non-trivial (boundary) compression disc”.

✿ Homeomorphisms.

- ⦿ Unless otherwise stated, homeomorphisms between 3-manifolds will be orientation-preserving.

- If X and Y are manifolds, and A_1, \dots, A_n and B_1, \dots, B_n are subspaces, a function

$$f: (X, A_1, \dots, A_n) \longrightarrow (Y, B_1, \dots, B_n)$$

is a function $X \rightarrow Y$ such that $f(A_i) \subseteq B_i$ for $1 \leq i \leq n$. We say that f is a homeomorphism if it is bijective and $f(A_i) = B_i$ for $1 \leq i \leq n$. When the number n is clear from the context or irrelevant, we will write (X, \mathbf{A}) as a shorthand for (X, A_1, \dots, A_n) .

- If X and Y are oriented, we denote by $\mathbf{H}((X, \mathbf{A}); (Y, \mathbf{B}))$ the set of homeomorphisms from (X, \mathbf{A}) to (Y, \mathbf{B}) , modulo isotopies through homeomorphisms of the same kind. Note that $\mathbf{H}((X, \mathbf{A}); (X, \mathbf{A}))$ has a natural group structure; we use the shorthand $\mathbf{H}(X, \mathbf{A})$ to refer to it. Moreover, if $S \subseteq X$ is a subspace, we denote by $\mathbf{H}_S(X, \mathbf{A})$ the group of self-homeomorphisms of (X, \mathbf{A}) fixing S pointwise, modulo isotopies through homeomorphisms of the same kind.
- If $f: (X, \mathbf{A}) \rightarrow (Y, \mathbf{B})$ is a function, we define its *trace*

$$f|_{\partial}: (\partial X, \mathbf{A} \cap \partial X) \longrightarrow (\partial Y, \mathbf{B} \cap \partial Y)$$

to be its restriction to the boundary. The trace behaves well with isotopies – that is, isotopic homeomorphisms have isotopic traces. As a consequence, when X and Y are oriented, we have a well-defined *trace map*

$$(-)|_{\partial}: \mathbf{H}((X, \mathbf{A}); (Y, \mathbf{B})) \longrightarrow \mathbf{H}((\partial X, \mathbf{A} \cap \partial X); (\partial Y, \mathbf{B} \cap \partial Y)),$$

which is functorial in the sense that it preserves composition. For the sake of convenience, if W is a set of (isotopy classes of) homeomorphisms, we define $W|_{\partial} = \{f|_{\partial} : f \in W\}$.

• Dehn twists.

- In the most general sense, a Dehn twist of a manifold X about a properly embedded two-sided codimension-one submanifold Z is a self-homeomorphism of X which is the identity outside $\mathcal{N}(Z)$.
- When Z is an essential two-sided curve in a surface X , the group of Dehn twists about Z modulo isotopies is isomorphic to \mathbb{Z} . If X is oriented there is a way to select a preferred generator of this group, which we call *the* Dehn twist about Z and denote by τ_Z . If X is orientable without a preferred orientation, fixing one arbitrarily allows us to pick the generating Dehn twists about different curves consistently. When X is non-orientable, instead, we arbitrarily call one of the generating Dehn twists τ_Z , so that the other one is τ_Z^{-1} .
- Similarly, if Z is an incompressible boundary incompressible annulus in an (oriented) 3-manifold X the group of Dehn twists about Z modulo isotopies is isomorphic to \mathbb{Z} . In this case, there is no preferred generator of this group, so we arbitrarily pick one to call τ_Z . We remark that, if z_1 and z_2 are the two boundary curves of Z , then $\tau_Z|_{\partial}: \partial X \rightarrow \partial X$ is equal to either $\tau_{z_1} \tau_{z_2}^{-1}$ or $\tau_{z_1}^{-1} \tau_{z_2}$, depending on the choice of τ_Z .

1.3 Classical algorithms on 3-manifolds

• **Triangulations of 3-manifolds.** We are, of course, interested in algorithms on 3-manifolds. The first issue we should address is, perhaps, what a 3-manifold *is* from

an algorithmic point of view. The simplest way to think about a 3-manifold M in this setting is seeing it as a collection of 3-simplices (or “tetrahedra”), with some pairs of faces identified by simplicial isomorphisms. Such a description is called a *triangulation* of M , and it fully encodes the topology of M in a discrete and combinatorial fashion. Of course, from the perspective of algorithmic implementation, the name “tetrahedra” is purely suggestive: 3-simplices are encoded as ordered 4-tuples of “vertices” – represented, for example, by integers – and an identification between faces can be specified by two ordered triples of vertices, each describing a face of a tetrahedron.

Some care must be taken in order to ensure that the topological space M obtained by gluing the tetrahedra is actually a manifold. We deal with orientation by orienting the tetrahedra, and requiring that gluing maps are orientation-reversing. Then M is a 3-manifold if and only if the link of each vertex is a disc – for vertices on the boundary – or a sphere. Note that both conditions can be easily checked algorithmically.

✿ **Simplicial subdivisions.** Ideally, we would want all the “objects” we have to deal with algorithmically (namely, maps and submanifolds) to be simplicial; in other words, we want to describe them discretely at the level of simplices. Fortunately, in the world of 3-manifolds, everything can be made simplicial, up to subdivision and isotopy.

Definition 1.1. Let \mathcal{T} be a triangulation of a 3-manifold M . A *subdivision* of \mathcal{T} is a triangulation \mathcal{T}' of a 3-manifold M' together with a choice, for each vertex v of \mathcal{T}' , of a tetrahedron $\Delta(v)$ of \mathcal{T} and four non-negative rational numbers $q_1(v)$, $q_2(v)$, $q_3(v)$ and $q_4(v)$, satisfying the following constraints.

- i For every vertex v of \mathcal{T}' , the equality $q_1(v) + q_2(v) + q_3(v) + q_4(v) = 1$ holds.
- ii Fix a vertex v of \mathcal{T}' , and denote by \mathbf{u}_1 , \mathbf{u}_2 , \mathbf{u}_3 and \mathbf{u}_4 the vertices of $\Delta(v)$, which we interpret as the four canonical base vectors in \mathbb{R}^4 . If we think of $\Delta(v)$ as the convex hull of its vertices, then there is a unique point

$$p(v) = \sum_{i=1}^4 q_i(v) \mathbf{u}_i \in \Delta(v) \subseteq M.$$

We require that, if v_1 , v_2 , v_3 and v_4 are vertices of a tetrahedron of \mathcal{T}' , then there is some tetrahedron of \mathcal{T} containing $p(v_1)$, $p(v_2)$, $p(v_3)$ and $p(v_4)$.

- iii Define a function $f: M' \rightarrow M$ by extending the map $v \mapsto p(v)$ linearly on the tetrahedra of \mathcal{T}' ; the previous constraint guarantees that f is well-defined. We require that f is a homeomorphism.

With slight abuse of notation, we will omit the choices of $\Delta(v)$ and $q_i(v)$, and simply call the triangulation \mathcal{T}' a subdivision of \mathcal{T} . With this definition, subdivisions can be described in a fully combinatorial fashion. Note that, with some work, the statement “ \mathcal{T}' is a subdivision of \mathcal{T} ” can be decided algorithmically.

The following two facts are crucial.

- Let \mathcal{T} be a triangulation of a 3-manifold M . Then every submanifold $N \subseteq M$ (of dimension 0, 1, 2, or 3) is isotopic to a simplicial submanifold of a subdivision of \mathcal{T} .
- Let \mathcal{T}_M and \mathcal{T}_N be triangulations of the 3-manifolds M and N respectively. Then every homeomorphism $f: M \rightarrow N$ is isotopic to a simplicial isomorphism between a subdivision of \mathcal{T}_M and a subdivision of \mathcal{T}_N .

Therefore, in every step of our algorithms, we will always assume that submanifolds of 3-manifolds are simplicial. More precisely, a submanifold of a 3-manifold M with a triangulation \mathcal{T} will be represented as a subdivision \mathcal{T}' of \mathcal{T} , together with a set of simplices of \mathcal{T}' (of the appropriate dimension). Similarly, a homeomorphism between 3-manifolds M and N with triangulations \mathcal{T}_M and \mathcal{T}_N respectively will be represented as subdivisions \mathcal{T}'_M and \mathcal{T}'_N of \mathcal{T}_M and \mathcal{T}_N respectively, together with a simplicial isomorphism between \mathcal{T}'_M and \mathcal{T}'_N .

✿ **Algorithmic operations.** A useful property of subdivisions is the following: every two subdivisions \mathcal{T}_1 and \mathcal{T}_2 of the same triangulation \mathcal{T} have a common subdivision \mathcal{T}_{12} , which can be computed algorithmically. Without going into too much detail, let us simply remark that this implies that we can – and will – always assume that different “objects” in the same 3-manifold are simplicial with respect to the same subdivision.

All the natural operations one may want to perform on 3-manifolds can be done algorithmically; here is a long – but by no means complete – list of examples:

- composing and inverting homeomorphisms;
- finding images of submanifolds under homeomorphisms;
- cutting along a properly embedded codimension-one submanifold;
- finding (simplicial) regular neighbourhoods of properly embedded submanifolds;
- isotoping properly embedded submanifolds so that they are in general position;
- finding intersections of properly embedded submanifolds in general position.

Some of these constructions are trivial, while some require a lot of work. For more details on triangulations and subdivisions of PL manifolds, we refer the reader to the excellent introductory book [22] by Rourke and Sanderson.

✿ **Infinite search template.** The infinite search template algorithm can be informally described as follows: if the elements of a set \mathcal{S} can be enumerated algorithmically and \mathcal{S} is guaranteed to be non-empty, then there is an algorithm to construct an element of \mathcal{S} . As obvious as it may sound, this template will often allow us to blur the distinction between “proving that something exists” and “being able to algorithmically construct it”.

As an example, consider two 3-manifolds M and N , with triangulations \mathcal{T}_M and \mathcal{T}_N respectively. We can algorithmically enumerate all triples $(\mathcal{T}'_M, \mathcal{T}'_N, f)$ where \mathcal{T}'_M is a subdivision of \mathcal{T}_M , \mathcal{T}'_N is a subdivision of \mathcal{T}_N , and f is a simplicial isomorphism between \mathcal{T}'_M and \mathcal{T}'_N . If we know that M and N are homeomorphic, then the algorithm enumerating all the triples will eventually find one, effectively constructing a homeomorphism from M to N . As a consequence, after proving that two 3-manifolds are homeomorphic, we will always assume to have a homeomorphism available for our algorithmic purposes.

Let us remark that what we have described above is *not* an algorithm to solve the homeomorphism problem – that is, to decide whether M and N are homeomorphic. In fact, the algorithm will terminate if M and N are guaranteed to be homeomorphic, but will run forever if they are not; at any given point in time, there is no way to know if the algorithm has not found a homeomorphism because we have not waited long enough, or because there is none.

✿ **Solved algorithmic problems.** In the realm of surfaces, all the elementary questions can be settled algorithmically.

Theorem 1.1. *Given a surface F , the following operations can be carried out algorithmically.*

- 1 *Deciding whether F is orientable or not.*
- 2 *Computing the unique integers $g \geq 0$ and $k \geq 0$ such that F is homeomorphic to $\Sigma_{g,k}$ if it is orientable, and to $U_{g,k}$ if it is not.*
- 3 *Given another surface F' , deciding if F and F' are homeomorphic.*
- 4 *Given a curve in F , deciding if it is trivial (that is, it bounds a disc in F) and if it is boundary parallel.*
- 5 *Given two (multi)curves in F , deciding if they are isotopic or not.*
- 6 *Given another surface F' and two homeomorphisms $F \rightarrow F'$, deciding whether they are isotopic or not.*

Of course, the situation for 3-manifolds is substantially more complicated. The homeomorphism problem has only been recently settled, thanks to the contributions of many authors, the last piece needed being the geometrisation theorem (see [2, 16, 23]). The isotopy problem for codimension-1 submanifolds has also proved to be quite difficult to investigate. The sole purpose of this article, after all, is to provide an algorithm solving the isotopy problem for one specific surface – namely, Σ_2 – embedded in one specific 3-manifold – namely, S^3 . There exist, however, classical algorithms to answer some basic questions about 3-manifolds.

Theorem 1.2. *Given a 3-manifold M , the following operations can be carried out algorithmically.*

- 1 *Deciding whether M is irreducible.*
- 2 *If M is irreducible, deciding whether M is boundary irreducible.*
- 3 *If M is irreducible, deciding whether ∂M admits a non-separating compression disc in M .*
- 4 *Deciding whether M is a handlebody (possibly a 3-ball); if it is, computing its genus.*
- 5 *Given a surface properly embedded in M , deciding if it is boundary parallel.*
- 6 *Given an orientable surface which is either properly embedded in M or embedded in ∂M , deciding whether it is incompressible.*
- 7 *If M is irreducible and boundary irreducible, given an orientable incompressible surface properly embedded in M , deciding whether it is boundary incompressible.*
- 8 *If M is irreducible, given two connected incompressible surfaces properly embedded in M , deciding whether they are isotopic.*

Proof. Statements 1, 2, 6, and 7 are proved by Matveev in Theorems 4.1.12, 4.1.13, 4.1.15, and 4.1.19 of [20] respectively. Statement 4 is solved by [11, Algorithm 9.3]. Concerning statement 5, recall that a surface $F \subseteq M$ is boundary parallel if and only if it cobounds a product $F \times I$ with some surface in ∂M , where $F \subseteq M$ is identified with $F \times \{0\}$. It is then enough to check whether F is separating in M and, if it is, whether there is a component N of $M \setminus F$ such that $(\text{clos}(N), \partial F)$ is homeomorphic to $(F \times I, \partial F \times \{0\})$; we can decide this thanks to Theorem 1.3 below.

In order to address statement **8**, let us loosely rephrase a result of Waldhausen, namely [27, Proposition 5.4]: if F_1 and F_2 are isotopic incompressible surfaces properly embedded in M intersecting transversely with $\partial F_1 = \partial F_2$, then a component of $F_1 \setminus F_2$ is isotopic to a component of $F_2 \setminus F_1$. First of all, we check if ∂F_1 and ∂F_2 are isotopic in ∂M . If they are not, then F_1 and F_2 are not isotopic. Otherwise, we can assume that $\partial F_1 = \partial F_2$. Then we iterate over all components of $M \setminus (F_1 \cup F_2)$ and check whether any of these is a product $G \times I$ with $G \times \{0\} \subseteq F_1$ and $G \times \{1\} \subseteq F_2$, once again by means of Theorem 1.3. If we find such a region, then either $G \times \{0\} = F_1$ and $G \times \{1\} = F_2$ – thus showing that F_1 and F_2 are isotopic – or we use the product $G \times I$ to isotope $G \times \{1\} \subseteq F_2$ across $G \times \{0\} \subseteq F_1$, so as to reduce the number of components of $F_1 \cap F_2$. Hence, after finitely many steps, we either prove that F_1 and F_2 are isotopic, or we cannot find any more product regions; in this case, the two surfaces are not isotopic.

Finally, we prove statement **3**. The crucial claim is the following: let $D \subseteq M$ be a separating compression disc for ∂M ; then M admits a non-separating compression disc for its boundary if and only if the same holds for one component of $M \setminus D$. In fact, let $E \subseteq M$ be a non-separating compression disc for ∂M . Since M is irreducible, we can arrange for D and E to intersect transversely in a collection of arcs. Denote by E_1, \dots, E_k the closures of the components of $E \setminus D$; they are all discs, and we can think of them as being properly embedded in $M \setminus D$. Restriction induces an isomorphism $H^1(M) \cong H^1(M \setminus D)$ on cohomology groups, sending $[E]$ – the cohomology class represented by E – to $[E_1] + \dots + [E_k]$. Since E is non-separating in M , it is non-trivial in cohomology. Therefore, the same must hold for one of the discs E_1, \dots, E_k , thus showing that at least one component of $M \setminus D$ admits a non-separating compression disc for the boundary. The algorithm then proceeds as follows. If M is boundary irreducible, then clearly the answer is no. Otherwise, let D be a compression disc for the boundary; if it is non-separating we are done. Otherwise, we run the algorithm on the components of $M \setminus D$, and return yes if and only if we find a non-separating disc in at least one of them. The algorithm will eventually terminate, since ∂M can only be compressed finitely many times. \square

Concerning statement **8**, let us remark that surfaces in S^3 are quite the opposite of incompressible. This is what makes classifying them a challenging task, even for surfaces of genus two. We refer the reader to Section 1.4 for a survey of the many different topological phenomena which can arise when embedding a genus-two surface in S^3 , despite the trivial topology of the ambient space.

✿ Homeomorphism problem for manifolds with boundary pattern. Most importantly, the homeomorphism problem has been solved for what Matveev calls “Haken 3-manifolds with boundary pattern” in [20]. It is impossible to overstate how crucial this result is for our classification algorithm. In fact, the reader will probably find us referring to the following theorem more frequently than any other, often implicitly.

Theorem 1.3. *Let M and M' be irreducible 3-manifolds with non-empty boundary, and let $p \subseteq \partial M$ and $p' \subseteq \partial M'$ be (possibly empty) unions of curves such that $\partial M \setminus p$ and $\partial M' \setminus p'$ are incompressible in M and M' respectively. There is an algorithm which, given as input (M, p) and (M', p') , decides whether they are homeomorphic.*

Theorem 1.3 is a restricted version of [20, Theorem 6.1.6] which deals with the more general settings in which p and p' are allowed to be arbitrary graphs (what Matveev calls “boundary patterns”). There is, however, a technical point that needs to be addressed. Matveev’s theorem only guarantees an algorithmic solution to the question “is there

a possibly orientation-reversing homeomorphism $(M, p) \rightarrow (M', p')$?" One could, in theory, go through the proofs in [20, Chapter 6] and convince themselves that only trivial modifications are needed to answer the oriented version of the homeomorphism question. There is, however, a trick which allows us to derive Theorem 1.3 directly from [20, Theorem 6.1.6]. Simply augment the boundary pattern p by adding the same small chiral graph to each component of $M \setminus p$, obtaining the boundary pattern $\Gamma \subseteq \partial M$; carry out the same procedure for p' to construct the boundary pattern $\Gamma' \subseteq M'$. Then (M, p) and (M', p') are orientation-preservingly homeomorphic if and only if (M, Γ) and (M, Γ') are possibly orientation-reversingly homeomorphic; a graphical explanation is provided in Figure 1.1.




Figure 1.1. The homeomorphism problem can be reduced to the unoriented homeomorphism problem by adding chiral graphs on the boundary.

1.4 Examples

Classification of genus-two surfaces in S^3 is significantly harder than the same task for tori. In fact, increasing the genus by one is enough for a set of “wild” topological phenomena to appear; this is in contrast with the relative “tameness” of tori embedded in S^3 .

- ❁ First of all, it is not always true that a genus-two surface $S \subseteq S^3$ bounds a handlebody. A class of examples, as seen in Figure 1.2, can be constructed by starting with a knotted torus T and adding a tube inside the solid torus component of $S^3 \setminus T$. Alternatively, one can start with two knotted tori and join them with a tube; see Figure 1.3 for an example.
- ❁ Even if one component of $S^3 \setminus S$ is a handlebody, the other is not necessarily boundary irreducible; an example is described in Figure 1.4, where two knotted tori are connected by a “trivial” tube.
- ❁ Moreover, even in the case where one component of $S^3 \setminus S$ is a handlebody and the other is boundary irreducible (like in Figure 1.5), it is not always easy to identify a canonical meridian curve on S ; we invite the reader to compare this with the existence

of a standard meridian in tori embedded in S^3 , which makes the classification problem significantly easier.

- ✿ Finally, for the genus-one case, one can completely bypass the matter of meridian curves by exploiting the fact that tori in S^3 are isotopic if and only if their complementary regions are homeomorphic. This result, which follows from the work of Gordon and Luecke (see [7]), does not hold for genus-two surfaces, even when they bound a handlebody on one side; for an example of this phenomenon is given in [17].




Figure 1.2. A genus-two surface in S^3 which does not bound a handlebody on either side.

2 Homeomorphisms of 3-manifolds

2.1 Rationale

Section 2 is dedicated to the study of the mapping class groups of 3-manifolds from an algorithmic point of view. These groups, of course, are often infinite, but they can be described with finite amount of data. To be more specific, it has been known for a long time (see [12, Corollary 27.6]) that the mapping class group of an irreducible sufficiently large 3-manifold contains a finite-index subgroup H which is generated by Dehn twists about annuli and tori. By going through the proofs in Johansson's book, one realises that, in fact, only finitely many Dehn twists are required to generate H . Therefore, theoretically, the mapping class group of a 3-manifold can be fully described by providing finitely many Dehn twists about annuli and tori, together with representatives of the cosets of H in the mapping class group.

There is, however, a substantial gap between acknowledging the existence of H and being able to algorithmically find generators of H and representatives of its cosets. The




Figure 1.3. Another genus-two surface in S^3 which does not bound a handlebody on either side.




Figure 1.4. A genus-two surface splitting S^3 into two components: one is a handlebody, but the other is not boundary irreducible.

proofs by Johannson are for the most part constructive, and one could carefully convert them into algorithms for finding the generating Dehn twists (whereas representatives of the cosets of H are somewhat trickier to extract from said proofs). This is exactly the path we will follow in Section 2, albeit with a few *caveats*.

- ⚙️ Instead of working in the full generality of what Johannson calls “3-manifolds with complete and useful boundary pattern”, we restrict our attention to *irreducible 3-manifold pairs*, following the work of Jaco and Shalen in [10]. We admit that our approach is perhaps less elegant and powerful than Johannson’s, but the decrease in flexibility is compensated by the (relative) conciseness of some of the algorithmic procedures we will describe.
- ⚙️ We make use of the geometrisation theorem (see [21, 24]) and Kuperberg’s excellent exposition (see [16]) in order to compute the (finite) mapping class group of simple 3-manifolds. This could probably be avoided, following Johannson’s inductive argument which relies on hierarchies. We decided to go for the shorter solution, rather than the one which was more faithful to Johannson’s original work.
- ⚙️ As will become clear in Section 4, we are only interested in the trace of the mapping




Figure 1.5. A genus-two surface splitting S^3 into two components: one is a handlebody, and the other is boundary irreducible.

class group, and care little about what happens in the interior of the 3-manifold. This is why Theorem 2.16 is stated the way it is. It is true that little additional effort would have been required to give a complete description of the mapping class group, but once again we decided to avoid exceeding in generality, in order to keep this section reasonably short.

In Section 2.2, we start by recalling the definition of JSJ system for an irreducible sufficiently large 3-manifold pair, following [10] as closely as possible. Sections 2.3, 2.4, and 2.5 address the computation of the mapping class groups and the homeomorphism problem for Seifert fibred-spaces, I -bundles, and simple manifolds respectively. In Section 2.6, we show how to actually compute the JSJ decomposition of a 3-manifold pair (with some additional assumptions). Finally, in Section 2.7, we put the results of the previous section together to deliver, as anticipated, a description of the mapping class group in the form of Theorem 2.16.

2.2 JSJ decomposition

This section is little more than a restatement of the definitions and the main theorem of [10, Chapter V, §6].

Definition 2.1. For $n \geq 2$, a *3-manifold n -tuple* is an n -tuple (M, R_1, \dots, R_{n-1}) where M is a 3-manifold, and R_1, \dots, R_{n-1} are surfaces in ∂M such that $R_i \cap R_j$ is a collection of curves for $1 \leq i < j \leq n-1$. We say *pair* and *triple* instead of 2-tuple and 3-tuple respectively.

For a 3-manifold pair (M, R) , we give the following definitions.

- We say that (M, R) is *irreducible* if M is irreducible and R is incompressible in M .
- We say that (M, R) is *sufficiently large* if M is.
- A surface $F \subseteq M$ is *properly embedded* in (M, R) if it is properly embedded in M and ∂F lies in R .
- Two disjoint surfaces F and F' in M with ∂F and $\partial F'$ lying in R are *parallel* in (M, R) if there is a 3-manifold $W \subseteq M$ homeomorphic to $F \times I$ such that $\partial_h W = F \cup F'$ and $\partial W \setminus \partial_h W \subseteq R$.

- A surface F in M is *boundary parallel* in (M, R) if it is parallel in (M, R) to a surface in R .

Our aim is to develop an algorithmic theory for the JSJ decomposition of irreducible sufficiently large 3-manifold pairs. Loosely speaking, as is the case for ordinary 3-manifolds, the decomposition for a pair (M, R) comes from an incompressible surface F properly embedded in (M, R) such that cutting M along it yields pieces which are in some sense “simpler” than the original pair. Note that these pieces immediately inherit a 3-manifold pair structure by considering the remnants of R as a surface in the boundary of $M \setminus F$. By doing so, however, we lose information about what parts of ∂M come from F . This is why, whenever we have a 3-manifold pair (M, R) with a properly embedded surface F , we naturally think of $M \setminus F$ as a 3-manifold triple. More precisely, we define $(M, R) \setminus F$ to be the 3-manifold triple (M', R', F') , where:

- $M' = M \setminus F = M \setminus \mathring{\mathcal{N}}(F)$;
- $R' = R \cap M'$;
- $F' = \mathcal{N}(F) \cap M'$.

We now proceed to define the types of pieces we allow in the JSJ decomposition of a 3-manifold pair.

Definition 2.2. A 3-manifold n -tuple (M, \mathbf{R}) is a *Seifert n -tuple* if $\partial M = R_1 \cup \dots \cup R_{n-1}$ and M admits a Seifert fibration such that R_1, \dots, R_{n-1} are unions of fibres. Such a fibration is called a *Seifert fibration* for (M, \mathbf{R}) . A 3-manifold n -tuple equipped with a Seifert fibration is called a *Seifert-fibred n -tuple*.

Definition 2.3. A 3-manifold triple is an *I -bundle triple* if it is homeomorphic to $(X, \partial_h X, \partial_v X)$ for some I -bundle X over a surface.

Definition 2.4. A 3-manifold triple (M, R, F) is a *simple triple* if the following hold:

- every incompressible torus properly embedded in M is parallel in M to a component of R or of F ;
- every incompressible annulus A properly embedded in M with $\partial A \subseteq \text{int}(R)$ is parallel in (M, R) to an annulus in R or in F .

We can finally define the JSJ system for a 3-manifold pair. As expected, the JSJ system is unique up to isotopy.

Definition 2.5. Let (M, R) be an irreducible sufficiently large 3-manifold pair. A *JSJ system* of (M, R) is a surface F properly embedded in (M, R) , possibly empty, satisfying the following properties.

- i** Each component of F is an incompressible torus or annulus, and it is not boundary parallel in (M, R) .
- ii** Each component of $(M, R) \setminus F$ is a Seifert triple, an I -bundle triple, or a simple triple.
- iii** The surface F is minimal with respect to inclusion among all surfaces properly embedded in (M, R) satisfying properties **i** and **ii**.

Theorem 2.1 ([10, Chapter V, §6, Generalised Splitting Theorem]). *Let (M, R) be an irreducible sufficiently large 3-manifold pair. Then (M, R) admits a JSJ system, which is unique up to isotopy in M fixing $\partial M \setminus \text{int}(R)$.*

We now embark on an in-depth study of the pieces of the JSJ decomposition, with the aim of achieving a better understanding of the original 3-manifold pair. Specifically, for a piece (N, R', F') of $(M, R) \setminus F$, we will be interested in the following algorithmic problems:

- solving the homeomorphism problem for (N, R', F') ;
- deciding whether a homeomorphism $F' \rightarrow F'$ extends to a self-homeomorphism of (N, R', F') ;
- describing the group $\mathbf{H}_{F'}(N, R')$ in a computationally feasible way (that is, with a finite amount of data).

2.3 Seifert pieces

We start with an analysis of Seifert pieces. In fact, we will work with general Seifert n -tuples instead of just triples. The reason will become apparent in Section 2.7, when we will need to deal with Seifert 4-tuples as well. We follow the convention that a *Seifert manifold* is a 3-manifold which admits a Seifert fibration, whereas a *Seifert-fibred manifold* is a 3-manifold endowed with a fixed Seifert fibration.

Now is a good time to briefly discuss a “computationally friendly” definition of Seifert fibration. For our algorithmic purposes, a Seifert fibration of a triangulated 3-manifold is given by a 2-subcomplex Z of (some subdivision of) M satisfying the following properties.

- i Every point of Z has a neighbourhood which looks like one of the local models in Figure 2.1; the set of points whose local models are of type (b) or (c) is a union of circles called *fibres*.
- ii The closure of each component of $M \setminus Z$ is a solid torus, and has at least one fibre on its boundary.
- iii No fibre bounds a disc in $M \setminus Z$.

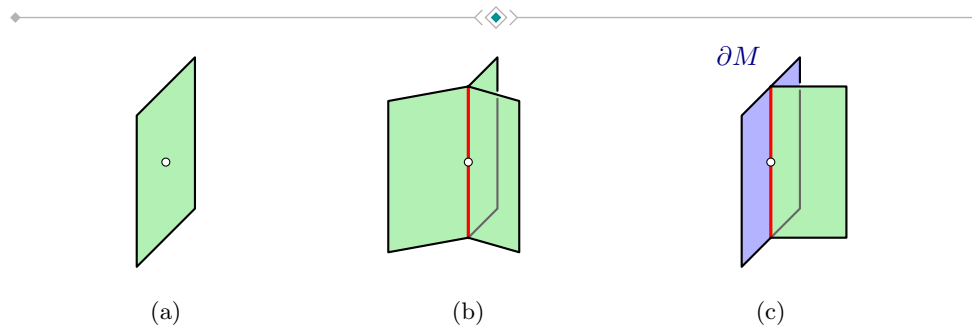


Figure 2.1. Local models for a 2-complex defining a Seifert fibration for a 3-manifold M . The 2-complex is coloured green, while the fibres are highlighted in red.

On one hand, if M is equipped with a Seifert fibration in the usual sense and $p: M \rightarrow B$ is the projection to the base surface, we can obtain a 2-complex $Z \subseteq M$ by taking the preimage under p of a suitable 1-complex cutting B into discs which

contain at most one singular point. Conversely, given a 2-complex Z as above, we can recover the Seifert fibration by fibring each complementary solid torus compatibly with the fibres of Z . Note that properties **i**, **ii**, and **iii** are easy to check algorithmically. As a consequence, given a Seifert manifold, we can find a Seifert fibration for it. This discussion naturally extends to a Seifert n -tuple (M, \mathbf{R}) ; in this case, we additionally require that ∂R_i is a union of fibres of Z for $1 \leq i \leq n-1$.

With the following we provide a version of [9, Lemma VI.19] which better suits our needs.

Lemma 2.2. *Let M and M' be Seifert-fibred manifolds. Let $f: M \rightarrow M'$ be a homeomorphism, and suppose that there is a non-empty finite union of fibres, fibred annuli, and fibred tori $X \subseteq \partial M$ such that the restriction of f to X is fibre-preserving. Then f can be isotoped, fixing X pointwise, to a fibre-preserving homeomorphism.*

Proof. The proof is essentially identical to that of Jaco, and we present it here for the sake of completeness. Let $p': M' \rightarrow B'$ be the projection to the base surface of the Seifert fibration, and let M' have n exceptional fibres. The proof proceeds by induction on $-\chi(B') + n$.

We can easily isotope f , fixing X , to be fibre-preserving on each component of ∂M which intersects X . Therefore, if $Y \subseteq \partial M$ is the union of all such components, then we can assume that $f|_Y: Y \rightarrow \partial M'$ is fibre-preserving, and prove that f can be isotoped, fixing Y , to a fibre-preserving homeomorphism.

If B' is a disc and $n \leq 1$, then M and M' are fibred solid tori of the same type, and it is easy to construct the desired isotopy. If instead B' is not a disc or $n \geq 2$, then let $y' \in \partial B'$ the image under p' of one of the fibres in $f(Y)$. There exists an arc a' properly embedded in B' connecting y' to a different point in $\partial B'$ and avoiding exceptional points, such that a' is not trivial in $\pi_1(B'_\bullet, \partial B')$, where B'_\bullet is the punctured surface obtained by removing the exceptional points from B' . Then $A' = (p')^{-1}(a')$ is an incompressible boundary incompressible vertical annulus properly embedded in M' . The annulus $A = f^{-1}(A')$ is incompressible and boundary incompressible in M . It need not be vertical, but at least one of the components of ∂A is a fibre, since it lies in Y . We can isotope A to be vertical in M , and by [12, Proposition 5.6]² we can assume that the isotopy fixes $A \cap Y$ pointwise. In fact, it is not hard to see that we can even take the isotopy to fix all of Y . As a consequence, we find that f restricts to a homeomorphism $M \setminus A \rightarrow M' \setminus A'$ which is fibre-preserving on $Y \cap (M \setminus A)$ as well as on the two annuli $\text{clos}(\partial(M \setminus A) \setminus \partial M)$. The conclusion now follows by induction. \square

The next two propositions generalise well known facts about Seifert spaces – namely uniqueness of the Seifert fibration and solution to the homeomorphism problem – to the setting of Seifert n -tuples.

Proposition 2.3. *Let (M, \mathbf{R}) be a Seifert $(n+1)$ -tuple with non-empty boundary. Then (M, \mathbf{R}) admits exactly one Seifert fibration up to isotopy in (M, \mathbf{R}) , unless the surfaces R_1, \dots, R_n are unions of tori and one of the following holds.*

- 1** *The 3-manifold M is a solid torus; in this case, every fibration of ∂M which does not contain a meridian of M extends to a unique Seifert fibration of M with at most one exceptional fibre; the fibration of ∂M containing a meridian of M does not extend to a Seifert fibration of M .*

²It is immediate to see that exceptions 5.1.1–5.1.5 listed by Johannson do not occur, since M has non-empty boundary and the case of a solid torus with a single exceptional fibre has already been addressed.

- ② The 3-manifold M is homeomorphic to $T^2 \times I$; in this case, every fibration of $T^2 \times \{0\}$ extends to a unique fibration of M .
- ③ The 3-manifold M is homeomorphic to $U_2 \times I$; in this case, M admits two Seifert fibrations, one over a Möbius band and one over a disc.

Proof. First of all, let us remark that, as soon as one of the surfaces R_1, \dots, R_n has an annulus component, then the Seifert fibration for (M, \mathbf{R}) is necessarily unique: given two such fibrations, the homeomorphism $\text{id}: M \rightarrow M$ is fibre-preserving on $X = \partial R_1 \cup \dots \cup \partial R_n$; by Lemma 2.2, we can isotope one fibration to the other in (M, R_1, \dots, R_n) .

When instead all the components of R_1, \dots, R_n are tori, it is clear that two Seifert fibrations for (M, \mathbf{R}) are isotopic in (M, \mathbf{R}) if and only if they are isotopic in M . The conclusion then follows from [9, Theorem VI.18]. \square

Remark 2.4. The two Seifert fibrations of $U_2 \times I$ are *inequivalent*, by which we mean that no self-homeomorphism of $U_2 \times I$ can send one to the other. In other words, if we fix a Seifert fibration for $U_2 \times I$, then every self-homeomorphism of this 3-manifold can be isotoped to be fibre-preserving.

Proposition 2.5. *There is an algorithm which, given as input two Seifert $(n+1)$ -tuples (M, \mathbf{R}) and (M', \mathbf{R}') , decides whether they are homeomorphic or not.*

Proof. Fix two Seifert fibrations for (M, \mathbf{R}) and (M', \mathbf{R}') . It is very well known (for a thorough and complete discussion, see [19, Theorem 10.4.19]) that computing the Seifert invariants for the fibrations of M and M' leads to a straightforward algorithm for deciding whether M and M' are homeomorphic. If they are not, then clearly neither are (M, \mathbf{R}) and (M', \mathbf{R}') . If M and M' are closed and homeomorphic, then we are done. On the other hand, if (M', \mathbf{R}') is one of the exceptions described in ① and ② of Proposition 2.3, then answering the homeomorphism problem is easy.

Therefore, let us assume that (M', \mathbf{R}') has non-empty boundary and is not one of these exceptions, and let $f: M \rightarrow M'$ be a homeomorphism. We can find a finite set \mathcal{F}_0 of fibre-preserving self-homeomorphisms of M' such that every permutation of the boundary components of M' is induced by an element of \mathcal{F}_0 . Let $\iota: M' \rightarrow M'$ be a fibre-preserving homeomorphism such that $\iota|_{\partial}$ acts as $-\text{id}$ on $H_1(\partial M')$, and define

$$\mathcal{F} = \bigcup_{g \in \mathcal{F}_0} \{g, \iota g\}.$$

We claim that (M, \mathbf{R}) and (M', \mathbf{R}') are homeomorphic if and only if $g^{-1}f: M \rightarrow M'$ can be isotoped to a homeomorphism $(M, \mathbf{R}) \rightarrow (M', \mathbf{R}')$ for some $g \in \mathcal{F}$ – and note that this condition is easy to check algorithmically. The reverse implication is trivial. Conversely, assume that there is a homeomorphism

$$f': (M, \mathbf{R}) \longrightarrow (M', \mathbf{R}').$$

We can isotope f' in M so that $f'f'^{-1}: M' \rightarrow M'$ is fibre-preserving. There is a homeomorphism $g \in \mathcal{F}$ such that $f'f'^{-1}g$ preserves each boundary component of M' , as well as the orientation of the fibres on $\partial M'$. But then the homeomorphism $f'f'^{-1}g$ acts as a power of the Dehn twist about a fibre on each boundary component of M' , hence it can be isotoped to a self-homeomorphism of (M', \mathbf{R}') . We immediately conclude that $g^{-1}f'$ can be isotoped to a homeomorphism $(M, \mathbf{R}) \rightarrow (M', \mathbf{R}')$. \square

We are now ready to give an algorithmic description of the mapping class group of (most) Seifert n -tuples with boundary.

Proposition 2.6. *There is an algorithm which, given as input a Seifert-fibred $(n+2)$ -tuple (M, \mathbf{R}, F) with non-empty boundary and not homeomorphic to one of the exceptions described in [1](#) and [2](#) of Proposition 2.3, returns*

- *a finite collection \mathcal{F} of self-homeomorphisms of $(\partial M, \mathbf{R})$ fixing F pointwise, and*
- *a finite collection $\mathcal{C} = \{(a_1, b_1), \dots, (a_m, b_m)\}$, where $a_1, \dots, a_m, b_1, \dots, b_m$ are pairwise disjoint fibres in $\text{int}(R_1) \cup \dots \cup \text{int}(R_n)$,*

such that

$$\mathbf{H}_F(M, \mathbf{R})|_{\partial} = \bigcup_{f \in \mathcal{F}} \langle \tau_{a_1} \tau_{b_1}^{-1}, \dots, \tau_{a_m} \tau_{b_m}^{-1} \rangle f$$

as a subgroup of $\mathbf{H}_F(\partial M, \mathbf{R})$.

Proof. Let $p: (M, \mathbf{R}, F) \rightarrow (B, \overline{\mathbf{R}}, \overline{F})$ be the projection to the base surface of the Seifert fibration, and consider a homeomorphism $f \in \mathbf{H}_F(M, \mathbf{R})$. We can isotope f to be fibre-preserving fixing F pointwise, using Proposition 2.3 and Remark 2.4 if F is empty and Lemma 2.2 otherwise. Denote by \overline{f} the induced self-homeomorphism of $(B, \overline{\mathbf{R}}, \overline{F})$, which is not necessarily orientation-preserving.

• **Identity on B .** As we will see, the role of \mathcal{F} is to encode the action of $\mathbf{H}_F(M, \mathbf{R})$ on the space $(\partial B, \overline{\mathbf{R}}, \overline{F})$. First of all, note that we can easily list all permutations of the components of ∂B which are induced by elements of $\mathbf{H}_F(M, \mathbf{R})$. Therefore, we can compute a set of representatives $\mathcal{F}_0 \subseteq \mathbf{H}_F(M, \mathbf{R})$ such that, up to replacing f with $f f_0$ for some $f_0 \in \mathcal{F}_0$, we can assume that f sends each component of ∂M to itself.

The homeomorphism f induces a self-homeomorphism of $(\partial B, \overline{\mathbf{R}}, \overline{F})$ which fixes \overline{F} pointwise and sends each component of ∂B to itself. If F is non-empty, the homeomorphism \overline{f} acts orientation-preservingly on each component of ∂B ; otherwise, it will consistently preserve or reverse the orientations of these components. Either way, there are only finitely many such self-homeomorphisms of $(\partial B, \overline{\mathbf{R}}, \overline{F})$, and they are all induced by elements of $\mathbf{H}_F(M, \mathbf{R})$. We can then compute a set $\mathcal{F}_1 \subseteq \mathbf{H}_F(M, \mathbf{R})$ such that, up to replacing f with $f f_1$ for some $f_1 \in \mathcal{F}_1$, we can assume that \overline{f} is the identity on ∂B .

In fact, up to further composing f with a self-homeomorphism of M fixing ∂M pointwise, we can assume that \overline{f} is the identity on all of B ; this follows from the fact that if a self-homeomorphism of B fixing ∂B lifts to a self-homeomorphism of M , then it lifts to a self-homeomorphism of M fixing ∂M .

• **Identity on almost all of ∂M .** Let \mathcal{X} be the set of components of $\partial M \setminus \text{int}(F)$. For each pair X, Y of different components in \mathcal{X} , let $A_{X,Y}$ be a vertical annulus connecting X and Y . We can of course assume that the boundary curves of all these annuli lie in $\text{int}(R_1) \cup \dots \cup \text{int}(R_n)$ and are pairwise disjoint. Moreover, for each component $X \in \mathcal{X}$, we can define k_X to be the unique integer such that $f|_X$ is isotopic to the composition of k_X Dehn twists of X about a fibre, through an isotopy which fixes ∂X pointwise. Up to composing with powers of Dehn twists about the annuli $A_{X,Y}$ – which induce the identity on B – we can assume that $k_X = 0$ for all components $X \in \mathcal{X}$ but one. In other words, we may assume that the restriction of f to $\partial M \setminus \text{int}(X)$ is the identity for some component $X \in \mathcal{X}$.

• **Identity on ∂M .** In fact, we now prove that f restricts to the identity on all of ∂M . Let N_1, \dots, N_m be fibred regular neighbourhoods of the exceptional fibres

of M . Without loss of generality, we may assume that f restricts to the identity on $N_1 \cup \dots \cup N_m$. Define $M' = M \setminus \text{int}(N_1 \cup \dots \cup N_m)$, which is an S^1 -bundle over the surface $B' = B \setminus \text{int}(p(N_1 \cup \dots \cup N_m))$; since B' has non-empty boundary, the bundle is uniquely determined by B' . The homeomorphism f restricts to $f': M' \rightarrow M'$ such that f' is the identity on $\partial M' \setminus \text{int}(X)$.

- Let us first deal with the case where B is orientable. Let S be a horizontal copy of B' in M' . At the level of homology, we have that

$$0 = f'_*[\partial S] = [\partial S] + k_X \cdot [c] = k_X \cdot [c] \in H_1(M'),$$

where $[c]$ is the homology class of a fibre. We conclude that $k_X = 0$ and that f' does indeed restrict to the identity on ∂M .

- Assume now that B – and hence B' – is non-orientable. Let \widetilde{M}' be the double covering of M' which is a product S^1 -bundle over the orientable double covering \widetilde{B}' of B' . Denote by $q: \widetilde{M}' \rightarrow M'$ the covering map, and by $\iota: \widetilde{M}' \rightarrow \widetilde{M}'$ the non-trivial deck transformation of q . The homeomorphism f' lifts to $\widetilde{f}': \widetilde{M}' \rightarrow \widetilde{M}'$, which is the identity on $\partial \widetilde{M}' \setminus q^{-1}(\text{int}(X))$; denote by \widetilde{X}_1 and \widetilde{X}_2 the two components of $q^{-1}(X)$, and by k_1 and k_2 the integers such that $\widetilde{f}'|_{\widetilde{X}_i}$ is the k_i -th power of the Dehn twist about a fibre for $i \in \{1, 2\}$. Let \widetilde{S} be a horizontal copy of \widetilde{B}' in \widetilde{M}' . Fixing an orientation of \widetilde{S} allows us to define canonically oriented horizontal curves a_1 and a_2 , which are the boundary components of \widetilde{S} intersecting \widetilde{X}_1 and \widetilde{X}_2 respectively. Moreover, let b_1 and b_2 be fibres of \widetilde{X}_1 and \widetilde{X}_2 , oriented so that they intersect a_1 and a_2 with positive sign. We can then define k_i as the integer such that $[\widetilde{f}'(a_i)] = [a_i] + k_i \cdot [b_i]$ in $H_1(\partial \widetilde{M}')$ for $i \in \{1, 2\}$. The homological argument described in the orientable case shows that $k_1 + k_2 = 0$. On the other hand, note that $[\iota(a_1)] = -[a_2]$ and $[\iota(b_1)] = -[b_2]$. Then the relation $\widetilde{f}'\iota = \iota\widetilde{f}'$ readily implies that $k_1 = k_2$, since

$$\begin{aligned} [\widetilde{f}'\iota(a_1)] &= -[\widetilde{f}'(a_2)] = -[a_2] - k_2 \cdot [b_2], \\ [\iota\widetilde{f}'(a_1)] &= [\iota(a_1)] + k_1 \cdot [\iota(b_2)] = -[a_2] - k_1 \cdot [b_2]. \end{aligned}$$

We conclude that \widetilde{f}' restricts to the identity on $\partial \widetilde{M}'$, and hence f' restricts to the identity on ∂M .

- **Conclusion.** We have decomposed f as a product

$$t \circ h \circ f_1 \circ f_0 \in \mathbf{H}_F(M, \mathbf{R}),$$

where $f_0 \in \mathcal{F}_0$, $f_1 \in \mathcal{F}_1$, $h \in \mathbf{H}_{\partial M}(M)$, and t is a product of powers of Dehn twists of M about the annuli $A_{X,Y}$ for $X, Y \in \mathcal{X}$. The algorithm will then return

$$\mathcal{F} = \{(f_0 f_1)|_{\partial} : f_0 \in \mathcal{F}_0, f_1 \in \mathcal{F}_1\} \quad \text{and} \quad \mathcal{C} = \{\partial A_{X,Y} : X, Y \in \mathcal{X} \text{ and } X \neq Y\}. \quad \square$$

As an easy consequence, we can also solve the extension problem for Seifert n -tuples.

Corollary 2.7. *Let (M, \mathbf{R}, F) and (M', \mathbf{R}', F') be Seifert-fibred $(n+2)$ -tuples not homeomorphic to one of the exceptions described in **1** and **2** of Proposition 2.3, and let $f: F \rightarrow F'$ be a fibre-preserving homeomorphism. There is an algorithm which, given as input the two $(n+2)$ -tuples and f , decides whether f extends to a (not necessarily fibre-preserving) homeomorphism*

$$(M, \mathbf{R}, F) \longrightarrow (M', \mathbf{R}', F').$$

Proof. It suffices to be able to solve the following two questions algorithmically.

- ❶ Are the triples (M, R, F) and (M', R', F') homeomorphic?
- ❷ Given a homeomorphism $g: F \rightarrow F$, does it extend to a self-homeomorphism of (M, R, F) ?

The first question can be easily answered, since we can compute the base surfaces of the I -bundles (M, R, F) and (M', R', F') and check if they are the same. As far as the second is concerned, we may assume that g is fibre-preserving. Moreover, since the group $\mathbf{H}(M, R, F)$ acts transitively on the components of F , we can further assume that g maps each component of F to itself. Now, the restriction of g to an annulus component A of F can either preserve or swap the boundary components of A . If g consistently preserves or swaps the boundary components of every component of F , then it extends to a self-homeomorphism of (M, R, F) ; otherwise, it does not. \square

Proposition 2.9. *There is an algorithm which, given as input an I -bundle triple (M, R, F) , returns*

- a finite collection \mathcal{F} of self-homeomorphisms of $(\partial M, R, F)$ fixing F pointwise, and
- a finite collection $\mathcal{C} = \{(a_1, b_1), \dots, (a_m, b_m)\}$, where $a_1, \dots, a_m, b_1, \dots, b_m$ are curves in $\text{int}(R)$ with $a_1 \cap b_1 = \dots = a_m \cap b_m = \emptyset$,

such that

$$\mathbf{H}_F(M, R)|_{\partial} = \bigcup_{f \in \mathcal{F}} \langle \tau_{a_1} \tau_{b_1}^{-1}, \dots, \tau_{a_m} \tau_{b_m}^{-1} \rangle f$$

as a subgroup of $\mathbf{H}_F(\partial M, R)$. Moreover, if (M, R, F) is the product I -bundle over $\Sigma_0, \Sigma_{0,1}, \Sigma_{0,2}$, or $\Sigma_{0,3}$, or the twisted I -bundle over $U_1, U_{1,1}$, or $U_{1,2}$, then the curves in \mathcal{C} are pairwise disjoint.

Proof. Let $p: (M, F) \rightarrow (B, \partial B)$ be the projection to the base surface of the I -bundle, and consider a homeomorphism $f \in \mathbf{H}_F(M, R)$. We can isotope f to be fibre-preserving fixing F pointwise. Denote by \bar{f} the induced self-homeomorphism of B , which is not necessarily orientation-preserving. We now distinguish two cases, depending on whether B is orientable or not.

❁ If B is orientable, then \bar{f} is necessarily orientation-preserving, unless $F = \emptyset$. If this is the case, let us choose a homeomorphism $\iota \in \mathbf{H}(M, R)$ such that the induced homeomorphism $\bar{\iota}: B \rightarrow B$ is orientation-reversing; otherwise, we set $\iota = \text{id}: M \rightarrow M$. Up to composing f with ι , we can assume that \bar{f} is orientation-preserving. There is a well-known explicit set of curves $c_1, \dots, c_m \subseteq B$ such that $\tau_{c_1}, \dots, \tau_{c_m}$ generate $\mathbf{H}_{\partial B}(B)$ – see, for instance, [5, Section 4.4.4]. As a consequence, the homeomorphism f belongs to the subgroup of $\mathbf{H}_F(M, R)$ generated by $\tau_{A_1}, \dots, \tau_{A_m}$, where $A_i = p^{-1}(c_i)$ for $1 \leq i \leq m$. The algorithm will then return

$$\mathcal{F} = \{\text{id}, \iota|_{\partial}\} \quad \text{and} \quad \mathcal{C} = \{\partial A_i : 1 \leq i \leq m\}.$$

Note that, if B is one of $\Sigma_0, \Sigma_{0,1}, \Sigma_{0,2}$, or $\Sigma_{0,3}$, then the curves c_1, \dots, c_m can be chosen to be disjoint.

❁ If B is non-orientable, denote by $\mathbf{H}_{\partial B}^{\pm}(B)$ the group of self-homeomorphisms of B fixing ∂B pointwise, modulo isotopies through homeomorphisms of the same kind. Chillingworth (in [3, Section 3]) and Kobayashi and Omori (in [14, Proposition 3.2]),

respectively for the cases $\partial B = \emptyset$ and $\partial B \neq \emptyset$, provide an explicit set of two-sided curves $c_1, \dots, c_m \subseteq B$ and a homeomorphism $y \in \mathbf{H}_{\partial B}^{\pm}(B)$ such that

$$\mathbf{H}_{\partial B}^{\pm}(B) = \langle \tau_{c_1}, \dots, \tau_{c_m}, y \rangle.$$

The homeomorphism y is a *Y-homeomorphism*, a description of which can be found in [18, Section 2]; in particular, we have that y^2 is a Dehn twist about a curve c_0 which can be computed explicitly. Easy algebraic manipulations show that

$$\mathbf{H}_{\partial B}^{\pm}(B) = \bigcup_{g \in \{\text{id}, y\}} \langle \tau_{c_1}, \dots, \tau_{c_m}, y\tau_{c_1}y^{-1}, \dots, y\tau_{c_m}y^{-1}, y^2 \rangle g.$$

Define $c_{m+i} = y(c_i)$ for $1 \leq i \leq m$, so that $y\tau_{c_i}y^{-1} = \tau_{c_{m+i}}$; let $h \in \mathbf{H}_F(M, R)$ be a homeomorphism such that $\bar{h} = y$. The algorithm will then return

$$\mathcal{F} = \{\text{id}, h|_{\partial}\} \quad \text{and} \quad \mathcal{C} = \{\partial p^{-1}(c_i) : 0 \leq i \leq 2m\}.$$

Note that, if B is one of $U_{1,0}$, $U_{1,1}$, or $U_{1,2}$, then the curves c_1, \dots, c_m can be chosen to be disjoint. \square

2.5 Simple pieces

We finally turn our attention to simple pieces. As anticipated, instead of using Johannson's original approach with hierarchies, we will prove the finiteness of the mapping class group of these pieces by means of hyperbolic geometry and the geometrisation theorem. While sacrificing generality, this strategy will allow us to bypass the algorithmically challenging induction procedure of Johannson.

In Sections 2.3 and 2.4 we dealt with Seifert n -tuples and I -bundle triples in full generality (save for a few exceptions). For simple triples, instead, we will restrict our attention to a subclass satisfying additional conditions. This will, of course, impose further restrictions on the type of 3-manifold pairs which our final algorithm will accept as input. However, we will maintain a level of generality which will be sufficient for our purposes. The additional constraints are listed in the following definition.

Definition 2.6. A 3-manifold triple (M, R, F) is *reasonable* if it satisfies the following properties:

- the pair (M, R) is irreducible;
- no component of R is a disc;
- the surface F is incompressible in M ;
- each component of F is a torus or an annulus, and is not boundary parallel in (M, R) ;
- each component of $\partial M \setminus \text{int}(R \cup F)$ is an annulus.

Part of the appeal of simple reasonable triples is that the two surfaces R and F cover the whole boundary of the 3-manifold M , save for a few explicit exceptions.

Lemma 2.10. *Let (M, R, F) be a reasonable 3-manifold triple. Then one of the following holds.*

- **i** *The triple (M, R, F) is not simple.*

- ii The surfaces R and F cover the boundary of M .
- iii The 3-manifold M is a solid torus, and $\partial M \setminus \text{int}(R \cup F)$ is a union of annuli which are incompressible in M .
- iv The 3-manifold M is homeomorphic to $T^2 \times I$ in such a way that $T^2 \times \{1\}$ is a component of R or F .

Proof. Suppose that (M, R, F) is simple, and that $\partial M \setminus \text{int}(R \cup F)$ is non-empty. Let A be a component of this surface; by assumption, A is an annulus. Let T be the boundary component of M containing A ; we show that T is a torus. Let A' be the component of $\partial M \setminus \text{int}(R)$ containing A . If A' is a torus, then $A' = T$ and we are done. Otherwise, the surface A' is an annulus. If A' were compressible in M , then one component of $\partial A' \subseteq \partial R$ would bound a disc in M ; since R is incompressible and has no disc components, this is impossible. Therefore, A' is incompressible. Since the triple (M, R, F) is simple, it follows that T is indeed a torus.

- If T is compressible, then M is in fact a solid torus, and the components of R, F , and $\partial M \setminus \text{int}(R \cup F)$ are parallel annuli in ∂M ; by assumption, these annuli are incompressible in M .
- If T is incompressible, then it must be parallel in (M, R) to a torus in R or in F . This immediately implies that M is homeomorphic to $T^2 \times I$, where $T^2 \times \{0\}$ is identified with T and $T^2 \times \{1\}$ is a component of either R or F . \square

As a final restriction, we only want to deal with simple triples which do not belong to the classes we have already analysed – namely, I -bundles and Seifert pieces.

Definition 2.7. A 3-manifold triple (M, R, F) is *strongly simple* if it is simple, it is not an I -bundle triple, and $(M, R, F, \partial M \setminus \text{int}(R \cup F))$ is not a Seifert 4-tuple.

Under the additional assumptions we have listed, we can now prove that the mapping class group of a simple triple is finite, and present an algorithm to compute it. At the same time, we will describe a solution to the homeomorphism problem for simple triples.

Proposition 2.11. *The following hold for reasonable strongly simple 3-manifold triples (M, R, F) and (M', R', F') .*

- 1 *There is an algorithm which, given as input the triples (M, R, F) and (M', R', F') , decides whether they are homeomorphic.*
- 2 *The group $\mathbf{H}(M, R, F)$ is finite. Moreover, there is an algorithm which, given as input the triple (M, R, F) , returns representatives of $\mathbf{H}(M, R, F)$.*

Proof. By Lemma 2.10, we immediately see that, if R and F do not cover the boundary of M , the 4-tuple $(M, R, F, \partial M \setminus \text{int}(R \cup F))$ admits a Seifert fibration. Therefore, we have that $\partial M = R \cup F$.

✿ **Boundary incompressibility of F .** We claim that F is boundary incompressible in (M, R) in the following sense: for every disc D in M whose boundary can be written as the union of two arcs a and b with $a \cap b = \partial a = \partial b$, $D \cap F = a$, and $D \cap R = b$, we have that a cuts a disc off of F .

In fact, suppose by contradiction that D is a disc as described above, with a an essential arc in an annulus component A of F . Let A' be the disc obtained by boundary compressing A along D , that is $A' = (A \setminus \partial_v \mathcal{N}(D)) \cup \partial_h \mathcal{N}(D)$. By irreducibility of

(M, R) , the disc A' cobounds a 3-ball B with some disc in R . If B is disjoint from D then A is boundary parallel. If instead B contains D then A is compressible, contradicting incompressibility of F .

✿ **The easy case.** Let us deal first with the case where every component of R is a torus; as a consequence, the same holds for the components of F . By assumption, we have that M is irreducible, boundary irreducible, atoroidal³, and not a Seifert manifold; additionally, every boundary component of M is a torus. By geometrisation, the interior of M admits a complete finite-volume hyperbolic structure. We can now invoke [16, Theorem 6.1] and conclude that:

- if (M', R', F') is another reasonable strongly simple 3-manifold triple such that each component of R' and F' is a torus, then we can algorithmically decide whether M and M' are homeomorphic or not;
- the group $\mathbf{H}(M)$ is finite and can be computed algorithmically.

Note that $\mathbf{H}(M, R, F)$ is a subgroup of $\mathbf{H}(M)$, and checking if an element of $\mathbf{H}(M)$ belongs to $\mathbf{H}(M, R, F)$ is trivial; this is enough to prove statement 2. As far as statement 1 is concerned, we simply check whether M and M' are homeomorphic. If they are, we find a homeomorphism $f: M \rightarrow M'$ and consider the (finitely many) elements of $\mathbf{H}(M')f$: if any of these gives a homeomorphism $(M, R, F) \rightarrow (M', R', F')$ then the two triples are homeomorphic, otherwise they are not.

✿ **Doubling M .** We now turn to the more involved case where not all components of R are tori. Let R_0 be the union of the components of R which are not tori. Let \overline{M} be the result of gluing two copies of M along R_0 . More precisely, we define

$$\overline{M} = M \times \{0, 1\} / \{(x, 0) \sim (x, 1) : x \in R_0\}.$$

By assumption, the surface R_0 is non-empty, so the 3-manifold \overline{M} is connected. Note that every boundary component of \overline{M} is either a copy of a torus component of R , a copy of a torus component of F , or the union of two copies of an annulus component of F . As a consequence, we see that every component of $\partial\overline{M}$ is a torus.

The surface R_0 has a canonical copy \overline{R}_0 in \overline{M} – specifically, the image in the quotient of $R_0 \times \{0\}$; this surface is incompressible in \overline{M} . More generally, if S is a surface in M , we let \overline{S} be the “doubled surface”, that is the image in the quotient of the surface $S \times \{0, 1\} \subseteq M \times \{0, 1\}$. Similarly, if $f: (M, R_0) \rightarrow (M, R_0)$ is a function, we let $\overline{f}: \overline{M} \rightarrow \overline{M}$ be the function defined by $\overline{f}(x, i) = (f(x), i)$ for $(x, i) \in M \times \{0, 1\}$.

Note that \overline{M} comes equipped with a natural orientation-reversing involution $\iota: \overline{M} \rightarrow \overline{M}$ defined by $\iota(x, i) = (x, 1 - i)$; for every homeomorphism $f: (M, R, F) \rightarrow (M, R, F)$, we have the relation $\overline{f}\iota = \iota\overline{f}$. Finally, let us remark that, since \overline{R}_0 is incompressible, the 3-manifold \overline{M} is sufficiently large.

✿ **The double is irreducible and boundary irreducible.** Irreducibility of \overline{M} follows immediately from the fact that M is irreducible and \overline{R}_0 is incompressible in \overline{M} .

Let now D be a disc properly embedded in \overline{M} . Up to isotopy, we can assume that D is in general position with respect to \overline{R}_0 . Additionally, let us isotope D so that the number of components of $D \cap \overline{R}_0$ is as small as possible. A standard innermost circle

³There are several slightly different definitions of “atoroidal” in the literature. In this article, by “ X is atoroidal” we mean that every incompressible torus in X is boundary parallel in X or, equivalently, that the triple $(X, \emptyset, \partial X)$ is simple. Note that other definitions could be inequivalent in general, but they all agree with ours when X is not a Seifert manifold.

argument on D shows that we can assume that all these components are arcs. If $D \cap \overline{R_0}$ is empty, then D can be interpreted as a compression disc for either R or F in M , hence it must be trivial in M and in \overline{M} as well. Otherwise, let a be an arc in $D \cap \overline{R_0}$ which is outermost in D ; the arc a cuts off a disc E from D such that $E \cap \overline{R_0} = a$. Note that $a \subseteq \overline{R_0}$, and $\partial D \setminus \text{int}(a)$ is an arc in \overline{F} . The disc E can then be interpreted as a boundary compression disc for F in (M, R) . By boundary incompressibility of F , it is easy to see that D can be isotoped to remove a from the intersection $D \cap \overline{R_0}$, thus contradicting our minimality assumption.

✿ **The double is strongly simple.** More precisely, we prove that the triple $(\overline{M}, \overline{R_1}, \overline{F})$ is strongly simple, where $R_1 = R \setminus R_0$ is the unions of the torus components of R . First, let us show that \overline{M} is not a Seifert manifold. If by contradiction \overline{M} admits a Seifert fibration, then up to isotopy we can assume that $\overline{R_0}$ is either horizontal or vertical. Since M is homeomorphic to either component of $\overline{M} \setminus \overline{R_0}$, it is easy to see that:

- if $\overline{R_0}$ is horizontal, then M inherits an I -bundle structure; the surfaces R and F have no torus components, and in fact they are the horizontal and vertical boundary respectively of the fibration inherited by M ;
- if $\overline{R_0}$ is vertical, then M is itself Seifert-fibred, in such a way that R and F are unions of fibres.

Both cases are incompatible with (M, R, F) being strongly simple.

We now address the “torus” condition for simple triples. Let T be an incompressible torus in \overline{M} . By the equivariant torus theorem (see [8, Corollary 4.6]), we can assume that T is either disjoint from or equal to $\iota(T)$.

- If T is disjoint from $\iota(T)$, then we can think of T as a torus in M ; incompressibility of $\overline{R_0}$ in \overline{M} implies that T is incompressible in M . Since (M, R, F) is a simple triple, T must be parallel in M to a component of R or F . But then T is boundary parallel in \overline{M} as well.
- If $T = \iota(T)$, then $T = \overline{A}$ for some annulus A properly embedded in (M, R) . The annulus A is incompressible in M since T is incompressible in \overline{M} . Then A must be parallel in (M, R) to an annulus in R or in F ; the first case is impossible, for otherwise T would be compressible in \overline{M} . Therefore, A is parallel in (M, R) to an annulus component A' of F . We conclude that $T = \overline{A}$ is parallel to the boundary component $\overline{A'}$ of \overline{M} .

Finally, let us consider the “annulus” condition. Let A be an incompressible annulus properly embedded in $(\overline{M}, \overline{R_1})$. If A is boundary compressible then it is boundary parallel. Otherwise, by the equivariant annulus theorem (see [15]), we can assume that A is either disjoint from or equal to $\iota(A)$.

- If A is disjoint from $\iota(A)$, then we can think of A as an incompressible annulus properly embedded in (M, R) . By simplicity of (M, R, F) , this implies that A is boundary parallel in M and, hence, in \overline{M} as well.
- If $A = \iota(A)$, then $A = \overline{A_0}$ for some incompressible annulus $A_0 \subseteq M$ whose boundary intersects two boundary components of M , which contradicts simplicity of (M, R, F) .

✿ **The double is hyperbolic.** It is not hard to verify that the triple $(\overline{M}, \overline{R_1}, \overline{F})$ is reasonable. Since, by definition, every component of $\overline{R_1}$ is a torus, we have already

proved statements **1** and **2** for this triple. In the process, we also noted that the interior of \overline{M} admits a complete finite-volume hyperbolic metric. In fact, for our purposes, it will be more convenient to think of \overline{M} as a compact hyperbolic 3-manifold whose boundary components are flat tori – in other words, as the result of truncating the cusps of a complete finite-volume non-compact hyperbolic 3-manifold. By combining Mostow-Prasad rigidity with [27, Theorem 7.1], we have the following.

Every self-homeomorphism of \overline{M} is isotopic to a unique isometry $\overline{M} \rightarrow \overline{M}$. **(1)**

As a consequence, let us remark that we can choose the hyperbolic metric on \overline{M} in such a way that ι is an isometry. In fact, the involution ι is isotopic to an isometry $\hat{\iota}: \overline{M} \rightarrow \overline{M}$. Note that $\hat{\iota}$ is itself an involution: the isometry $\hat{\iota} \circ \hat{\iota}$ is isotopic to $\iota \circ \iota = \text{id}$ and, therefore, equal to the identity. We now apply [25, Theorem B]⁴ to deduce that $\hat{\iota} = h\iota h^{-1}$ for some homeomorphism $h: \overline{M} \rightarrow \overline{M}$. Then ι is an isometry with respect to the pull-back of the hyperbolic metric on \overline{M} by h .

✿ Solving the homeomorphism problem. Consider another reasonable strongly simple 3-manifold triple (M', R', F') ; suppose that, additionally, there is at least one component of R' which is not a torus – otherwise (M, R, F) and (M', R', F') could not possibly be homeomorphic. Define \overline{M}' , ι' , and R'_1 like we did with \overline{M} , ι , and R_1 . Recall that \overline{M}' can be endowed with a finite-volume hyperbolic metric with flat boundary, with respect to which ι' is an isometry. Naturally, statement **(1)** also holds for homeomorphisms $\overline{M} \rightarrow \overline{M}'$.

If the triples $(\overline{M}, \overline{R}_1, \overline{F})$ and $(\overline{M}', \overline{R}'_1, \overline{F}')$ are not homeomorphic – which we can algorithmically decide thanks to statement **1** – then clearly neither are (M, R, F) and (M', R', F') . Otherwise, let

$$f: (\overline{M}, \overline{R}_1, \overline{F}) \longrightarrow (\overline{M}', \overline{R}'_1, \overline{F}')$$

be a homeomorphism. We claim that the following are equivalent:

i there is a homeomorphism $g: (M, R, F) \rightarrow (M', R', F')$ such that

$$\bar{g}: (\overline{M}, \overline{R}_1, \overline{F}) \longrightarrow (\overline{M}', \overline{R}'_1, \overline{F}')$$

is isotopic to f ;

ii the homeomorphism $f^{-1}\iota'f\iota: \overline{M} \rightarrow \overline{M}$ is isotopic to the identity.

It is clear that **i** implies **ii**. Conversely, if **ii** holds, we can assume up to isotopy that f restricts to an isometry on the interior of \overline{M} . But then $f^{-1}\iota'f\iota$ is an isometry of \overline{M} which is isotopic to the identity and, hence, equal to it. In other words, the following diagram commutes:

$$\begin{array}{ccc} (\overline{M}, \overline{R}_1, \overline{F}) & \xrightarrow{f} & (\overline{M}', \overline{R}'_1, \overline{F}') \\ \downarrow \iota & & \downarrow \iota' \\ (\overline{M}, \overline{R}_1, \overline{F}) & \xrightarrow{f} & (\overline{M}', \overline{R}'_1, \overline{F}') \end{array}$$

We then conclude that f is induced by a homeomorphism $g: M \rightarrow M'$; it is easy to see that g must send F to F' and, therefore, R to R' .

⁴Using Tollefson's notation, the case where $\beta \neq \text{id}$ does not apply here, since the fundamental group of complete finite volume hyperbolic 3-manifolds has trivial centre.

In order to decide whether (M, R, F) and (M', R', F') are homeomorphic, we can then carry out the following procedure. Let f_0 be any homeomorphism from $(\overline{M}, \overline{R_1}, \overline{F})$ to $(\overline{M'}, \overline{R'_1}, \overline{F'})$. Thanks to statement **1**, we can compute representatives for the finite group $\mathbf{H}(\overline{M'}, \overline{R'_1}, \overline{F'})$. For each element f of $\mathbf{H}(\overline{M'}, \overline{R'_1}, \overline{F'})f_0$, we check whether $f^{-1}f_0$ is isotopic to the identity of \overline{M} (we can do this algorithmically since $\mathbf{H}(\overline{M}, \overline{R_1}, \overline{F})$ is finite and we can compute representatives for it). If this happens for some f , then (M, R, F) and (M', R', F') are homeomorphic, otherwise they are not.

✿ **Computing the mapping class group.** By the very same argument as the one presented in the previous step, we can compute representatives of the subgroup

$$\{f \in \mathbf{H}(\overline{M}, \overline{R_1}, \overline{F}) : f \text{ is isotopic to } \overline{g} \text{ for some } g \in \mathbf{H}(M, R, F)\} \leq \mathbf{H}(\overline{M}, \overline{R_1}, \overline{F}).$$

We claim that the group homomorphism $\mathbf{H}(M, R, F) \rightarrow \mathbf{H}(\overline{M}, \overline{R_1}, \overline{F})$ sending g to \overline{g} is injective; in other words, if \overline{g} is isotopic to the identity in \overline{M} then g is isotopic to the identity in (M, R, F) .

✿ First of all, if \overline{g} is isotopic to the identity, then we can isotope g in (M, R, F) so that it restricts to the identity on $R_1 \cup F$. As a consequence, the homeomorphism \overline{g} restricts to the identity on $\partial\overline{M}$.

✿ Let $\pi: \widetilde{M} \rightarrow \overline{M}$ be the universal covering; we will think of \widetilde{M} as a closed subset of hyperbolic 3-space \mathbb{H}^3 , whose complement is a collection of open horoballs. We define a retraction $r: \mathbb{H}^3 \rightarrow \widetilde{M}$ as follows. If x is a point of \widetilde{M} we set $r(x) = x$. If instead x is contained in a horoball B , we set $r(x) = \gamma \cap \partial B$, where γ is the geodesic emanating from the centre of B and passing through x . Finally, if x and y are two (not necessarily distinct) points of \widetilde{M} , we define $\gamma_{x,y}: [0, 1] \rightarrow \mathbb{H}^3$ to be the geodesic such that $\gamma_x(0) = x$ and $\gamma_{x,y}(1) = y$.

✿ Let $\widetilde{g}: \widetilde{M} \rightarrow \widetilde{M}$ be a homeomorphism that is a lift of \overline{g} . Since \overline{g} is isotopic to the identity in \overline{M} and restricts to the identity on $\partial\overline{M}$, we can choose \widetilde{g} so that it restricts to the identity on $\partial\widetilde{M}$. Let us define a homotopy $\widetilde{g}_t: \widetilde{M} \rightarrow \widetilde{M}$ as $\widetilde{g}_t(x) = r(\gamma_{x, \widetilde{g}(x)}(t))$, so that $\widetilde{g}_0 = \text{id}$ and $\widetilde{g}_1 = \widetilde{g}$. It is easy to check that \widetilde{g}_t is invariant under deck transformations of π , hence it descends to a homotopy $\overline{g}_t: \overline{M} \rightarrow \overline{M}$ between the identity and \overline{g} . By definition, for each $t \in [0, 1]$ the homeomorphism \overline{g}_t fixes $\partial\overline{M}$ pointwise.

✿ Since $\iota\overline{g} = \overline{g}$ and the homotopy \overline{g}_t has been defined only in terms of the metric, we must have that $\iota\overline{g}_t = \overline{g}_t$ for every $t \in [0, 1]$. In particular, we find that $\overline{g}_t(\overline{R_0}) \subseteq \overline{R_0}$, from which we deduce that $g|_{R_0}$ is homotopic to the identity in R_0 , through a homotopy that fixes ∂R_0 . By a classical result of Epstein (namely, [4, Theorem 6.4]), this implies that $g|_{R_0}$ is actually isotopic to the identity in R_0 , and we can therefore assume that $g|_{\partial}: \partial M \rightarrow \partial M$ is the identity.

✿ Note that the homotopy \widetilde{g}_t preserves the two components of $\widetilde{M} \setminus \overline{R_0}$ for each $t \in [0, 1]$. Therefore, there is an induced homotopy $g_t: M \rightarrow M$ between g and the identity, which restricts to the identity on ∂M for each $t \in [0, 1]$. By [27, Theorem 7.1], this implies that g is isotopic to the identity in (M, R, F) .

As a final remark, note that, given $f \in \mathbf{H}(\overline{M}, \overline{R_1}, \overline{F})$, we can algorithmically find $g \in \mathbf{H}(M, R, F)$ such that f is isotopic to \overline{g} , provided that one exists. Therefore, we have presented a complete algorithm to compute representatives of $\mathbf{H}(M, R, F)$. \square

2.6 Computing the JSJ decomposition

After analysing the types of pieces individually, we now show how to actually compute the JSJ decomposition. The next three propositions explain how to recognise Seifert

pieces, I -bundle pieces and (strongly) simple pieces respectively. Once we have these three algorithms, computing the JSJ decomposition is straightforward, as the proof of Theorem 2.15 witnesses.

Proposition 2.12. *There is an algorithm which, given a 3-manifold $(n+1)$ -tuple (M, \mathbf{R}) as input, decides whether it is a Seifert $(n+1)$ -tuple or not.*

Proof. If $R_1 \cup \dots \cup R_n \neq \partial M$, or one of the surfaces R_1, \dots, R_n has a component which is not a torus or an annulus, then the answer is no. We can use [11, Algorithm 8.2] to decide whether M is a Seifert manifold or not. If it is not then we are done. If M is homeomorphic to a solid torus, a product $T^2 \times I$, or a twisted I -bundle $U_2 \tilde{\times} I$, then we can easily decide if (M, \mathbf{R}) is a Seifert $(n+1)$ -tuple or not. Otherwise, find the unique Seifert fibration for M ; if it is compatible with the surfaces R_1, \dots, R_n then the answer is yes; otherwise, the answer is no. \square

Proposition 2.13. *There is an algorithm which, given a 3-manifold triple (M, R, F) as input, decides whether it is an I -bundle triple or not.*

Proof. If any of the following hold, then (M, R, F) is not an I -bundle triple:

- at least one component of F is not an annulus;
- R has zero or more than two components;
- R has two non-homeomorphic components;
- there are one component of R and one component of F which are disjoint.

Otherwise, let B be one component of R if R has two components, or the non-orientable surface doubly covered by R if R has only one component; if (M, R, F) is an I -bundle triple, then it is the (product or twisted) I -bundle over B . Let $p \subseteq F$ be the union of the core curves of the annulus components of F . Let X be the I -bundle over B , and denote by $q \subseteq \partial X$ the union of the core curves of the annuli making up the vertical boundary of X . Then (M, R, F) is an I -bundle triple if and only if the 3-manifolds with boundary pattern (M, p) and (X, q) are homeomorphic. \square

Proposition 2.14. *There is an algorithm which, given a reasonable 3-manifold triple (M, R, F) as input, decides whether it is a simple triple or not.*

Proof. Let us first deal with the cases where M is homeomorphic to a solid torus or to $T^2 \times I$.

- In a solid torus, every incompressible annulus is boundary parallel. Therefore, there are only finitely many incompressible annuli properly embedded in (M, R) up to isotopy. We can enumerate them and check whether any of them is not parallel to a component of R or F .
- When M is homeomorphic to $T^2 \times I$, the triple (M, R, F) is not simple if and only if one of the following happens:
 - $T^2 \times \{i\} \subseteq R$ and $T^2 \times \{i\} \neq R$ for some $i \in \partial I$;
 - there are isotopic curves $c_0, c_1 \subseteq T^2$ such that $c_i \times \{i\} \subseteq \partial R$ for each $i \in \partial I$;
 - there are two annulus components R_1 and R_2 of $R \cap (T^2 \times \{i\})$ for some $i \in \partial I$, such that neither component of $(T^2 \times \{i\}) \setminus \text{int}(R_1 \cup R_2)$ is contained in F .

From now on, we will assume that M is neither a solid torus nor homeomorphic to $T^2 \times I$. Moreover, thanks to Lemma 2.10, we can assume that $\partial M = R \cup F$. Let $p = R \cap F$; using the terminology of [20], we have that the 3-manifold with boundary pattern (M, p) is boundary irreducible.

✿ **Detecting essential tori.** We can use [20, Lemma 6.4.7] to decide if M contains an incompressible torus which is not parallel to a boundary component of M . If it does, then (M, R, F) is not simple. Likewise, if M has a torus boundary component which is not a component of R or F , then the triple (M, R, F) is not simple. If none of these two condition is satisfied, then (M, R, F) satisfies the “torus” condition for being simple, and we only need to check for incompressible annuli.

✿ **Detecting essential annuli.** We can use [20, Lemma 6.4.8] to decide if M contains an incompressible annulus A with the following properties:

- the boundary of A lies in $\text{int}(R)$;
- A is not parallel to a union of annulus components of R and F ;
- every non-trivial boundary compression disc for A intersects F .

If there is such an annulus, then (M, R, F) is not simple. If there is an annulus component of R whose two boundary curves lie in different components of F , then (M, R, F) is not simple. If none of these two conditions holds, then (M, R, F) satisfies the “annulus” condition for being simple, and is therefore simple. \square

Theorem 2.15. *Let (M, R) be an irreducible sufficiently large 3-manifold pair. Suppose that no component of R is a disc and that $\partial M \setminus \text{int}(R)$ is a (possibly empty) collection of annuli. There is an algorithm which, given as input the pair (M, R) , returns the JSJ system F of (M, R) . Moreover, each component of $(M, R) \setminus F$ is reasonable.*

Proof. It is easy to check that, under the assumptions made in the statement, the pieces of the JSJ decomposition of (M, R) are reasonable. We can search through all surfaces F' properly embedded in (M, R) whose components are annuli and tori, until we find one satisfying properties **i** and **ii** of Definition 2.5 and, additionally, the property that each component of $(M, R) \setminus F'$ is reasonable. All these conditions can be checked algorithmically, and we are guaranteed to find at least one such surface F' . Then, for every union of components of F' , we check whether it still satisfies the above properties and, among those that do, we return one that is minimal. \square

Let us remark that the additional constraints we put on the pair (M, R) – namely the fact that R has no disc components and $\partial M \setminus \text{int}(R)$ is union of annuli – are a consequence of the assumption of reasonableness we require for simple pieces.

2.7 Piecing things together

For the sake of convenience, we will slightly deviate from the statement of Theorem 2.1 and classify pieces of the JSJ decomposition as follows. If (M, R) is an irreducible sufficiently large 3-manifold pair and F is its JSJ system, we say that a component (N, R', F') of $(M, R) \setminus F$ is:

- a *Seifert piece* if $(N, R', F', \partial N \setminus \text{int}(R' \cup F'))$ is a Seifert 4-tuple;
- an *I-bundle piece* if (N, R', F') is an I -bundle triple;

- a *strongly simple piece* if (N, R', F') is a strongly simple triple.

It is clear that every component of $(M, R) \setminus F$ is either a Seifert piece, an I -bundle piece, or a strongly simple piece. Note that the three possibilities are not mutually exclusive: it is possible for a component to be a Seifert piece and an I -bundle piece at the same time.

Theorem 2.16. *Let (M, R) be an irreducible sufficiently large 3-manifold pair. Suppose that no component of R is a disc and that $\partial M \setminus \text{int}(R)$ is a (possibly empty) collection of annuli. There is an algorithm which, given as input the pair (M, R) , returns*

- a finite collection \mathcal{F} of self-homeomorphisms of $(\partial M, R)$, and
- a finite collection $\mathcal{C} = \{(a_1, b_1), \dots, (a_m, b_m)\}$, where $a_1, \dots, a_m, b_1, \dots, b_m$ are curves in R with $a_1 \cap b_1 = \dots = a_m \cap b_m = \emptyset$,

such that

$$\mathbf{H}(M, R)|_{\partial} = \bigcup_{f \in \mathcal{F}} \langle \tau_{a_1} \tau_{b_1}^{-1}, \dots, \tau_{a_m} \tau_{b_m}^{-1} \rangle f$$

as a subgroup of $\mathbf{H}(\partial M, R)$. Moreover, if every I -bundle piece in the JSJ decomposition of (M, R) is a product I -bundle over $\Sigma_0, \Sigma_{0,1}, \Sigma_{0,2}$, or $\Sigma_{0,3}$, or the twisted I -bundle over $U_1, U_{1,1}$, or $U_{1,2}$, then the curves in \mathcal{C} are pairwise disjoint.

Proof. We use Theorem 2.15 to compute the JSJ system F of (M, R) , and note that all the pieces are reasonable.

• **Exceptional cases.** First of all, let us deal with the cases where the JSJ decomposition of (M, R) is trivial or contains exceptional Seifert pieces.

• Suppose that F is empty. Then (M, R, \emptyset) is either a strongly simple piece, an I -bundle piece, or a Seifert piece. If it is strongly simple, then we can just return $\mathcal{F} = \mathbf{H}(M, R)|_{\partial}$ and $\mathcal{C} = \emptyset$ thanks to Proposition 2.11. If $(M, R, \partial M \setminus \text{int}(R))$ is a Seifert triple, then we conclude immediately by Proposition 2.6. Finally, if (M, R, \emptyset) is an I -bundle triple, then we apply Proposition 2.9.

• Suppose that a component (N, R', F') is homeomorphic to $(T^2 \times I, \emptyset, T^2 \times \partial I)$. The two boundary components of N must come from the same torus in F , otherwise the surface F would not be minimal. We then have that M is homeomorphic to the mapping torus of some homeomorphism $T^2 \rightarrow T^2$, and we simply return $\mathcal{F} = \{\text{id}\}$ and $\mathcal{C} = \emptyset$.

From now on, we can therefore assume that F is non-empty and that there are no $(T^2 \times I, \emptyset, T^2 \times \partial I)$ components in $(M, R) \setminus F$. As a consequence of Proposition 2.3, it is easy to see that every Seifert piece either has a unique Seifert fibration, or it has two and it is homeomorphic to $(U_2 \tilde{\times} I, \emptyset, \partial(U_2 \tilde{\times} I))$. Since the two fibrations of $U_2 \tilde{\times} I$ are inequivalent, let us arbitrarily and consistently pick one for every $(U_2 \tilde{\times} I, \emptyset, \partial(U_2 \tilde{\times} I))$ component of $(M, R) \setminus F$ – say the one over the Möbius band. Then every Seifert piece in the JSJ decomposition of (M, R) has a distinguished fibration, either the unique one – for pieces which are not homeomorphic to the twisted I -bundle over a Klein bottle – or the one we have selected. Moreover, every homeomorphism between Seifert pieces can be isotoped to be fibre-preserving.

• **JSJ graph.** We closely follow the construction of the JSJ graph given by Kuperberg in [16, Section 6.3], with slight changes to accommodate the presence of boundary. More precisely, we define a graph Γ as follows.

- There is one vertex for each component of $(M, R) \setminus F$.

- There is one edge for each component of F . Each edge joins the two (not necessarily distinct) vertices corresponding to the JSJ pieces whose boundaries this component lies on.

An *automorphism* of the JSJ graph Γ is an automorphism $\varphi: \Gamma \rightarrow \Gamma$ of the underlying graph, together with some additional data:

- for each vertex (N, R', F') representing a strongly simple piece, a homeomorphism $\varphi|_N$ from (N, R', F') to the component of $(M, R) \setminus F$ associated to the vertex $\varphi(N, R', F')$, defined up to isotopy in $\varphi(N, R', F')$;
- for each edge X , a homeomorphism $\varphi|_X$ from X to the component of F associated to the edge $\varphi(X)$, defined up to isotopy in $\varphi(X)$; note that if X is an annulus, there are only two possible values for $\varphi|_X$.

Moreover, we enforce the following consistency condition for each strongly simple piece (N, R', F') : for each edge X having (N, R', F') as an endpoint, we ask that the restriction of $\varphi|_N$ to X is isotopic to $\varphi|_X$ in $\varphi(X)$. If both endpoints of X are equal to (N, R', F') , then we ask that this condition holds for both copies of X in F' .

• **Finitely many automorphisms.** Every homeomorphism $f: (M, R) \rightarrow (M, R)$ preserving F induces an automorphism φ_f of Γ in a natural way. We claim that the set

$$\Phi = \{\varphi_f : f \text{ is a self-homeomorphism of } (M, R) \text{ preserving } F\}$$

is finite. In fact, the only issue that needs addressing is the potentially infinite number of homeomorphisms between JSJ tori connecting two Seifert pieces. Let T be such a torus, connecting Seifert pieces (N_1, R_1, S_1) and (N_2, R_2, S_2) . Let c_1 and c_2 be curves in T which are fibres of (N_1, R_1, S_1) and (N_2, R_2, S_2) respectively; note that c_1 and c_2 are not isotopic in T , otherwise we could remove T from the JSJ system F . Consider a homeomorphism $f: (M, R) \rightarrow (M, R)$ preserving F . Necessarily, the curves c_1 and c_2 will be sent by f to curves in $\varphi_f(T)$ which are isotopic to fibres c'_1 and c'_2 of $\varphi_f(N_1, R_1, S_1)$ and $\varphi_f(N_2, R_2, S_2)$ respectively. Since the curves $c_1, c_2, c'_1,$ and c'_2 do not depend on f , and there are only finitely many isotopy classes of homeomorphisms $T \rightarrow \varphi_f(T)$ sending c_1 to c'_1 and c_2 to c'_2 , we conclude that there are only finitely many choices for $\varphi_f|_T$.

We remark that this argument is still valid if (N_1, R_1, S_1) and (N_2, R_2, S_2) are the same piece; in this case, we simply have to consider the possibility that f sends c_1 to c'_2 and c_2 to c'_1 .

• **Computable automorphisms.** Our argument for the finiteness of Φ actually provides an algorithm to compute a finite superset Φ' of Φ . As a consequence, in order to compute Φ , we simply need an algorithm to decide whether an automorphism of Γ is induced by a self-homeomorphism of (M, R) preserving F .

Let φ be an automorphism of Γ . If (N, R', F') is not homeomorphic to $\varphi(N, R', F')$ for some piece (N, R', F') then clearly φ is not induced by a homeomorphism. Otherwise, for each piece (N, R', F') which is not strongly simple, let us ask the following question: is there a homeomorphism $f_N: (N, R', F') \rightarrow \varphi(N, R', F')$ such that, for each component X of F' , the homeomorphisms $f_N|_X$ and $\varphi|_X$ are isotopic in $\varphi(X)$? If the answer – which we can compute thanks to Corollary 2.7 and Proposition 2.8 – is no, then φ cannot be induced by a homeomorphism. Otherwise, we easily see that φ belongs to Φ . In fact, we can isotope the homeomorphisms f_N so that they agree on their shared boundary F , and then combine them to produce a homeomorphism $f: (M, R) \rightarrow (M, R)$ preserving F ; here, for ease of notation, we have used f_N to refer to the homeomorphism $\varphi|_N$ for strongly simple pieces (N, R', F') .

We can therefore algorithmically construct a finite set \mathcal{F}_0 by picking, for each $\varphi \in \Phi$, a self-homeomorphism of (M, R) preserving F and inducing φ . This finite set has the property that every homeomorphism $f: (M, R) \rightarrow (M, R)$ can be isotoped so that, for some $f_0 \in \mathcal{F}_0$, the homeomorphism ff_0 is the identity on F and on all the strongly simple pieces of the JSJ decomposition. In fact, it is enough to isotope f so to preserve F and pick $f_0 \in \mathcal{F}_0$ such that $\varphi_{f_0} = \varphi_{f^{-1}}$, so that ff_0 induces the trivial automorphism of Γ . Up to further isotoping f , we can then assume that ff_0 is the identity on F and on strongly simple pieces, as desired.

✿ **Finitely many Dehn twists.** Finally, let us show how to compute the sets \mathcal{F} and \mathcal{C} . For each Seifert or I -bundle piece (N, R', F') , we apply Propositions 2.6 and 2.9 to find

- a finite collection \mathcal{F}_N of self-homeomorphisms of $(\partial N, R')$ fixing F' pointwise, and
- a finite collection $\mathcal{C}_N = \{(a_{N,1}, b_{N,1}), \dots, (a_{N,m_N}, b_{N,m_N})\}$ where $a_{N,1}, \dots, a_{N,m_N}, b_{N,1}, \dots, b_{N,m_N}$ are curves in $\partial N \setminus (F' \cup \partial R')$ with $a_{N,1} \cap b_{N,1} = \dots = a_{N,m_N} \cap b_{N,m_N} = \emptyset$,

such that

$$\mathbf{H}_{F'}(N, R')|_{\partial} = \bigcup_{f \in \mathcal{F}_N} \langle \tau_{a_{N,i}} \tau_{b_{N,i}}^{-1} : 1 \leq i \leq m_N \rangle f$$

as a subgroup of $\mathbf{H}_{F'}(\partial N, R')$. Moreover, if the piece (N, R', F') is a product I -bundle over $\Sigma_0, \Sigma_{0,1}, \Sigma_{0,2}$, or $\Sigma_{0,3}$, or a twisted I -bundle over $U_1, U_{1,1}$, or $U_{1,2}$, then the curves in \mathcal{C}_N are pairwise disjoint. Let us extend the homeomorphisms in \mathcal{F}_N to self-homeomorphisms of $(\partial M, R)$ by setting them equal to the identity on $\partial M \setminus N$. Finally, if some curve in \mathcal{C}_N lies in $\partial M \setminus R$, we can replace it with a parallel curve in R without changing the isotopy class of the corresponding Dehn twist. It is then easy to see that the algorithm can return the collections

$$\mathcal{F} = \{f_0|_{\partial} : f_0 \in \mathcal{F}_0\} \cup \bigcup_N \mathcal{F}_N \quad \text{and} \quad \mathcal{C} = \bigcup_N \mathcal{C}_N. \quad \square$$

Remark 2.17. Translated in the language of Matveev's boundary patterns, Theorem 2.16 can be used to compute the mapping class group of a Haken 3-manifold (M, p) with boundary pattern such that p is a collection of simple closed curves and no component of $\partial M \setminus p$ is a disc.

As a consequence of Theorem 2.16 and the work of Gordon and Luecke (see [7]), we have the following. Note that statement 3 can also be proved using an oriented version of Matveev's boundary patterns.

Corollary 2.18. *Let $M \subseteq S^3$ be the complement of a non-trivial knot.*

- 1 *Every self-homeomorphism of M sends the meridian curve of ∂M to itself.*
- 2 *The group $\mathbf{H}(M)|_{\partial}$ is either trivial or generated by the homeomorphism*

$$(-\text{id}): \partial M \longrightarrow \partial M$$

inducing multiplication by -1 on $H_1(\partial M)$.

- 3 *There is an algorithm which, given M as input, returns representatives of $\mathbf{H}(M)|_{\partial}$.*

Proof. Statement ❶ follows from [7]; statement ❷ is an immediate consequence. As far as statement ❸ is concerned, we can apply Theorem 2.16 to the pair $(M, \partial M)$. Since homeomorphisms of the form $\tau_a \tau_b^{-1}$ cannot act as $(-\text{id})$ on $H_1(\partial M)$, it is enough to check whether the collection \mathcal{F} returned by the algorithm contains any non-trivial homeomorphism. \square

3 An algorithmic problem on free groups

3.1 Introducing band systems

We now take a break from topology and enter the realm of combinatorics and free groups. The bulk of Section 3 will be devoted to proving Theorem 3.2. The reader should however not be surprised by how seemingly distant the statement of the theorem is to anything concerning isotopies of surfaces. In fact, an easy consequence of the theorem – namely Corollary 3.3 – will allow us to answer an algorithmic question about extensions of homeomorphisms to the interior of handlebodies.

We start by introducing the notation we will need to talk about cancellation of words in free groups, namely (marked) band systems and bundling maps. Section 3.2 is entirely dedicated to the proof of Theorem 3.2, while the topological corollary is presented in Section 3.3.

Definition 3.1. Let n be a positive integer. A *band system*⁵ of length $2n$ is a collection

$$B = \{(a_1, b_1), \dots, (a_n, b_n)\}$$

such that:

- ❶ $\{a_1, b_1, \dots, a_n, b_n\} = \{1, \dots, 2n\}$;
- ❷ for each $1 \leq i \leq n$, the inequality $a_i < b_i$ holds;
- ❸ there is no pair of indices i, j such that $a_i < a_j < b_i < b_j$.

The individual pairs (a_i, b_i) are called *bands*, with a_i being the *left endpoint* and b_i being the *right endpoint*.

Figure 3.1a shows a graphical depiction of the band system

$$B = \{(1, 8), (2, 5), (3, 4), (6, 7)\} \text{ of length } 2n = 8,$$

hopefully justifying the name. If we think of a pair $(a, b) \in B$ as a physical band connecting the elements a and b , then condition ❸ ensures that these bands do not intersect.

Definition 3.2. A *marked band system* is a pair (B, p) where B is a band system of length – say – $2n$, and $p = (p_1, \dots, p_r)$ is a weakly increasing sequence of numbers with

$$p_i \in \left\{ \frac{1}{2}, 1 + \frac{1}{2}, \dots, 2n + \frac{1}{2} \right\} \quad \text{for } 1 \leq i \leq r.$$

⁵It was brought to our attention by an anonymous referee that what we call “band systems” here were already known in the literature under the name “crossingless matchings”. However, to the best of our knowledge, the theory and results developed in this section are still novel.

Figure 3.1b depicts the same band system as Figure 3.1a, equipped with the marking

$$p = \left(\frac{1}{2}, 3 + \frac{1}{2}, 7 + \frac{1}{2} \right).$$

We think of an element $p_i = a + 1/2$ as a “separator” between the integers a and $a + 1$.




Figure 3.1. A band system (a) and a marked band system (b).

Let (B, p) be a marked band system. The *maximal bundle*⁶ of (B, p) is, loosely speaking, the marked band system (B', p') constructed from (B, p) by bundling together parallel bands as much as possible, while making sure that the bundles don't cross any marked spot. Formally, let $2n$ be the length of B . Consider the equivalence relation \sim on $\{1, \dots, 2n\}$ generated by

$$a_1 \sim a_2 \text{ and } b_1 \sim b_2 \text{ whenever } \begin{cases} (a_1, b_1), (a_2, b_2) \in B, \\ a_2 = a_1 + 1, b_2 = b_1 - 1, \text{ and} \\ a_1 + \frac{1}{2}, b_1 - \frac{1}{2} \notin p. \end{cases}$$

Let $E = \{1, \dots, 2n\} / \sim$ be the quotient set, and let $2n'$ be its cardinality (which is easily seen to be even). Since equivalence classes consist of consecutive integers, there is a natural bijection between E and $\{1, \dots, 2n'\}$ such that classes containing smaller integers are mapped to smaller integers. Let

$$\iota: \{1, \dots, 2n\} \longrightarrow \{1, \dots, 2n'\}$$

be the composition of the projection map $\{1, \dots, 2n\} \rightarrow E$ with said bijection; we will call ι the *bundling map*. We can then define

$$B' = \{(\iota(a), \iota(b)) : (a, b) \in B\}.$$

As far as the marking is concerned, if $p = (p_1, \dots, p_r)$, for each $1 \leq i \leq r$ we define

$$p'_i = \begin{cases} \frac{1}{2} & \text{if } p_i = \frac{1}{2}, \\ \iota\left(p_i - \frac{1}{2}\right) + \frac{1}{2} & \text{otherwise,} \end{cases}$$

⁶Note that this notion of “bundle” has nothing to do with that of “fibre bundle” used elsewhere in the article.

and set $p' = (p'_1, \dots, p'_r)$. It is not hard to verify that (B', p') is a marked band system.

As an example, let us look at Figure 3.2. The marked band system on the top is defined by

$$B = \{(1, 12), (2, 11), (3, 8), (4, 7), (5, 6), (9, 10)\} \text{ of length } 2n = 12,$$

$$p = \left(2 + \frac{1}{2}, 6 + \frac{1}{2}, 9 + \frac{1}{2}, 12 + \frac{1}{2}\right).$$

In order to construct the maximal bundle, we group together the bands $(1, 12)$ and $(3, 4)$, since they are parallel and do not cross any marked spot. We do the same for bands $(3, 8)$ and $(4, 7)$; note that the band $(5, 6)$ does not belong to the same bundle, since the spot $6 + 1/2$ is marked. The maximal bundle of (B, p) is then (B', p') , where

$$B' = \{(1, 8), (2, 5), (3, 4), (6, 7)\} \text{ of length } 2n' = 8,$$

$$p' = \left(1 + \frac{1}{2}, 4 + \frac{1}{2}, 6 + \frac{1}{2}, 8 + \frac{1}{2}\right).$$




Figure 3.2. A marked band system (on the top) and its maximal bundle (on the bottom). Bands of the same colour get bundled together; vertical arrows represent the bundling map ι .

Let us now remark a few properties of the bundling map ι .

- The bundling map sends bands of B to bands of B' . From now on, with slight abuse of notation, for each band $(a, b) \in B$ we will denote by $\iota(a, b)$ the band $(\iota(a), \iota(b)) \in B'$.

- If $\iota(a_1, b_1) = \iota(a_2, b_2)$ where (a_1, b_1) and (a_2, b_2) are bands of B with $a_1 < a_2 < b_2 < b_1$, then $(a_1 + i, b_1 - i) \in B$ and $\iota(a_1 + i, b_1 - i) = \iota(a_1, b_1)$ for each $0 \leq i \leq a_2 - a_1$. Moreover, no element of p lies between a_1 and a_2 or between b_2 and b_1 .

The following facts are completely elementary, but we prefer to state them here so as not to clutter the proof of Theorem 3.2.

Lemma 3.1. *Let (B, p) be a marked band system.*

- 1 If (a, b) is a band of B , then there is a band $(c, c + 1) \in B$ for some $a \leq c < b$.
- 2 Let (B', p') be the maximal bundle of (B, p) , with bundling map ι . If (a_1, b_1) and (a_2, b_2) are bands of B with $a_1 < a_2 < b_2 < b_1$, then either
 - $\iota(a_1, b_1) = \iota(a_2, b_2)$,
 - there is an element of p lying between a_1 and a_2 or between b_2 and b_1 , or
 - there is a band $(c, c + 1) \in B$ for some

$$c \in \{a_1, \dots, a_2 - 1\} \cup \{b_2, \dots, b_1 - 1\}.$$

Proof. In order to prove statement 1, simply consider the narrowest band amongst those having left endpoint in $\{a, \dots, b - 1\}$. We now turn to statement 2. If there is a band $(a, b) \in B$ with $a_1 < a < b < a_2$ or $b_2 < a < b < b_1$ then we can apply statement 1 and immediately conclude. Otherwise, for every band $(a, b) \in B$ we have that $a_1 \leq a \leq a_2$ if and only if $b_2 \leq b \leq b_1$. But then, using the notations from the definition of maximal bundle, it is clear that either

$$(a_1, b_1) \sim (a_1 + 1, b_1 - 1) \sim \dots \sim (a_2, b_2),$$

or there is marked spot (in the marking p) lying between a_1 and a_2 or between b_2 and b_1 . \square

Let (B', p') be a marked band system of length $2n'$. We say that (B', p') is *maximal* if it is its own maximal bundle. Of course, if (B', p') is maximal, there are infinitely many marked band systems (B, p) having (B', p') as their maximal bundle; in fact, these marked band systems can be parametrised quite nicely. An *unbundling map* for (B', p') is a function $\varphi: \{1, \dots, 2n'\} \rightarrow \mathbb{Z}_{>0}$ such that $\varphi(a') = \varphi(b')$ for each band $(a', b') \in B'$. There is a one-to-one correspondence between unbundling maps for (B', p') and marked band systems having (B', p') as their maximal bundle.

- Given a marked band system (B, p) whose maximal bundle is (B', p') , we can define an unbundling map φ as $\varphi(a') = |\iota^{-1}(a')|$, where ι is the bundling map for (B, p) .
- Conversely, let $\varphi: \{1, \dots, 2n'\} \rightarrow \mathbb{Z}_{>0}$ be an unbundling map for (B', p') . Let $2n = \varphi(1) + \dots + \varphi(2n')$; for each $1 \leq a \leq 2n$, define $\iota(a)$ to be the unique integer in $\{1, \dots, 2n'\}$ such that

$$\varphi(1) + \dots + \varphi(\iota(a) - 1) < a \leq \varphi(1) + \dots + \varphi(\iota(a)).$$

It is now easy to construct the unique marked band system (B, p) having (B', p') as maximal bundle and ι as bundling map. In fact, we can set

$$B = \bigcup_{(a', b') \in B'} \left\{ (a + i, b - i) : \begin{array}{l} \{a + 1, \dots, a + \varphi(a')\} = \iota^{-1}(a'), \\ \{b - \varphi(a'), \dots, b - 1\} = \iota^{-1}(b'), \\ 1 \leq i \leq \varphi(a') \end{array} \right\}$$

and $p = (p_1, \dots, p_r)$, where $p' = (p'_1, \dots, p'_r)$ and

$$p_i = \begin{cases} \frac{1}{2} & \text{if } p'_i = \frac{1}{2}, \\ \max\left(\iota^{-1}\left(p'_i - \frac{1}{2}\right)\right) + \frac{1}{2} & \text{otherwise} \end{cases} \quad \text{for } 1 \leq i \leq r.$$

Figure 3.3 shows an instance of this construction. The marked band system of length $2n' = 4$ on the top is defined by $B' = \{(1, 4), (2, 3)\}$ and $p' = (1 + 1/2)$, and is clearly maximal. We choose the unbundling map φ such that $\varphi(1) = \varphi(4) = 4$ and $\varphi(2) = \varphi(3) = 2$. If we replace each band $(a, b) \in B'$ with $\varphi(a)$ parallel bands, we obtain the marked band system (B, p) displayed on the bottom, where

$$B = \{(1, 12), (2, 11), (3, 10), (4, 9), (5, 8), (6, 7)\} \text{ of length } 2n = 12,$$

$$p = \left(4 + \frac{1}{2}\right).$$




Figure 3.3. A maximal marked band system (on the top) and the marked band system associated to an unbundling map φ (on the bottom).

Given an unbundling map φ for (B', p') , it is sometimes convenient to define the corresponding *cumulative unbundling map* $\widehat{\varphi}: \{1, \dots, 2n' + 1\} \rightarrow \mathbb{Z}_{\geq 0}$ as

$$\widehat{\varphi}(a') = \varphi(1) + \dots + \varphi(a' - 1).$$

With this definition, we can compactly write $\iota^{-1}(a')$ as $\{\widehat{\varphi}(a') + 1, \dots, \widehat{\varphi}(a' + 1)\}$, without having to refer to the band system (B, p) associated to φ .

3.2 Band systems and free groups

Let us introduce some notation about free groups which we will employ in this section. Let g be a positive integer. Fix g symbols $\{x_1, \dots, x_g\}$ and denote by \mathbb{F}_g the free group over them. If s is a word of length m in the symbols x_1, \dots, x_g and their inverses, we use the following notations:

- s_i is the i -th symbol of s , where $1 \leq i \leq m$;
- $s_{[i:j]}$ is the subword $s_i s_{i+1} \dots s_j$, where $1 \leq i \leq j \leq m$;
- s^{-1} is the word such that $(s^{-1})_i = (s_{m+1-i})^{-1}$ for $1 \leq i \leq m$;
- s^k is the concatenation of k copies of s , where $k \geq 0$; if $k < 0$, we define $s^k = (s^{-1})^{-k}$;
- s^∞ the infinite word obtained by concatenating infinitely many copies of s ; this does not correspond to an element of \mathbb{F}_g , and we will only ever make use of finite subwords of infinite words.

Let $s = s_1 s_2 \dots s_{2n}$ be a word which reduces to 1 in \mathbb{F}_g . Then it must do so via a sequence of cancellations of adjacent symbols of the form xx^{-1} or $x^{-1}x$ (by “adjacent” we mean “adjacent after performing the previous cancellations in the sequence”). For any such sequence, we can define a band system

$$B = \{(a, b) : 1 \leq a < b \leq 2n, s_a \text{ cancels with } s_b\}.$$

We say that B is a *cancellation band system* for s ; cancellation band systems need not be unique, but each word reducing to 1 has at least one of them. For example, Figure 3.4 shows how the band system

$$B = \{(1, 6), (2, 3), (4, 5), (7, 10), (8, 9)\}$$

is a cancellation band system for the word $yy^{-1}yx^{-1}xy^{-1}xyy^{-1}x^{-1}$ in the free group generated by the symbols x and y .




Figure 3.4. A cancellation band system for the word $yy^{-1}yx^{-1}xy^{-1}xyy^{-1}x^{-1}$.

The main technical result of this section is the following.

Theorem 3.2. Let $w_0, \dots, w_l, t_1, \dots, t_l \in \mathbb{F}_g$. Consider the set

$$\mathcal{A} = \left\{ (k_1, \dots, k_l) \in \mathbb{Z}^l : w_0 t_1^{k_1} w_1 t_2^{k_2} w_2 \cdots t_l^{k_l} w_l = 1 \right\}.$$

There is an algorithm which, given as input w_0, \dots, w_l and t_1, \dots, t_l , returns a collection of sets A_1, \dots, A_N such that $\mathcal{A} = A_1 \cup \dots \cup A_N$, and each A_j is of the form

$$A_j = \left\{ \mathbf{z}_j + \mathbf{M}_j \mathbf{v} : \mathbf{v} \in \mathbb{Z}_{\geq 0}^{d_j} \right\}$$

for some vector $\mathbf{z}_j \in \mathbb{Z}^l$ and some matrix $\mathbf{M}_j \in \mathbb{Z}^{l \times d_j}$.

Proof. Let w_0, \dots, w_l and t_1, \dots, t_l be given as reduced words over $\{x_1, \dots, x_g\}$. For a vector $\mathbf{k} \in \mathbb{Z}^l$, denote by $w(\mathbf{k})$ the word $w_0 t_1^{k_1} w_1 \cdots t_l^{k_l} w_l$.

✿ **Cyclically reduced words.** For each $1 \leq i \leq l$ we can write $t_i = u_i t'_i u_i^{-1}$, where t'_i is a cyclically reduced word and u_i is a reduced word. Define:

- $w'_0 = w_0 u_1$;
- $w'_i = u_i^{-1} w_i u_{i+1}$ for $1 \leq i < l$;
- $w'_l = u_l^{-1} w_l$.

Then

$$\mathcal{A} = \left\{ (k_1, \dots, k_l) \in \mathbb{Z}^l : w'_0 (t'_1)^{k_1} w'_1 \cdots (t'_l)^{k_l} w'_l = 1 \right\}.$$

Therefore, up to replacing each w_i with w'_i and each t_i with t'_i , we can assume that each t_i is cyclically reduced.

✿ **Only non-negative powers.** For each $\epsilon = (\epsilon_1, \dots, \epsilon_l) \in \{-1, 1\}^l$ define

$$\mathcal{A}^\epsilon = \left\{ (\epsilon_1 k_1, \dots, \epsilon_l k_l) : \begin{array}{l} (k_1, \dots, k_l) \in \mathbb{Z}_{\geq 0}^l, \\ w_0 t_1^{\epsilon_1 k_1} w_1 \cdots t_l^{\epsilon_l k_l} w_l = 1 \end{array} \right\},$$

so that

$$\mathcal{A} = \bigcup_{\epsilon \in \{-1, 1\}^l} \mathcal{A}^\epsilon.$$

Denote by \mathcal{A}^+ the set

$$\mathcal{A}^{(1, \dots, 1)} = \left\{ (k_1, \dots, k_l) \in \mathbb{Z}_{\geq 0}^l : w_0 t_1^{k_1} w_1 \cdots t_l^{k_l} w_l = 1 \right\}.$$

Suppose we have an algorithm to decompose \mathcal{A}^+ as described in the statement, and fix $\epsilon = (\epsilon_1, \dots, \epsilon_l) \in \{-1, 1\}^l$. Then we can algorithmically decompose

$$\left\{ (k_1, \dots, k_l) \in \mathbb{Z}_{\geq 0}^l : w_0 (t_1^{\epsilon_1})^{k_1} w_1 \cdots (t_l^{\epsilon_l})^{k_l} w_l = 1 \right\} = A_1 \cup \dots \cup A_N,$$

where each A_j is of the form

$$A_j = \left\{ \mathbf{z}_j + \mathbf{M}_j \mathbf{v} : \mathbf{v} \in \mathbb{Z}_{\geq 0}^{d_j} \right\}$$

for some vector $\mathbf{z}_j \in \mathbb{Z}^l$ and some matrix $\mathbf{M}_j \in \mathbb{Z}^{l \times d_j}$. But then we have the decomposition

$$\mathcal{A}^\epsilon = A'_1 \cup \dots \cup A'_N,$$

where, for each $1 \leq j \leq N$:

- $A_j = \{ \mathbf{z}'_j + \mathbf{M}'_j \mathbf{v} : \mathbf{v} \in \mathbb{Z}_{\geq 0}^{d_j} \};$
- \mathbf{z}'_j is obtained from \mathbf{z}_j by multiplying the i -th coordinate by ϵ_i for $1 \leq i \leq l$;
- \mathbf{M}'_j is obtained from \mathbf{M}_j by multiplying the i -th row by ϵ_i for $1 \leq i \leq l$.

Since there are only finitely many choices of $\epsilon \in \{-1, 1\}^l$, it is enough to describe an algorithm to compute the set \mathcal{A}^+ .

• **Non-empty words.** Suppose we have an algorithm which works when every t_i is a non-empty word (that is, $t_i \neq 1$ as an element of \mathbb{F}_g). If t_i is the empty word for some $1 \leq i \leq l$, by induction on l we can assume that the set

$$\left\{ (k_1, \dots, k_{i-1}, k_{i+1}, \dots, k_l) \in \mathbb{Z}_{\geq 0}^{l-1} : w_0 t_1^{k_1} w_1 \cdots t_{i-1}^{k_{i-1}} (w_{i-1} w_i) t_{i+1}^{k_{i+1}} w_{i+1} \cdots t_l^{k_l} w_l = 1 \right\}$$

can be algorithmically decomposed as $A_1 \cup \dots \cup A_N$, where each A_j is of the form

$$A_j = \left\{ \mathbf{z}_j + \mathbf{M}_j \mathbf{v} : \mathbf{v} \in \mathbb{Z}_{\geq 0}^{d_j} \right\}$$

for some vector $\mathbf{z}_j \in \mathbb{Z}^{l-1}$ and some matrix $\mathbf{M}_j \in \mathbb{Z}^{(l-1) \times d_j}$. For each $1 \leq j \leq N$, write

$$\mathbf{z}_j = \begin{pmatrix} \mathbf{z}_j^{(1)} \\ \mathbf{z}_j^{(2)} \end{pmatrix} \quad \text{and} \quad \mathbf{M}_j = \begin{pmatrix} \mathbf{M}_j^{(1)} \\ \mathbf{M}_j^{(2)} \end{pmatrix},$$

where $\mathbf{z}_j^{(1)} \in \mathbb{Z}^{i-1}$, $\mathbf{z}_j^{(2)} \in \mathbb{Z}^{l-i}$, and similarly $\mathbf{M}_j^{(1)} \in \mathbb{Z}^{(i-1) \times d_j}$, $\mathbf{M}_j^{(2)} \in \mathbb{Z}^{(l-i) \times d_j}$. Define

$$\mathbf{z}'_j = \begin{pmatrix} \mathbf{z}_j^{(1)} \\ 0 \\ \mathbf{z}_j^{(2)} \end{pmatrix} \in \mathbb{Z}^l \quad \text{and} \quad \mathbf{M}'_j = \begin{pmatrix} \mathbf{M}_j^{(1)} & \mathbf{0} \\ \mathbf{0} & 1 \\ \mathbf{M}_j^{(2)} & \mathbf{0} \end{pmatrix} \in \mathbb{Z}^{l \times (d_j+1)}.$$

We get the required decomposition

$$\mathcal{A}^+ = \bigcup_{j=1}^N \left\{ \mathbf{z}'_j + \mathbf{M}'_j \mathbf{v} : \mathbf{v} \in \mathbb{Z}_{\geq 0}^{d_j+1} \right\}.$$

Therefore, we can assume that each t_i is non-empty.

• **Fixed cancellation bundle.** Let $\mathbf{k} = (k_1, \dots, k_l)$ be a vector in \mathcal{A}^+ . Consider a cancellation band system B for the word $w(\mathbf{k})$. Moreover, define the *block marking* $p = (p_1, \dots, p_{2l+2})$ to separate different “blocks” of $w(\mathbf{k})$, by setting

$$\begin{aligned} p_{2i-1} &= |w_0| + k_1 |t_1| + |w_1| + \dots + k_{i-1} |t_{i-1}| + \frac{1}{2} & \text{for } 1 \leq i \leq l+1, \\ p_{2i} &= |w_0| + k_1 |t_1| + |w_1| + \dots + |w_{i-1}| + \frac{1}{2} & \text{for } 1 \leq i \leq l+1. \end{aligned}$$

Denote by (B', p') the maximal bundle of the marked band system (B, p) ; we say that (B', p') is a *cancellation bundle* for \mathbf{k} if it can be obtained in this way (that is, as the

maximal bundle of a cancellation band system for $w(\mathbf{k})$ endowed with the block marking described above). Let

$$\iota: \{1, \dots, 2n\} \longrightarrow \{1, \dots, 2n'\}$$

be the bundling map, where $2n$ and $2n'$ are the lengths of B and B' respectively. Let us label the integers $\{1, \dots, 2n'\}$ with the symbols $W_0, \dots, W_l, T_1, \dots, T_l$ depending on which “block” of $w(\mathbf{k})$ they come from. More precisely, we label

- the integers between p'_{2i+1} and p'_{2i+2} with the symbol W_i for $0 \leq i \leq l$, and
- the integers between p'_{2i} and p'_{2i+1} with the symbol T_i for $1 \leq i \leq l$.

We will use the notation $i \prec X$ to signify that the integer $i \in \{1, \dots, 2n'\}$ has label $X \in \{W_0, \dots, W_l, T_1, \dots, T_l\}$. See Figure 3.5 for a graphical representation of a cancellation bundle.

The crucial fact we seek to prove now is that $2n'$ is bounded above by a constant which does not depend on \mathbf{k} . This is a consequence of the following two statements.

- 1 For each $1 \leq i \leq l$, there is no band of B' with both endpoints labelled T_i . In fact, suppose that $(a', b') \in B'$ is such a band, and let $(a, b) \in B$ with $\iota(a, b) = (a', b')$. By statement 1 of Lemma 3.1, there is a band $(c, c+1) \in B$ with $a \leq c < b$. But this is impossible, since positions a and b (and therefore c and $c+1$) belong to the block $t_i^{k_i}$; the word t_i is cyclically reduced, therefore there can be no cancellation between the symbols in positions c and $c+1$.
- 2 For each $1 \leq i < j \leq l$, there is at most one band of $(a', b') \in B'$ with $a' \prec T_i$ and $b' \prec T_j$. In fact, suppose that (a'_1, b'_1) and (a'_2, b'_2) are two such bands, with $a'_1 < a'_2 < b'_2 < b'_1$. Let (a_1, b_1) and (a_2, b_2) be bands of B with $\iota(a_1, b_1) = (a'_1, b'_1)$ and $\iota(a_2, b_2) = (a'_2, b'_2)$. Consider statement 2 of Lemma 3.1. Clearly,

$$\iota(a_1, b_1) = (a'_1, b'_1) \neq (a'_2, b'_2) = \iota(a_2, b_2).$$

Moreover, since $a'_1, a'_2 \prec T_i$ and $b'_1, b'_2 \prec T_j$, we have that

$$p_{2i} < a_1 < a_2 < p_{2i+1} \leq p_{2j} < b_2 < b_1 < p_{2j+1}.$$

As a consequence, there must be a band $(c, c+1) \in B$ for some

$$c \in \{a_1, \dots, a_2 - 1\} \cup \{b_2, \dots, b_1 - 1\},$$

but this is once again impossible since t_i and t_j are cyclically reduced words.

As anticipated, these two facts imply the existence of an upper bound C for the length $2n'$ of B' which is independent of \mathbf{k} . An estimate for C can be found as follows. Say that a band $(a', b') \in B'$ is:

- of *type 1* if $a' \prec W_i$ or $b' \prec W_i$ for some $0 \leq i \leq l$;
- of *type 2* if it is not of type 1, that is if $a' \prec T_i$ and $b' \prec T_j$ for some $1 \leq i, j \leq l$.

An easy counting argument shows that there are at most $|w_0| + \dots + |w_l|$ bands of type 1 and at most $l(l-1)/2$ bands of type 2. Therefore, a suitable upper bound for $2n'$ is

$$2n' \leq C = 2(|w_0| + \dots + |w_l|) + l(l-1).$$

Let us now forget our choice of \mathbf{k} , and fix any maximal marked band system (B', p') such that B' has length $2n' \leq C$ and $p' = (p'_1, \dots, p'_{2l+2})$ with $p'_1 = 1/2$ and $p'_{2l+2} =$




Figure 3.5. The cancellation bundle (on the bottom) induced by a cancellation band system (on the top). In this specific example, we have that $|w_0| = 2$, $|t_1| = 2$, $|w_1| = 5$, $|t_2| = 4$, and $|w_2| = 1$. The word $w(\mathbf{k})$ on the top corresponds to the vector $\mathbf{k} = (2, 1)$, and reduces to 1 via the cancellation band system represented in the picture – although this is of course impossible to verify without knowing the actual the words involved. The marked spots separate the different “blocks” of $w(\mathbf{k})$, namely w_0 , t_1^2 , w_1 , t_2 , and w_2 . The maximal bundling procedure yields the marked band system depicted on the bottom, with vertical arrows representing the bundling map. Here, the integers in $\{1, \dots, 10\}$ are labelled with W_0 , T_1 , W_1 , T_2 , or W_2 according to the block they come from.

$2n' + 1/2$. Since there are only finitely many such marked band systems, and since every $\mathbf{k} \in \mathcal{A}^+$ has a cancellation bundle of this form, it is enough to describe an algorithm to compute the set

$$\left\{ \mathbf{k} = (k_1, \dots, k_l) \in \mathbb{Z}_{\geq 0}^l : w_0 t_1^{k_1} w_1 t_2^{k_2} w_2 \cdots t_l^{k_l} w_l = 1, \right. \\ \left. (B', p') \text{ is a cancellation bundle for } \mathbf{k} \right\},$$

which will be henceforth referred to as $\mathcal{A}_{(B', p')}^+$.

Let us remark that the labelling procedure described above can be carried out without any reference to \mathbf{k} ; in fact, only the integer $2n'$ and the marking p' are required. Therefore, we will from now on assume that the integers in $\{1, \dots, 2n'\}$ are labelled with symbols $W_0, \dots, W_l, T_1, \dots, T_l$ according to the marking p' .

✿ Fixed unbundling class. Let us dwell some more on what it means for a vector $\mathbf{k} \in \mathbb{Z}_{\geq 0}^l$ to have (B', p') as a cancellation bundle. Recall that band systems having (B', p') as their maximal bundle are parametrised by unbundling maps for (B', p') . It is easy to see that an unbundling map φ for (B', p') describes a cancellation band system for \mathbf{k} if and only if the following conditions are satisfied:

i $\sum_{a' \prec W_i} \varphi(a') = |w_i|$ for $0 \leq i \leq l$;

ii $\sum_{a' \prec T_i} \varphi(a') = k_i |t_i|$ for $1 \leq i \leq l$;

iii $w(\mathbf{k})_{[\widehat{\varphi}(a')+1: \widehat{\varphi}(a'+1)]} = (w(\mathbf{k})_{[\widehat{\varphi}(b')+1: \widehat{\varphi}(b'+1)]})^{-1}$ for each $(a', b') \in B'$.

Moreover, condition **ii** implies that φ describes a cancellation band system for at most one vector $\mathbf{k}(\varphi) \in \mathbb{Z}_{\geq 0}^l$, whose coordinates can be easily computed as

$$k_i(\varphi) = \frac{\sum_{a' \prec T_i} \varphi(a')}{|t_i|} \quad \text{for } 1 \leq i \leq l \quad (2)$$

(recall that $|t_i| \neq 0$ for each i).

Let us say that an unbundling map φ for (B', p') is *compatible with the labelling* if

$$\sum_{a' \prec W_i} \varphi(a') = |w_i| \text{ for } 0 \leq i \leq l \quad \text{and} \quad \sum_{a' \prec T_i} \varphi(a') \equiv 0 \pmod{|t_i|} \text{ for } 1 \leq i \leq l.$$

If, additionally, we have

$$w(\mathbf{k}(\varphi))_{[\widehat{\varphi}(a')+1: \widehat{\varphi}(a'+1)]} = (w(\mathbf{k}(\varphi))_{[\widehat{\varphi}(b')+1: \widehat{\varphi}(b'+1)]})^{-1} \quad \text{for each } (a', b') \in B',$$

then we say that φ is *cancelling*. We can then write the following equivalent definition for $\mathcal{A}_{(B', p')}^+$:

$$\mathcal{A}_{(B', p')}^+ = \left\{ \mathbf{k}(\varphi) : \begin{array}{l} \varphi \text{ is an unbundling map for } (B', p') \text{ which is} \\ \text{compatible with the labelling and cancelling} \end{array} \right\}.$$

Let us define an equivalence relation on the set of unbundling maps for (B', p') which are compatible with the labelling. Given two unbundling maps φ and ψ , we say that $\varphi \sim \psi$ if the following properties hold for each $1 \leq a' \leq 2n'$:

✿ $\varphi(a') = \psi(a')$ whenever a' is an endpoint of a band of type 1;

- $\varphi(a') \equiv \psi(a') \pmod{|t_i|}$ whenever $a' \prec T_i$ for some $1 \leq i \leq l$.

Fix an equivalence class Ξ . After choosing a representative $\varphi \in \Xi$, with slight abuse of notation, for each $1 \leq a' \leq 2n'$ let us denote by $\Xi(a')$:

- the integer $\varphi(a')$ if $a' \prec W_i$ for some $0 \leq i \leq l$;
- the residue class of $\varphi(a')$ modulo $|t_i|$ if $a' \prec T_i$ for some $1 \leq i \leq l$.

In both cases, the value of $\Xi(a')$ does not depend on the choice of φ . The values of $\Xi(a')$ for $1 \leq a' \leq 2n'$ uniquely identify Ξ , hence there are only finitely many equivalence classes of unbundling maps. As a consequence, it is enough to describe an algorithm to compute the set

$$\{\mathbf{k}(\varphi) : \varphi \in \Xi \text{ is cancelling}\},$$

which will be henceforth referred to as $\mathcal{A}_{(B', p'), \Xi}^+$.

• **Final computation.** The reason why fixing the unbundling class Ξ is beneficial is that conditions of the form

$$w(\mathbf{k}(\varphi))_{[\widehat{\varphi}(a')+1 : \widehat{\varphi}(a'+1)]} = \left(w(\mathbf{k}(\varphi))_{[\widehat{\varphi}(b')+1 : \widehat{\varphi}(b'+1)]} \right)^{-1} \quad (3)$$

for bands $(a', b') \in B'$ become mutually independent. More precisely, whether condition (3) is satisfied for a band (a', b') or not only depends on the value of $\varphi(a')$, as long as $\varphi \in \Xi$; we will now clarify and justify this claim.

Let $a' \in \{1, \dots, 2n'\}$ be an integer, and let $X \in \{W_0, \dots, W_l, T_1, \dots, T_l\}$ be its label. Define

$$h(a') = \sum_{\substack{b' \prec X \\ b' < a'}} \Xi(b').$$

Note that, when $X = T_i$ for some $1 \leq i \leq l$, the integer $h(a')$ is only well-defined modulo $|t_i|$; by convention, we will assume that $0 \leq h(a') < |t_i|$. We can now rewrite the subwords involved in (3): for each $\varphi \in \Xi$ we have

$$w(\mathbf{k}(\varphi))_{[\widehat{\varphi}(a')+1 : \widehat{\varphi}(a'+1)]} = \begin{cases} (w_i)_{[h(a')+1 : h(a')+\varphi(a')]} & \text{if } a' \prec W_i, \\ (t_i^\infty)_{[h(a')+1 : h(a')+\varphi(a')]} & \text{if } a' \prec T_i. \end{cases}$$

Consider a band $(a', b') \in B'$. It is now clear that whether (a', b') satisfies (3) or not only depends on the value of $\varphi(a') = \varphi(b')$. Our next task is to show how to compute the set

$$S(a', b') = \{s \in \mathbb{Z}_{>0} : (a', b') \text{ satisfies (3) if and only if } \varphi(a') = s\},$$

and we do so by analysing two cases.

• Suppose that (a', b') is a band of type 1, with $a' \prec W_i$ for some $0 \leq i \leq l$ (the case where $b' \prec W_i$ is identical). Then $\varphi(a') = \Xi(a')$ is fixed, and condition (3) reads

$$\begin{aligned} (w_i)_{[h(a')+1 : h(a')+\Xi(a')]} &= \left((w_j)_{[h(b')+1 : h(b')+\Xi(a')]} \right)^{-1} && \text{if } b' \prec W_j, \text{ or} \\ (w_i)_{[h(a')+1 : h(a')+\Xi(a')]} &= \left((t_j^\infty)_{[h(b')+1 : h(b')+\Xi(a')]} \right)^{-1} && \text{if } b' \prec T_j. \end{aligned}$$

Either way, whether the equality of words holds or not only depends on Ξ . We have that $S(a', b') = \{\Xi(a')\}$ if the two words are equal, and $S(a', b') = \emptyset$ if they are not.

✿ Suppose that (a', b') is a band of type 2, with $a' \prec T_i$ and $b' \prec T_j$ for some $1 \leq i < j \leq l$. Then all the possible values for $\varphi(a') = \varphi(b')$ are of the form $q + Lv$, where

- $L = \text{lcm}(|t_i|, |t_j|)$,
- v is a non-negative integer, and
- q is the only integer in $\{1, \dots, L\}$ such that $q \equiv \Xi(a') \pmod{|t_i|}$ and $q \equiv \Xi(b') \pmod{|t_j|}$;

if no such q exists then the class Ξ is empty, and we conclude that $\mathcal{A}_{(B', p'), \Xi}^+ = \emptyset$. Otherwise, all that is left to do is determining for which values of $v \geq 0$ the band (a', b') satisfies (3). We can rewrite

$$(t_i^\infty)_{[h(a')+1: h(a')+\varphi(a')]} = ((t_i^\infty)_{[h(a')+1: h(a')+L]})^v (t_i^\infty)_{[h(a')+1: h(a')+q]},$$

and similarly for $(t_j^\infty)_{[h(b')+1: h(b')+\varphi(b')]}$. Therefore, condition (3) reads

$$\begin{aligned} ((t_i^\infty)_{[h(a')+1: h(a')+L]})^v (t_i^\infty)_{[h(a')+1: h(a')+q]} = \\ ((t_j^\infty)_{[h(b')+q+1: h(b')+q+L]})^{-v} ((t_j^\infty)_{[h(b')+1: h(b')+q]})^{-1}. \end{aligned}$$

There are three cases.

- If

$$(t_i^\infty)_{[h(a')+1: h(a')+q]} \neq ((t_j^\infty)_{[h(b')+1: h(b')+q]})^{-1},$$

then (3) is not satisfied for any value of v , and $S(a', b') = \emptyset$.

- If

$$(t_i^\infty)_{[h(a')+1: h(a')+q]} = ((t_j^\infty)_{[h(b')+1: h(b')+q]})^{-1}$$

but

$$(t_i^\infty)_{[h(a')+1: h(a')+L]} \neq ((t_j^\infty)_{[h(b')+q+1: h(b')+q+L]})^{-1},$$

then (3) is satisfied only for $v = 0$, leading to $S(a', b') = \{q\}$.

- If

$$(t_i^\infty)_{[h(a')+1: h(a')+L]} = ((t_j^\infty)_{[h(b')+q+1: h(b')+q+L]})^{-1},$$

then (3) is satisfied for every value of $v \geq 0$, so $S(a', b') = \{q + Lv : v \geq 0\}$.

In conclusion, we have shown how to compute a set $S(a', b')$ for each band $(a', b') \in B'$ with the following property: a function $\varphi: \{1, \dots, 2n'\} \rightarrow \mathbb{Z}_{>0}$ is a cancelling representative of Ξ if and only if $\varphi(a') = \varphi(b') \in S(a', b')$ for every $(a', b') \in B'$. If $S(a', b') = \emptyset$ for some band (a', b') then clearly $\mathcal{A}_{(B', p'), \Xi}^+ = \emptyset$. Otherwise, each set $S(a', b')$ can be written in the form

$$S(a', b') = \{q(a', b') + L(a', b')v : v \in \mathbb{Z}_{\geq 0}\}$$

for some integers $q(a', b') > 0$ and $L(a', b') \geq 0$. Therefore, every cancelling unbanding map $\varphi \in \Xi$ is described by a (not necessarily unique) vector $\mathbf{v} \in \mathbb{Z}_{\geq 0}^{n'}$ whose coordinates $v_{(a', b')}$ are indexed by bands $(a', b') \in B'$, such that

$$\varphi(a') = \varphi(b') = q(a', b') + L(a', b')v_{(a', b')} \quad \text{for each } (a', b') \in B'.$$

Proof. Fix a basepoint $x_0 \in \partial V$ once and for all; additionally, let us arbitrarily pick orientations for a_1, \dots, a_{2m} . Consider a complete system of oriented meridians $b_1, \dots, b_g \subseteq \partial V$. For each b_r , we give a few definitions. Let $f(b_r)$ intersect $a_1 \cup \dots \cup a_{2m}$ transversely at points $p_{r,1}, \dots, p_{r,l_r}$, numbered in the order they appear on $f(b_r)$. For $j = 1, \dots, l_r$, let (when $j = l_r$, we take $j + 1$ to mean 1):

- $\beta_{r,j}$ be the subarc of $f(b_r)$ going from $p_{r,j}$ to $p_{r,j+1}$;
- $\gamma_{r,j} \subseteq \partial V$ be an arc joining x_0 to $p_{r,j}$;
- $i(r,j)$ be the only integer in $\{1, \dots, 2m\}$ such that $p_{r,j} \in a_{i(r,j)}$;
- $w_{r,j}$ be the element of $\pi_1(V, x_0)$ represented by $\gamma_{r,j} * \beta_{r,j} * \gamma_{r,j+1}^{-1}$, where “ $*$ ” denotes the concatenation of paths;
- $t_{r,j}$ be the element of $\pi_1(V, x_0)$ represented by $\gamma_{r,j} * a_{i(r,j)}^\epsilon * \gamma_{r,j}^{-1}$, where we consider $a_{i(r,j)}$ as a loop based at $p_{r,j}$, and ϵ is 1 or -1 depending on the “sign” of the intersection between $f(b_r)$ and $a_{i(r,j)}$ at $p_{r,j}$, as depicted in Figure 3.6.




Figure 3.6. Choice of ϵ at the intersection $p_{r,j}$ between $f(b_r)$ and $a_{i(r,j)}$.

It is not hard to see that, given integers h_1, \dots, h_{2m} , the conjugacy class of the curve $\tau_{a_1}^{h_1} \dots \tau_{a_{2m}}^{h_{2m}} f(b_r)$ in $\pi_1(V, x_0)$ is represented by the element

$$t_{r,1}^{h_{i(r,1)}} w_{r,1} \dots t_{r,l_r}^{h_{i(r,l_r)}} w_{r,l_r};$$

we hope that Figure 3.7 will be convincing enough for the reader. Consider the set

$$\mathcal{A}_r = \left\{ (k_1, \dots, k_{l_r}) \in \mathbb{Z}^{l_r} : t_{r,1}^{k_1} w_{r,1} \dots t_{r,l_r}^{k_{l_r}} w_{r,l_r} = 1 \right\}.$$

By the previous remark, we have that $\tau_{a_1}^{h_1} \dots \tau_{a_{2m}}^{h_{2m}} f(b_r)$ is a meridian if and only if $(h_{i(r,j)})_{j=1}^{l_r} \in \mathcal{A}_r$. Consequently, $\tau_{a_1}^{h_1} \dots \tau_{a_{2m}}^{h_{2m}} f$ extends to a homeomorphism of V if and only if $(h_{i(r,j)})_{j=1}^{l_r} \in \mathcal{A}_1 \cap \dots \cap \mathcal{A}_g$.

By Theorem 3.2, the sets \mathcal{A}_r can be algorithmically decomposed as

$$\mathcal{A}_r = A_{r,1} \cup \dots \cup A_{r,N_r},$$

where

$$A_{r,q} = \left\{ \mathbf{z}_{r,q} + \mathbf{M}_{r,q} \mathbf{v} : \mathbf{v} \in \mathbb{Z}_{\geq 0}^{d_{r,q}} \right\}$$




Figure 3.7. The effect of applying a Dehn twist about a to the curve $f(b)$; indices were omitted for clarity.

for some vector $\mathbf{z}_{r,q} \in \mathbb{Z}^{l_r}$ and some matrix $\mathbf{M}_{r,q} \in \mathbb{Z}^{l_r \times d_{r,q}}$. Therefore, we finally find that there exist integers h_1, \dots, h_m such that

$$\tau_{a_1}^{h_1} \dots \tau_{a_m}^{h_m} \tau_{a_{m+1}}^{-h_1} \dots \tau_{a_{2m}}^{-h_m} f$$

extends to a homeomorphism of V if and only if, for some choice of indices $q_1 \in \{1, \dots, N_1\}, \dots, q_g \in \{1, \dots, N_g\}$, the following system of linear equations has at least one solution in the variables $\mathbf{v}_1 \in \mathbb{Z}_{\geq 0}^{d_{1,q_1}}, \dots, \mathbf{v}_g \in \mathbb{Z}_{\geq 0}^{d_{g,q_g}}, h_1, \dots, h_m \in \mathbb{Z}$:

$$\begin{cases} \mathbf{z}_{r,q_r} + \mathbf{M}_{r,q_r} \mathbf{v}_r = (h_{i(r,j)})_{j=1}^{l_r} & \text{for } r = 1, \dots, g, \\ h_r = -h_{m+r} & \text{for } r = 1, \dots, m. \end{cases}$$

This condition can be checked algorithmically (a classical result proved in [6] by Gathen and Sieveking) for every choice of indices q_1, \dots, q_g , providing an algorithmic solution to the question in the statement. \square

4 The classification algorithm

4.1 Outline

Section 4 will be entirely devoted to providing a proof of the following.

Theorem 4.1. *There is an algorithm to decide whether two genus-two oriented surfaces embedded in S^3 are isotopic or not.*

More precisely, let S_1 and S_2 be genus-two surfaces embedded in S^3 , and let us fix an orientation for S_1 and for S_2 . We say that S_1 and S_2 are isotopic (as oriented surfaces) if there is an isotopy $f_t: S^3 \rightarrow S^3$ such that f_0 is the identity and f_1 maps S_1 to S_2 orientation-preservingly. Since every self-homeomorphism of S^3 is isotopic to the identity, the oriented surfaces S_1 and S_2 are isotopic if and only if there exists a homeomorphism $f: S^3 \rightarrow S^3$ sending S_1 to S_2 orientation-preservingly.

Another equivalent definition is the following. Let M_1 be the closure of the component of $S^3 \setminus S_1$ which lies on the positive side of S_1 , and let N_1 be the closure of the other component; we will sometimes call M_1 and N_1 the *sides* of S_1 . Define M_2 and N_2 in a similar way. Then S_1 and S_2 are isotopic if and only if there exist homeomorphisms $f: M_1 \rightarrow M_2$ and $g: N_1 \rightarrow N_2$ such that $f|_{\partial}$ and $g|_{\partial}$ are isotopic as homeomorphisms $S_1 \rightarrow S_2$.

Remark 4.2. A non-oriented version of Theorem 4.1 states that there is an algorithm to decide whether two genus-two surfaces embedded in S^3 are isotopic or not. Of course, the non-oriented version is an immediate corollary of the oriented one. In fact, in order to decide whether the surfaces S_1 and S_2 are isotopic without any constraint on orientation, we can simply check if they are orientation-preservingly isotopic for at least one choice of orientations.

From now on, we will take $S_1, S_2, M_1, N_1, M_2,$ and N_2 to refer to the objects we have already introduced; in particular, we will assume orientations for S_1 and S_2 have been fixed. Our aim will be to provide a proof of Theorem 4.1. The full algorithm is quite involved, and for the sake of convenience we will split it into several cases. Before we start describing the actual algorithm, let us briefly outline the structure of the following sections.

- ⚙️ Several cases of the proof of Theorem 4.1 will be addressed by finding *canonical* compression discs for S_1 and S_2 , then compressing S_1 and S_2 along these discs and – carefully – reducing the question to the isotopy problem of tori in S^3 . In Section 4.2 we provide a general framework for dealing with these canonical compression discs; by abstracting the repetitive parts of the algorithm away, we will later be able to focus on what is meaningfully different in each of these cases.
- ⚙️ In Section 4.3 we deal with the problem of finding such canonical compression discs. The recipe we present relies on showing the existence of discs satisfying properties which are *stable under boundary compressions*; if these properties are suitably chosen, uniqueness will easily follow.
- ⚙️ After the introductory sections, we finally start describing the classification algorithm. Sections 4.4 and 4.5 deal with the case where S_1 and S_2 are *partially compressible on one side*. Specifically, the case where one side has at least one – and, *a posteriori*, only one – non-separating compression disc is addressed in Section 4.4; Section 4.5 explains the strategy to follow when one side only has separating compression discs. In both cases, we heavily rely on the tools developed in Sections 4.2 and 4.3.
- ⚙️ We are then left to address the case where one side of S_1 is a handlebody and the other is boundary irreducible. The algorithm we use here will depend on the JSJ decomposition of the boundary irreducible component. If all the I -bundle pieces are “small”, then the algorithmic problem on free groups discussed in Section 3 will quickly lead to a solution of the isotopy problem, as described in Section 4.6. Otherwise, only two cases can arise: either a product I -bundle over a punctured torus appears in the JSJ decomposition – and we deal with this in Section 4.7 – or one of the pieces is a twisted I -bundle over a punctured Klein bottle; the latter situation is addressed in Section 4.8. In both cases, an *ad hoc* discussion is enough to settle the isotopy problem.

4.2 General strategy for canonical compression discs

We say that non-trivial compression discs D_1 for S_1 and D_2 for S_2 are *canonical* if every self-homeomorphism of S^3 sending S_1 to S_2 orientation-preservingly sends D_1 to D_2 , up to isotopy preserving S_2 . As anticipated, we will often rely on compressing the surfaces along canonical compression discs in order to reduce the isotopy question to the more manageable problem for tori in S^3 . Let us begin our discussion with a definition.

Definition 4.1. Let X_1 and X_2 be 3-manifolds, and let $p_1 \subseteq \partial X_1$ and $p_2 \subseteq \partial X_2$ be simple closed curves. Let X'_1 and X'_2 be the 3-manifolds obtained by attaching a 2-handle,

respectively, to X_1 along p_1 and to X_2 along p_2 . Define the *handle extension map*

$$\text{HExt}_{p_1, p_2} : \mathbf{H}((X_1, p_1); (X_2, p_2)) \longrightarrow \mathbf{H}(X'_1; X'_2)$$

by extending homeomorphisms $(X_1, p_1) \rightarrow (X_2, p_2)$ to the 2-handles.

It is easy to see that this map is well-defined, in that its output does not depend on the choice of representative for the isotopy class nor on the specific extension to the 2-handles. Moreover, the handle extension map is functorial in the following sense: for 3-manifolds X_1, X_2, X_3 , curves $p_1 \subseteq \partial X_1, p_2 \subseteq \partial X_2, p_3 \subseteq \partial X_3$, and homeomorphisms

$$f : (X_1, p_1) \longrightarrow (X_2, p_2), \quad g : (X_2, p_2) \longrightarrow (X_3, p_3),$$

we have the equality

$$\text{HExt}_{p_1, p_3}(gf) = \text{HExt}_{p_2, p_3}(g) \circ \text{HExt}_{p_1, p_2}(f).$$

Suppose now that we can find canonical compression discs D_1 and D_2 for S_1 and S_2 respectively. Up to flipping the orientations of S_1 and S_2 , we can assume that D_1 and D_2 lie in M_1 and M_2 respectively. Let $p_1 \subseteq S_1$ and $p_2 \subseteq S_2$ be the boundary curves of D_1 and D_2 respectively. Let $P_1 = M_1 \setminus D_1$ (which may be a disconnected 3-manifold) and $Q_1 = \text{clos}(S^3 \setminus P_1)$. Note that Q_1 can also be obtained by attaching the 2-handle $\text{clos}(Q_1 \setminus N_1)$ to N_1 along the curve p_1 . The 3-manifolds P_1 and Q_1 have the same boundary T_1 , which is either a torus or the union of two tori. Similarly, define $P_2 = M_2 \setminus D_2, Q_2 = \text{clos}(S^3 \setminus P_2)$, and $T_2 = \partial P_2 = \partial Q_2$.

Proposition 4.3. *In the situation described in the previous paragraph, the oriented surfaces S_1 and S_2 are isotopic if and only if there is a homeomorphism $f : (N_1, p_1) \rightarrow (N_2, p_2)$ such that $\text{HExt}_{p_1, p_2}(f)|_{\partial}$, seen as a map from T_1 to T_2 , extends to a homeomorphism $P_1 \rightarrow P_2$.*

Proof. If S_1 and S_2 are isotopic, let $f : S^3 \rightarrow S^3$ be a homeomorphism sending S_1 to S_2 orientation-preservingly. Since D_1 and D_2 are canonical, we can isotope f so that it sends D_1 to D_2 (and, hence, p_1 to p_2). Then $f|_{N_1} : (N_1, p_1) \rightarrow (N_2, p_2)$ is a homeomorphism such that $\text{HExt}_{p_1, p_2}(f|_{N_1})|_{\partial}$ extends to $f|_{P_1} : P_1 \rightarrow P_2$.

Conversely, let $f : (N_1, p_1) \rightarrow (N_2, p_2)$ and $g : P_1 \rightarrow P_2$ be homeomorphisms such that $\text{HExt}_{p_1, p_2}(f)$ and g have the same trace $T_1 \rightarrow T_2$. Then $\text{HExt}_{p_1, p_2}(f)$ and g can be combined to construct a self-homeomorphism of S^3 sending S_1 to S_2 orientation-preservingly. \square

Proposition 4.3 will occasionally be useful on its own. However, for most of our applications, we will refer to the following result. The proof exploits the functoriality of the handle extension map – together with the output of Theorem 2.16 – to translate the isotopy problem into a group-theoretic question which, while unsolvable in general, can be easily answered in the special cases we need.

Proposition 4.4. *With the same notations as above, we can algorithmically decide whether the oriented surfaces S_1 and S_2 are isotopic provided that:*

- i** the 3-manifold pair $(N_1, \partial N_1 \setminus p_1)$ is irreducible;
- ii** P_1 is either a (possibly trivial) knot complement, or the union of two non-trivial knot complements.

Proof. Since $(N_1, \partial N_1 \setminus p_1)$ is irreducible, we can decide whether (N_1, p_1) is homeomorphic to (N_2, p_2) or not. If it is not, then clearly S_1 and S_2 are not isotopic. Otherwise, let $f_0: (N_1, p_1) \rightarrow (N_2, p_2)$ be a homeomorphism. Similarly, since P_1 is a union of possibly trivial knot complements, we can decide if P_1 and P_2 are homeomorphic. If they are not, then once again S_1 and S_2 are not isotopic. Otherwise, let $g_0: P_1 \rightarrow P_2$ be a homeomorphism. By Proposition 4.3, we can solve the isotopy problem if we can answer the following question: is there a homeomorphism $f \in \mathbf{H}(N_2, p_2)$ such that

$$\mathrm{HExt}_{p_2, p_2}(f)|_{\partial} \circ (\mathrm{HExt}_{p_1, p_2}(f_0)|_{\partial} \circ g_0|_{\partial}^{-1}) \in \mathbf{H}(P_2)|_{\partial}?$$

By applying Theorem 2.16, we can further reduce the question to the following: given homeomorphisms $f_1, \dots, f_n \in \mathbf{H}(Q_2)|_{\partial}$ which are products of powers of Dehn twists of T_2 , and $h \in \mathbf{H}(T_2)$, does there exist $f \in \langle f_1, \dots, f_n \rangle$ such that $fh \in \mathbf{H}(P_2)|_{\partial}$?

✿ **When P_2 and Q_2 are solid tori.** Let $\ell, m \subseteq T_2$ be an oriented longitude and an oriented meridian of P_2 respectively; denote by $[\ell]$ and $[m]$ the corresponding homology classes in $H_1(T_2)$. The Dehn twists of T_2 which extend to Q_2 generate a subgroup of $\mathbf{H}(T_2)$ which is isomorphic to \mathbb{Z} – namely, the homeomorphisms that fix $[\ell]$. We can therefore compute an integer k_0 such that⁸ the group $\langle f_1, \dots, f_n \rangle$ is generated by the self-homeomorphism of T_2 sending $[\ell]$ to $[\ell]$ and $[m]$ to $k_0[\ell] + [m]$. It is immediate to check that the answer to the isotopy question is positive if and only if h sends $[m]$ to $k[\ell] \pm [m]$ where k is a multiple of k_0 .

✿ **When P_2 is a solid torus but Q_2 is not.** In this case, the group $\langle f_1, \dots, f_n \rangle$ is actually trivial, since $\mathbf{H}(Q_2)|_{\partial} \leq \langle -\mathrm{id} \rangle$ and Dehn twists cannot act as $(-\mathrm{id})$ on $H_1(T_2)$. Therefore, it is sufficient to check whether $h \in \mathbf{H}(P_2)|_{\partial}$ or not.

✿ **When P_2 is a non-trivial knot complement.** By Corollary 2.18, the group $\mathbf{H}(P_2)|_{\partial}$ is finite and can be computed. Therefore, it is sufficient to be able to decide if $h \in \langle f_1, \dots, f_n \rangle$. Note that, in this case, Q_2 is a solid torus. As argued above, the homeomorphisms f_1, \dots, f_n belong to the subgroup of $\mathbf{H}(T_2)$ fixing the homology class of the meridian of Q_2 . This group is isomorphic to \mathbb{Z} , and we can compute a generator of $\langle f_1, \dots, f_n \rangle$. By studying the action of h on $H_1(T_2)$, we can easily decide whether h belongs to $\langle f_1, \dots, f_n \rangle$ or not.

✿ **When P_2 is the union of two non-trivial knot complements.** Like before, it is enough to be able to decide if $h \in \langle f_1, \dots, f_n \rangle$. In this case, the 3-manifold Q_2 is homeomorphic to the complement of the 2-component unlink in S^3 . Denote by T and T' the two components of T_2 . Let $m \subseteq T$ and $m' \subseteq T'$ the two unique curves which are homologically trivial in Q_2 but not in T_2 , and fix an orientation for each of them. Clearly, every self-homeomorphism of Q_2 preserves $m \cup m'$ as a set up to isotopy. Moreover, Dehn twists of T_2 cannot swap T and T' or invert the orientation of m or m' . Therefore, the homeomorphisms f_1, \dots, f_n belong to the subgroup of $\mathbf{H}(T_2)$ fixing the homology classes $[m]$ and $[m']$ in $H_1(T_2)$. This group is isomorphic to $\mathbb{Z} \times \mathbb{Z}$; hence, by studying the action of h on $H_1(T_2)$, we can decide whether h belongs to $\langle f_1, \dots, f_n \rangle$ or not. \square

4.3 Compression discs for the boundary

The reader should by now reasonably believe that canonical compression discs will play a crucial role in our classification algorithm. A good strategy to prove that two discs – say –

⁸More explicitly, if f_i sends $[m]$ to $k_i[\ell] + [m]$ for $1 \leq i \leq n$, then $k_0 = \mathrm{gcd}(k_1, \dots, k_n)$. In fact, since the homeomorphisms f_1, \dots, f_n are actually products of two Dehn twists about disjoint curves, the integers k_i will always be equal to $-1, 0$, or 1 .

in M_1 and M_2 are canonical consists in characterising them as the unique discs satisfying some property which only depends on the intrinsic topologies of M_1 and M_2 . The results in this section will provide us with useful tools to prove such a characterisation.

Let (M, R) be a 3-manifold pair. Consider a property \mathcal{P} for discs properly embedded in M which is invariant under isotopies in (M, R) , such as “being separating”. We say that \mathcal{P} is *stable under boundary compressions* if

- for every disc D properly embedded in (M, R) , and
- for every possibly trivial boundary compression disc $E \subseteq M$ for D with $\partial E \subseteq D \cup R$,

we have that at least one of the two discs obtained by boundary compressing D along E satisfies \mathcal{P} . The crucial fact is that, (very) loosely speaking, discs satisfying a property which is stable under boundary compressions can be made disjoint. More precisely, we have the following.

Proposition 4.5. *Let M be an irreducible 3-manifold, and let $R \subseteq \partial M$ be a surface. Let F be a disc properly embedded in (M, R) . Suppose that there is a disc properly embedded in (M, R) not isotopic to F in M and satisfying some property \mathcal{P} which is stable under boundary compressions. Then there is a disc properly embedded in (M, R) which is not isotopic to F in M , is disjoint from F , and satisfies \mathcal{P} .*

Proof. Among all discs properly embedded in (M, R) which are in general position with respect to F , are not isotopic to F , and satisfy \mathcal{P} , pick D to minimise the number of components of $F \cap D$; we will show that D is in fact disjoint from F . Since M is irreducible, a standard innermost circle argument shows that $F \cap D$ is a collection of arcs. Suppose that this intersection is non-empty, and let $a \subseteq F \cap D$ be an outermost arc in F . The arc a cuts a disc E off of F , such that E is a (possibly trivial) boundary compression disc for D and $\partial E \subseteq D \cup R$. Let D_1 and D_2 be the two discs obtained by boundary compressing D along E , with D_1 satisfying \mathcal{P} .

If D_1 is isotopic to F , then D can in fact be isotoped to be disjoint from F , contradicting our minimality assumption. Otherwise, D_1 is a disc satisfying \mathcal{P} which is not isotopic to F , and its intersection with F has strictly fewer components than $F \cap D$. Again, this contradicts the minimality of D . It follows that D must be disjoint from F , as required. \square

The next proposition shows that some of the properties we care about are, in fact, stable under boundary compressions.

Proposition 4.6. *Let M be an irreducible 3-manifold, and let $R \subseteq \partial M$ be a surface. The following properties for discs properly embedded in (M, R) are stable under boundary compressions:*

- 1 *being non-separating;*
- 2 *having boundary which is a non-trivial curve in R .*

Proof. Let D be a disc properly embedded in (M, R) , and let $E \subseteq M$ be a possibly trivial boundary compression disc for D with $\partial E \subseteq D \cup \text{int}(R)$. Denote by D_1 and D_2 the two discs obtained by boundary compressing D along E .

- 1 The disc D is non-separating if and only if $[D] \in H^1(M)$ is non-trivial. But $[D] = [D_1] + [D_2] \in H^1(M)$, therefore if D is non-separating then one of D_1 and D_2 must be as well.

- 2 Suppose that D_1 cobounds a ball $B \subseteq M$ with some disc in R . If B is disjoint from E , then D is actually isotopic to D_2 , so ∂D_2 is trivial in R if and only if ∂D is. If instead B contains E then it also contains D_2 , and it is easy to see that ∂D is in trivial in R . \square

This is a good time to remark that, when we say that a disc properly embedded in a 3-manifold is separating, we mean that it splits the 3-manifold itself, and not only the boundary, in two connected components. However, as the following proposition shows, there is no difference for 3-manifolds which can be embedded in S^3 . Since all the 3-manifolds we will work with have this property – and, usually, already lie inside S^3 – we will freely use the term “separating” to denote discs which separate the 3-manifold they lie in and its boundary.

Proposition 4.7. *Let M be a 3-manifold embedded in S^3 with connected boundary, and let $D \subseteq M$ be a properly embedded disc. Then D separates M if and only if ∂D separates ∂M .*

Proof. Clearly if D separates M then ∂D separates ∂M . Conversely, suppose that ∂D separates ∂M but D does not separate M . Let $P = M \setminus D$ be the result of cutting M along D ; by assumption, we have that P is connected, but ∂P is the union of two components. We can interpret $\text{clos}(M \setminus P)$ as a 1-handle which, when attached to P , yields the original 3-manifold M . Let $a \subseteq M$ be the co-core arc of this 1-handle, and let b be an arc properly embedded in P connecting the two endpoints of a . Then $a \cup b$ is a closed curve in M which intersects each boundary component of P exactly once. This provides a contradiction, since M is embedded in S^3 . \square

4.4 One non-separating compression disc

We finally start presenting the actual classification algorithm. As anticipated, this section addresses the case where one side of S_1 – say M_1 – has a properly embedded non-separating disc D_1 but is not a handlebody; see Figure 4.1 for an example. Define p_1 , P_1 , Q_1 , and T_1 as described in Section 4.2 (even though we don’t know that D_1 belongs to a pair of canonical discs yet). The torus T_1 bounds a solid torus in S^3 . But P_1 cannot be a solid torus, for otherwise M_1 would be a handlebody. Hence, the solid torus must be the other component of $S^3 \setminus T_1$, namely Q_1 . In other words, P_1 is the complement of a non-trivial knot and, as such, is boundary irreducible. As described in Figure 4.2, the 3-manifold M_1 is obtained by attaching a 1-handle to P_1 , while drilling an arc from the solid torus Q_1 yields N_1 . The following lemma applied to $(P_1, \partial P_1)$ implies, *a posteriori*, that no arbitrary choice was made when selecting the disc D_1 . We remark that the lemma is stated in greater generality than we actually need here, but we will need this more general version later.

Lemma 4.8. *Let (P, R) be an irreducible 3-manifold pair. Let M be the result of attaching a 1-handle H to P along two discs in $\text{int}(R)$, and define $R' = \partial M \setminus (\partial P \setminus R)$. Then the co-core of H is the unique compression disc for R' which does not separate M .*

Proof. Let F be the co-core of H , and suppose there is another non-separating compression disc D for R' which is not isotopic to F . By Propositions 4.5 and 4.6, we may assume that D is disjoint from F and, therefore, contained in P . Since the pair (P, R) is irreducible, D must cobound a ball in P with some disc $D' \subseteq R$. But D is non-separating in M , and this is only possible if D' contains exactly one of the attaching discs of H . This would imply that D is isotopic to F , contrary to our assumption. \square




Figure 4.1. The surface S_1 admits a non-separating compression disc D_1 .

Now, if M_2 is a handlebody or S_2 has no non-separating compression discs in M_2 then clearly S_1 and S_2 are not isotopic. Otherwise, by imitating the above procedure for S_2 instead of S_1 , we define p_2 , D_2 , P_2 , Q_2 , and T_2 . Since D_1 (respectively D_2) can be intrinsically defined as the unique non-separating compression disc for S_1 (respectively S_2) in M_1 (respectively M_2), every homeomorphism $M_1 \rightarrow M_2$ must send D_1 to a disc isotopic to D_2 . If the 3-manifold pair $(N_1, \partial N_1 \setminus p_1)$ is irreducible, we immediately conclude by applying Proposition 4.4; note that this is the case for the surface depicted in Figure 4.1, since the pair $(N_1, \partial N_1 \setminus p_1)$ shown in Figure 4.2b is irreducible.

Otherwise, the surface $\partial N_1 \setminus p_1$ is compressible in N_1 , and we can find a compression disc $E_0 \subseteq N_1$ for it. The disc E_0 is properly embedded in the solid torus Q_1 ; if it is a meridian disc of Q_1 , then let us define $E_1 = E_0$. Otherwise, E_0 must cut a ball off of Q_1 ; this ball necessarily contains the curve p_1 , for otherwise E_0 would be a trivial compression disc for $\partial N_1 \setminus p_1$. We can therefore find a meridian disc E_1 of Q_1 which is disjoint from B (and, hence, from p_1).

Either way, we have found a disc E_1 properly embedded in $(N_1, \partial N_1 \setminus p_1)$ which is non-separating in Q_1 . Similarly, let E_2 be a disc properly embedded in $(N_2, \partial N_2 \setminus p_2)$ which is non-separating in Q_2 ; see the top left of Figure 4.3 for an example. We claim that S_1 and S_2 are isotopic if and only if P_1 is homeomorphic to P_2 and (N_1, p_1) is homeomorphic to (N_2, p_2) . The forward implication is trivial. Conversely, suppose that the 3-manifolds with boundary pattern (N_1, p_1) and (N_2, p_2) are homeomorphic. Note that ∂E_2 does not separate $\partial N_2 \setminus p_2$; hence, as shown in Figure 4.3, it is easy to construct a self-homeomorphism h of (N_2, p_2) which maps E_2 to itself with the opposite orientation, and is the identity in a neighbourhood of p_2 . Then the Dehn twists about E_2 , together with $\text{HExt}_{p_2, p_2}(h)$, generate the whole mapping class group of the solid torus Q_2 . As a




Figure 4.2. The 3-manifold M_1 (a) is homeomorphic to a knot complement P_1 with a 1-handle attached. The 3-manifold N_1 (b) is homeomorphic to a solid torus Q_1 with an arc drilled out.

consequence, the map

$$\mathbf{HExt}_{p_1, p_2} : \mathbf{H}((N_1, p_1); (N_2, p_2)) \longrightarrow \mathbf{H}(Q_1; Q_2)$$

is surjective. If, moreover, the knot complements P_1 and P_2 are homeomorphic, then every homeomorphism $P_1 \rightarrow P_2$ extends to a homeomorphism $Q_1 \rightarrow Q_2$ which, in turn, is the image under \mathbf{HExt}_{p_1, p_2} of some homeomorphism $(N_1, p_1) \rightarrow (N_2, p_2)$. By Proposition 4.3, this shows that S_1 and S_2 are isotopic.

Finally, we explain how to algorithmically decide whether (N_1, p_1) and (N_2, p_2) are homeomorphic or not. Note that $N_1 \setminus E_1$ is a knot complement, since E_1 is a meridian disc for Q_1 ; obviously, the same holds for $N_2 \setminus E_2$.

- ✿ If N_1 and N_2 are not handlebodies, then $N_1 \setminus E_1$ and $N_2 \setminus E_2$ are non-trivial knot complements, and p_1 and p_2 are meridian curves of $N_1 \setminus E_1$ and $N_2 \setminus E_2$ respectively. Lemma 4.8 then implies that E_1 and E_2 are the unique non-separating compression discs for $\partial N_1 \setminus p_1$ and $\partial N_2 \setminus p_2$ in N_1 and N_2 respectively. As a consequence, every homeomorphism $(N_1, p_1) \rightarrow (N_2, p_2)$ maps E_1 to E_2 ; we deduce that (N_1, p_1) and (N_2, p_2) are homeomorphic if and only if $N_1 \setminus E_1$ and $N_2 \setminus E_2$ are.
- ✿ If N_1 and N_2 are handlebodies, we claim that (N_1, p_1) and (N_2, p_2) are always homeomorphic. Note that attaching a 2-handle to $N_1 \setminus E_1$ along p_1 yields a 3-ball. In other words, the curve p_1 is a longitude of the solid torus $N_1 \setminus E_1$. As a consequence, there exists a homeomorphism $(N_1 \setminus E_1, p_1) \rightarrow (N_2 \setminus E_2, p_2)$, and every such homeomorphism extends to a homeomorphism $(N_1, p_1) \rightarrow (N_2, p_2)$.

4.5 One separating compression disc

We now deal with the case where one of the sides of S_1 – again, say M_1 – has compressible boundary, but all the compression discs are separating; an example of this situation is depicted in Figure 4.4. Let D_1 be a compression disc for S_1 in M_1 . Define p_1, P_1, Q_1 ,




Figure 4.3. There exists a self-homeomorphism h of (N_2, p_2) which flips the disc E_2 and restricts to the identity in a neighbourhood of p_2 . This homeomorphism can be constructed by cutting N_2 along E_2 , swapping the two sides of E_2 by sliding in the shaded region of ∂N_2 , and finally gluing the two sides back together.

and T_1 as described in Section 4.2. In this case, the 3-manifold P_1 is the union of two connected components $P_{1,1}$ and $P_{1,2}$, with boundary tori $T_{1,1}$ and $T_{1,2}$ respectively, so that $T_1 = T_{1,1} \cup T_{1,2}$.

The torus $T_{1,1}$ bounds a solid torus \mathbb{T} in S^3 . But $P_{1,1}$ cannot be a solid torus, since S_1 has no non-separating compression discs in M_1 . Hence, the solid torus \mathbb{T} must be the closure of the other component of $S^3 \setminus T_{1,1}$. In particular, we have that $P_{1,2} \subseteq \mathbb{T}$. Now, the torus $T_{1,2}$ must be compressible in \mathbb{T} , since it is not π_1 -injective. It cannot be compressible in $P_{1,2}$ though, because S_1 has no non-separating compression discs in M_1 . Therefore, it is easy to see that $P_{1,2}$ is contained in a ball and, moreover, it is a (non-trivial) knot complement. Clearly, the same holds for $P_{1,1}$. From this discussion, it follows that Q_1 is homeomorphic to the complement in S^3 of the 2-component unlink. As described in Figure 4.5, the 3-manifold M_1 is obtained by joining $P_{1,1}$ and $P_{1,2}$ by a 1-handle, while drilling an arc from Q_1 yields N_1 .

The following lemma implies, *a posteriori*, that no arbitrary choice was made when




Figure 4.4. The surface S_1 admits a separating compression disc D_1 .

selecting the disc D_1 .

Lemma 4.9. *Let P_1 and P_2 be irreducible boundary irreducible 3-manifolds, and let M be the result of joining them by a 1-handle H . Then the co-core of H is the unique non-trivial compression disc for ∂M in M .*

Proof. Let F be the co-core of H , and suppose there is another (non-trivial) compression disc D for ∂M which is not isotopic to F ; in particular, D is not boundary-parallel. By Propositions 4.5 and 4.6 applied to the pair $(M, \partial M)$, we may assume that D is disjoint from F and therefore, without loss of generality, contained in P_1 . Since P_1 is irreducible and boundary irreducible, D must cobound a ball in P_1 with some disc $D' \subseteq \partial P_1$. But D is not boundary parallel in M , and this is only possible if D' contains one of the attaching discs of H . This would imply that D is isotopic to F , contrary to our assumption. \square

Now, if M_2 is boundary irreducible or S_2 has a non-separating compression disc in M_2 , then clearly S_1 and S_2 are not isotopic. Otherwise, by imitating the above procedure for S_2 instead of S_1 , we define $p_2, D_2, P_2, Q_2, T_2, P_{2,1}, P_{2,2}, T_{2,1}$, and $T_{2,2}$. Since D_1 (respectively D_2) can be intrinsically defined as the unique compression disc for S_1 (respectively S_2) in M_1 (respectively M_2), every homeomorphism $M_1 \rightarrow M_2$ must send D_1 to a disc isotopic to D_2 . If the 3-manifold pair (N_1, p_1) is irreducible, we immediately conclude by applying Proposition 4.4.

Otherwise, we can assume that the surfaces $\partial N_1 \setminus p_1$ and $\partial N_2 \setminus p_2$ are compressible in N_1 and N_2 respectively. We will show that, in this case, the surfaces S_1 and S_2 are isotopic if and only if P_1 and P_2 are homeomorphic. The forward implication is trivial. Conversely, suppose that P_1 and P_2 are homeomorphic. Let E be a non-trivial compression disc for $\partial N_1 \setminus p_1$ in N_1 . The disc E is properly embedded in Q_1 , and its boundary lies in a component of T_1 , say $T_{1,1}$.

- ⚙️ If ∂E is non-trivial in $T_{1,1}$, then $N_1 \setminus E$ is homeomorphic to a solid torus \mathbb{T} via a homeomorphism sending p_1 to a trivial curve in $\partial \mathbb{T}$. Therefore, N_1 is a handlebody and p_1 is a meridian curve of N_1 .
- ⚙️ If ∂E is trivial in $T_{1,1}$, up to isotoping E , we can assume that $\partial E = p_1$. The sphere $S = E \cup D_1$ is embedded in Q_1 , and does not bound a ball, hence it separates Q_1 into




Figure 4.5. The 3-manifold M_1 (a) is homeomorphic to two knot complements $P_{1,1}$ and $P_{1,2}$ joined by a 1-handle. The 3-manifold N_1 (b) is homeomorphic to the complement Q_1 of a 2-unlink, from which an arc connecting the two boundary components has been drilled out.

two components, each homeomorphic to a punctured solid torus. As a consequence, the disc E separates N_1 into two components, each homeomorphic to a solid torus. We conclude that N_1 is a handlebody, and p_1 is a meridian curve of N_1 .

Either way, we have shown that (N_1, p_1) and (N_2, p_2) are both homeomorphic to a handlebody endowed with a separating meridian curve as a boundary pattern. Let $m_{1,1} \subseteq T_{1,1}$ and $m_{1,2} \subseteq T_{1,2}$ be the unique curves which are homologically trivial in Q_1 but not in T_1 ; define $m_{2,1} \subseteq T_{2,1}$ and $m_{2,2} \subseteq T_{2,2}$ analogously. By construction, the curves $m_{1,1}$, $m_{1,2}$, $m_{2,1}$ and $m_{2,2}$ are meridians of the knot complements $P_{1,1}$, $P_{1,2}$, $P_{2,1}$ and $P_{2,2}$ respectively. By Corollary 2.18, every homeomorphism $P_1 \rightarrow P_2$ sends $m_{1,1} \cup m_{1,2}$ to $m_{2,1} \cup m_{2,2}$ up to isotopy. Moreover, it is not hard to see that every homeomorphism $T_1 \rightarrow T_2$ with this property is induced by $\text{HExt}_{p_1, p_2}(f)$ for some $f \in \mathbf{H}((N_1, p_1); (N_2, p_2))$; see Figure 4.6 for an explanation. Therefore, the trace of every homeomorphism $P_1 \rightarrow P_2$ extends to a homeomorphism $Q_1 \rightarrow Q_2$ which is the image under HExt_{p_1, p_2} of some homeomorphism $(N_1, p_1) \rightarrow (N_2, p_2)$. Thanks to Proposition 4.3, this is enough to conclude that S_1 and S_2 are isotopic.

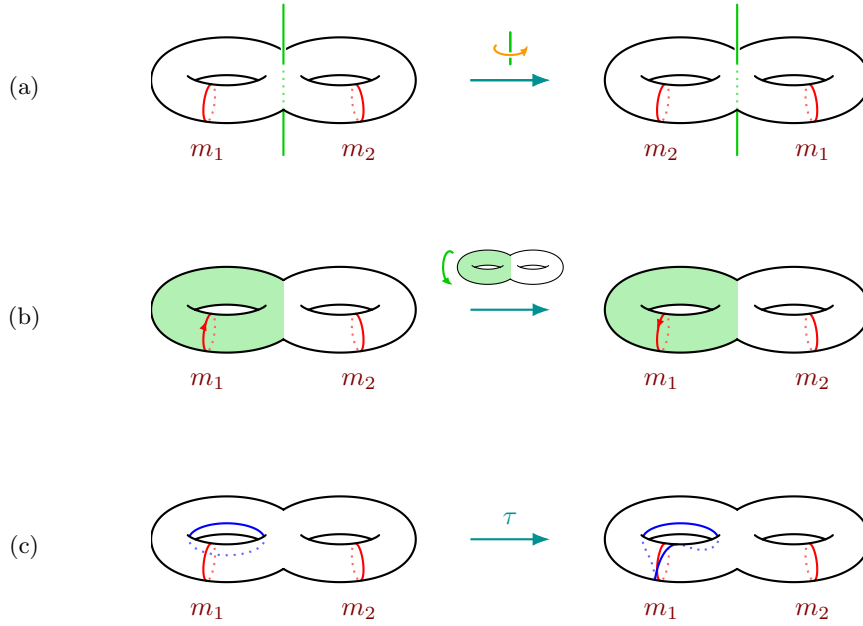


Figure 4.6. Three self-homeomorphisms of a handlebody N endowed with a separating meridian curve p as a boundary pattern (the boundary of the green region in (b)): (a) swaps the two meridian curves m_1 and m_2 ; (b) inverts the orientation of m_1 and leaves m_2 unchanged; (c) is a Dehn twist about the meridian disc bounded by m_1 . If Q denotes the 3-manifold obtained by attaching a 2-handle to N along p and $T = \partial Q$, then the images of (a), (b), and (c) under $(-)|_{\partial} \circ \text{HExt}_{p,p}$ generate the subgroup of $\mathbf{H}(T)$ of homeomorphisms preserving the isotopy class of $m_1 \cup m_2$.

4.6 Small I -bundles in the JSJ decomposition

Finally, we deal with the case where S_1 is not partially compressible on one side; in other words, we assume that each component of $S^3 \setminus S_1$ is either boundary irreducible or a handlebody. Since S_1 is compressible in S^3 , at least one of the sides – say N_1 – must be a handlebody. If M_1 is a handlebody too, then $S^3 = M_1 \cup N_1$ is a Heegaard splitting of S^3 . By a theorem of Waldhausen (see [26]), all genus-two Heegaard surfaces of S^3 are isotopic, hence S_1 and S_2 are isotopic if and only if M_2 and N_2 are handlebodies.

We can therefore assume that M_1 is boundary irreducible. If M_2 is not homeomorphic to M_1 or N_2 is not a handlebody, then clearly S_1 and S_2 are not isotopic. Otherwise, M_2 is boundary irreducible too. Let F be the JSJ system of M_2 . The boundary curves of the annuli components of F split the surface S_2 into subsurfaces R_1, \dots, R_r (namely, the components of $S_2 \setminus \partial F$) such that $\chi(R_1) + \dots + \chi(R_r) = -2$. A simple analysis shows that there are only a handful of possible surface types for each component R_i :

- the annulus $\Sigma_{0,2}$, with $\chi(\Sigma_{0,2}) = 0$;

- the pair of pants $\Sigma_{0,3}$, with $\chi(\Sigma_{0,3}) = -1$;
- the sphere with four punctures $\Sigma_{0,4}$, with $\chi(\Sigma_{0,4}) = -2$;
- the punctured torus $\Sigma_{1,1}$, with $\chi(\Sigma_{1,1}) = -1$;
- the torus with two punctures $\Sigma_{1,2}$, with $\chi(\Sigma_{1,2}) = -2$;
- the closed surface of genus two Σ_2 , with $\chi(\Sigma_2) = -2$.

We are interested in the types of I -bundles which can occur in the JSJ decomposition of M_2 .

- 1 If one of the components is a product bundle $\Sigma_{1,1} \times I$, then $S_2 \setminus \partial F$ is the union of two surfaces homeomorphic to $\Sigma_{1,1}$ (that is, the horizontal boundary of the I -bundle) and a positive number of annuli.
- 2 If one of the components is a twisted bundle $U_{2,1} \tilde{\times} I$, then $S_2 \setminus \partial F$ is the union of a surface homeomorphic to $\Sigma_{1,2}$ (that is, the horizontal boundary of the I -bundle) and a positive number of annuli.
- 3 Otherwise, each I -bundle component of the JSJ decomposition of M_2 is either a product I -bundle over $\Sigma_{0,2}$ or $\Sigma_{0,3}$, or a twisted I -bundle over $U_{1,1}$ or $U_{1,2}$.

We will now provide a solution to the isotopy problem for the last case, deferring the analysis of the first two to Sections 4.7 and 4.8.

Let $f_0: M_1 \rightarrow M_2$ be a homeomorphism. By Theorem 2.16, we can algorithmically compute

- a finite collection \mathcal{F} of self-homeomorphisms of ∂M_2 , and
- a finite collection $a_1, \dots, a_m, b_1, \dots, b_m$ of pairwise disjoint curves in ∂M_2 ,

such that

$$\mathbf{H}(M_1; M_2)|_{\partial} = \bigcup_{f \in \mathcal{F}} \langle \tau_{a_1} \tau_{b_1}^{-1}, \dots, \tau_{a_m} \tau_{b_m}^{-1} \rangle f f_0|_{\partial}.$$

Let $g_0: N_1 \rightarrow N_2$ be a homeomorphism. We have that S_1 and S_2 are isotopic if and only if some element of $\mathbf{H}(M_1; M_2)|_{\partial} g_0|_{\partial}^{-1}$, seen as a self-homeomorphism of ∂N_2 , extends to a homeomorphism $N_1 \rightarrow N_2$. For each $f \in \mathcal{F}$, thanks to Corollary 3.3, we can algorithmically decide whether there is an element of

$$\langle \tau_{a_1} \tau_{b_1}^{-1}, \dots, \tau_{a_m} \tau_{b_m}^{-1} \rangle f (f_0 g_0^{-1})|_{\partial}$$

which extends to a self-homeomorphism of N_2 . Since \mathcal{F} is finite, this is enough to solve the isotopy problem for S_1 and S_2 .

4.7 Product I -bundle over punctured torus

Only two very special cases are left – namely, those where the JSJ decomposition of M_2 contains a “large” I -bundle piece. In this section, we address the case where a component Z of $M_2 \setminus F$ is homeomorphic to $\Sigma_{1,1} \times I$, where the horizontal boundary $\partial_h Z$ consists of – say – the surfaces R_1 and R_2 ; see Figure 4.7 for an example. Due to the Euler characteristic constraint, the complement $\text{clos}(S_2 \setminus \partial_h Z)$ is an annulus. By gluing this annulus to the vertical boundary of Z , we get a torus $T_2 = (S_2 \setminus \partial_h Z) \cup \partial_v Z$ embedded in S^3 . This torus separates S^3 into two components: let P_2 be the one lying inside M_2 , and $Q_2 = N_2 \cup Z$ be the other one. If P_2 were a solid torus, then $\partial R_1 \subseteq \partial P_2$ would

be a longitude of P_2 , since it bounds a surface in the complement of P_2 . But then the meridian disc of P_2 would be a boundary compression disc for $\partial_v Z$ in M_2 , contradicting the fact that $\partial_v Z$ is an annulus in a JSJ system and, hence, not boundary parallel. Therefore, Q_2 must be a solid torus, and P_2 is the complement of a non-trivial knot. The situation is described in Figure 4.8.




Figure 4.7. The surface S_2 splits S^3 into two components: one is a handlebody, and the other has a JSJ piece Z which is an I -bundle over a punctured torus.

Let K_2 be a section of the projection $Z \rightarrow \Sigma_{1,1}$ lying in $Z \setminus \partial_h Z$ (in other words, a surface of the form $\Sigma_{1,1} \times \{1/2\} \subseteq \Sigma_{1,1} \times I$). Note that K_2 is properly embedded in Q_2 , and $Q_2 \setminus K_2$ is the handlebody N_2 .




Figure 4.8. The torus T_2 splits S^3 into two components. One of them, namely P_2 (the “outside” in this picture), is a non-trivial knot complement; the other, namely Q_2 , is a solid torus. The surface $\partial_h Z$ is a union of two punctured tori, it lies inside Q_2 , and splits it into two components: the handlebody N_2 and the product bundle Z .

Lemma 4.10. *Let K be a punctured torus properly embedded in a solid torus \mathbb{T} such that $\mathbb{T} \setminus K$ is a handlebody. Then:*

- ① *the curve $\partial K \subseteq \partial \mathbb{T}$ is a meridian;*

- 2 every punctured torus K' properly embedded in \mathbb{T} such that $\mathbb{T} \setminus K'$ is a handlebody is isotopic to K .

Proof. Statement 1 follows immediately from the fact that the meridian in $\partial\mathbb{T}$ is the only non-trivial curve which is trivial in $H_1(\mathbb{T})$. As far as statement 2 is concerned, there is a “standard” embedding of a punctured torus in \mathbb{T} , as depicted in Figure 4.9. We aim to show that K is isotopic to this standard punctured torus. Since K is not π_1 -injective, it is compressible in \mathbb{T} . Let $D \subseteq \mathbb{T}$ be a compression disc for K . We now analyse two cases, depending on whether ∂D separates K or not (although, *a posteriori*, we could always pick ∂D to be non-separating).




Figure 4.9. A “standard” punctured torus K embedded in a solid torus \mathbb{T} .

- ❖ If ∂D is non-separating in K , then compressing K along D yields a meridian disc $E \subseteq \mathbb{T}$. The punctured torus K can be recovered by removing $\partial_v \mathcal{N}(a)$ and adding $\partial_h \mathcal{N}(a)$ to E , where a is a suitable arc in \mathbb{T} such that $a \cap E = \partial a$. Note that $\mathbb{T} \setminus E$ is a 3-ball, and $P = (\mathbb{T} \setminus E) \setminus \mathring{\mathcal{N}}(a)$ is a knot complement. Moreover, if we attach a 1-handle to P , we obtain a 3-manifold which is homeomorphic to the handlebody $\mathbb{T} \setminus K$. Looking at fundamental groups, this implies that $\pi_1(P) * \mathbb{Z} = \mathbb{Z} * \mathbb{Z}$, hence the knot complement P must in fact be a solid torus. This readily implies that the arc a must be trivial, by which we mean that it must cobound a disc D' with an arc in E , such that $\text{int}(D') \cap E = \emptyset$. There is only one meridian disc E up to isotopy, and there are two trivial arcs, one on each side of E . But K is completely determined by E and a , and it is easy to see that both choices of a yield standard punctured tori in \mathbb{T} .
- ❖ If ∂D is separating in K , then compressing K along D yields the union of a meridian disc E and a torus T . The punctured torus K can be recovered by removing $\partial_v \mathcal{N}(a)$ and adding $\partial_h \mathcal{N}(a)$ to $E \cup T$, where a is a suitable arc in \mathbb{T} which has one endpoint in E , one in T , and is otherwise disjoint from $E \cup T$. Let Q be the closure of the component of $\mathbb{T} \setminus T$ which is disjoint from E . If Q is a solid torus, then K admits a compression disc whose boundary does not separate it; by our analysis of the first case, it follows that K is the standard punctured torus. Otherwise, Q is a non-trivial knot complement, and $P = (\mathbb{T} \setminus E) \setminus \text{int}(Q)$ is a punctured solid torus. Note that if we join the knot complement Q and the solid torus $P \setminus \mathring{\mathcal{N}}(a)$ with a 1-handle, we obtain a 3-manifold which is homeomorphic to the handlebody $\mathbb{T} \setminus K$. Looking at fundamental groups, this implies that $\pi_1(Q) * \mathbb{Z} = \mathbb{Z} * \mathbb{Z}$, contradicting the fact that Q is a non-trivial knot complement. \square

We know that M_1 is homeomorphic to M_2 ; in particular, it has the same JSJ decomposition. Therefore, we can define T_1, P_1, Q_1 and K_1 like we did with T_2, P_2, Q_2 and K_2 . These definitions are all canonical up to isotopy (note that there is exactly one

component in the JSJ decomposition of M_2 which is an I -bundle $\Sigma_{1,1} \times I$). Consequently, every homeomorphism $S^3 \rightarrow S^3$ sending S_1 to S_2 orientation-preservingly can be isotoped so that it sends P_1 to P_2 , Q_1 to Q_2 , and K_1 to K_2 . Clearly, if S_1 and S_2 are isotopic then the 3-manifolds with boundary pattern $(P_1, \partial K_1)$ and $(P_2, \partial K_2)$ are homeomorphic; we claim that the converse is also true. In fact, let $f: (P_1, \partial K_1) \rightarrow (P_2, \partial K_2)$ be a homeomorphism. Since ∂K_1 and ∂K_2 are meridians of the solid tori Q_1 and Q_2 respectively, f extends to a homeomorphism $f: S^3 \rightarrow S^3$ sending Q_1 to Q_2 . By Lemma 4.10, we can isotope f so that it sends K_1 to K_2 . Since $N_1 = Q_1 \setminus \mathring{N}(K_1)$ and $N_2 = Q_2 \setminus \mathring{N}(K_2)$, the homeomorphism f sends N_1 to N_2 and, hence, S_1 to S_2 orientation-preservingly.

The knot complements P_1 and P_2 are irreducible and boundary irreducible, therefore we can algorithmically decide whether $(P_1, \partial K_1)$ and $(P_2, \partial K_2)$ are homeomorphic or not. As we have shown, this provides a solution to the isotopy problem for S_1 and S_2 .

4.8 Twisted I -bundle over punctured Klein bottle

Finally, suppose there is a component Z of $M_2 \setminus F$ homeomorphic to $U_{2,1} \tilde{\times} I$, where the horizontal boundary $\partial_h Z$ is a torus with two punctures K_2 embedded in S_2 . Due to the Euler characteristic constraint, the complement $\text{clos}(S_2 \setminus K_2)$ is an annulus. We know that M_1 is homeomorphic to M_2 ; in particular, it has the same JSJ decomposition. Therefore, we can define K_1 like we did with K_2 . There are two cases.

✿ **When K_2 is incompressible in N_2 .** By gluing the annulus $\text{clos}(S_2 \setminus K_2)$ to the vertical boundary of Z , we get a torus $T_2 = (S_2 \setminus K_2) \cup \partial_v Z$ embedded in S^3 . This torus separates S^3 into two components: let P_2 be the one lying inside M_2 , and $Q_2 = N_2 \cup Z$ be the other one. Note that K_2 is properly embedded in Q_2 , and moreover it is incompressible in Q_2 , since it is incompressible on both sides (that is, in N_2 and in Z). This implies that Q_2 cannot be a solid torus, because $\pi_1(K_2) = \mathbb{Z} * \mathbb{Z} * \mathbb{Z}$ does not embed in \mathbb{Z} . Therefore, Q_2 is a non-trivial knot complement and P_2 is a solid torus.

If K_1 is compressible in N_1 , then clearly S_1 and S_2 are not isotopic. Otherwise, let us define T_1 , P_1 and Q_1 like we did with T_2 , P_2 and Q_2 . These definitions are all canonical up to isotopy (note that there is exactly one component in the JSJ decomposition of M_2 which is an I -bundle $U_{2,1} \tilde{\times} I$). It is then easy to see that S_1 and S_2 are isotopic if and only if there is a homeomorphism $S^3 \rightarrow S^3$ sending Q_1 to Q_2 and K_1 to K_2 . If Q_1 is not homeomorphic to Q_2 then once again S_1 and S_2 are not isotopic. Otherwise, by Corollary 2.18, we can algorithmically produce a list \mathcal{F} of representatives of isotopy classes of homeomorphisms $Q_1 \rightarrow Q_2$. Note that every $f \in \mathcal{F}$ extends to a homeomorphism $S^3 \rightarrow S^3$ by Corollary 2.18. If for some $f \in \mathcal{F}$ the surfaces $f(K_1)$ and K_2 are isotopic in Q_2 (which we can check, since the two surfaces are incompressible), then S_1 and S_2 are isotopic. Otherwise, they are not.

✿ **When K_2 is compressible in N_2 .** Let $D \subseteq N_2$ be a compression disc for K_2 . Let $N'_2 = N_2 \setminus D$, and let $K'_2 = \partial N'_2 \setminus (\partial N_2 \setminus K_2)$ be the result of compressing K_2 along D . If D is separating in the handlebody N_2 , then one of the two components of N'_2 is a solid torus whose boundary is fully contained in K'_2 . Therefore, up to replacing D with the meridian disc of this solid torus, we can assume that D is non-separating.

✿ Let us first assume that K'_2 is incompressible in N'_2 . Lemma 4.8 implies that D is the unique non-separating compression disc for K_2 in N_2 , up to isotopy in N_2 ; let $D_2 = D$. By repeating the above procedure for K_1 in N_1 , we can check whether K_1 admits a unique non-separating compression disc in N_1 . If this is not the case, then clearly

S_1 and S_2 are not isotopic. Otherwise, let $D_1 \subseteq N_1$ be this unique non-separating compression disc. By construction, the discs D_1 and D_2 are canonical.

- ✿ Assume now that K'_2 is compressible in N'_2 . Since K'_2 is an annulus, this immediately implies that the annulus $\text{clos}(\partial N_2 \setminus K_2)$ is compressible in N_2 or, equivalently, that the boundary components of K_2 are meridian curves of N_2 . Let D_2 be a disc properly embedded in N_2 whose boundary ∂D_2 is a boundary component of K_2 ; obviously, this disc is unique up to isotopy. If the boundary components of K_1 do not bound discs in N_1 , then clearly S_1 and S_2 are not isotopic. Otherwise, let D_1 be a disc properly embedded in N_1 whose boundary is a boundary component of K_1 . As we said before, the discs D_1 and D_2 are canonical.

Either way, we have found two canonical non-separating compression discs D_1 and D_2 for S_1 and S_2 respectively. Moreover, M_1 is boundary irreducible and the 3-manifold $N_1 \setminus D_1$ is a solid torus. By Proposition 4.4, we can algorithmically decide whether S_1 and S_2 are isotopic.

References

- [1] James W. Alexander. “On the Subdivision of 3-Space by a Polyhedron”. In: *Proceedings of the National Academy of Sciences of the United States of America* 10.1 (1924), pp. 6–8.
- [2] Matthias Aschenbrenner, Stefan Friedl and Henry Wilton. “Decision problems for 3-manifolds and their fundamental groups”. In: *Interactions between low-dimensional topology and mapping class groups*. Vol. 19. Geometry and Topology Monographs. 2015, pp. 201–236.
- [3] David R. J. Chillingworth. “A finite set of generators for the homeotopy group of a non-orientable surface”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 65.2 (1969), pp. 409–430.
- [4] David B. A. Epstein. “Curves on 2-manifolds and isotopies”. In: *Acta Mathematica* 115 (1966), pp. 83–107.
- [5] Benson Farb and Dan Margalit. *A primer on mapping class groups (PMS-49)*. Princeton Mathematical Series. Princeton, NJ: Princeton University Press, 2011.
- [6] Joachim von zur Gathen and Malte Sieveking. “A bound on solutions of linear integer equalities and inequalities”. In: *Proceedings of the American Mathematical Society* 72.1 (1978), pp. 155–158.
- [7] Cameron McA. Gordon and John Luecke. “Knots are determined by their complements”. In: *Journal of the American Mathematical Society* 2.2 (1989), pp. 371–415.
- [8] Wolfgang H. Holzmann. “An equivariant torus theorem for involutions”. In: *Transactions of the American Mathematical Society* 326.2 (1991), pp. 887–906.
- [9] William Jaco. *Lectures on Three-Manifold Topology*. Vol. 43. American Mathematical Society (AMS), 1980, p. 251.
- [10] William Jaco and Peter B. Shalen. “Seifert fibered spaces in 3-manifolds”. In: *Geometric Topology*. Elsevier, 1979, pp. 91–99.
- [11] William Jaco and Jeffrey L. Tollefson. “Algorithms for the complete decomposition of a closed 3-manifold”. In: *Illinois Journal of Mathematics* 39.3 (1995).

- [12] Klaus Johannson. *Homotopy equivalences of 3-manifolds with boundaries*. Vol. 761. Lecture Notes in Mathematics. Springer, Berlin, 1979, pp. ii+303.
- [13] Rob Kirby. “Problems in low-dimensional topology”. In: *Proceedings of Georgia Topology Conference, Part 2*. 1995.
- [14] Ryoma Kobayashi and Genki Omori. “An infinite presentation for the mapping class group of a non-orientable surface with boundary”. In: *Osaka Journal of Mathematics* 59.2 (2022), pp. 269–314.
- [15] Tsuyoshi Kobayashi. “Equivariant annulus theorem for 3-manifolds”. In: *Proceedings of the Japan Academy, Series A, Mathematical Sciences* 59.8 (1983).
- [16] Greg Kuperberg. “Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization”. In: *Pacific Journal of Mathematics* 301.1 (2019), pp. 189–241.
- [17] Jung Hoon Lee and Sangyop Lee. “Inequivalent handlebody-knots with homeomorphic complements”. In: *Algebraic & Geometric Topology* 12.2 (2012), pp. 1059–1079.
- [18] William B. R. Lickorish. “Homeomorphisms of non-orientable two-manifolds”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 59.2 (1963), pp. 307–317.
- [19] Bruno Martelli. *An introduction to geometric topology*. 2022.
- [20] Sergei V. Matveev. *Algorithmic topology and classification of 3-manifolds*. 2nd ed. Algorithms and computation in mathematics v. 9. Springer, 2007. 492 pp.
- [21] John W. Morgan and Frederick Tsz-Ho Fong. *Ricci flow and geometrization of 3-manifolds*. Vol. 53. University Lecture Series. American Mathematical Society, Providence, RI, 2010, pp. x+150.
- [22] Colin P. Rourke and Brian J. Sanderson. *Introduction to piecewise-linear topology*. Springer Study Edition. Reprint. Springer-Verlag, Berlin-New York, 1982, pp. viii+123.
- [23] Peter Scott and Hamish Short. “The homeomorphism problem for closed 3-manifolds”. In: *Algebraic & Geometric Topology* 14.4 (2014), pp. 2431–2444.
- [24] William P. Thurston. “Three-dimensional manifolds, Kleinian groups and hyperbolic geometry”. In: *American Mathematical Society. Bulletin. New Series* 6.3 (1982), pp. 357–381.
- [25] Jeffery L. Tollefson. “Involutions of sufficiently large 3-manifolds”. In: *Topology* 20.4 (1981), pp. 323–352.
- [26] Friedhelm Waldhausen. “Heegaard-Zerlegungen der 3-sphäre”. In: *Topology* 7.2 (1968), pp. 195–203.
- [27] Friedhelm Waldhausen. “On Irreducible 3-Manifolds Which are Sufficiently Large”. In: *The Annals of Mathematics* 87.1 (1968), p. 56.