Seifert Fibered Spaces: Triangulation and Recognition



Adele Jackson
Lincoln College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Michaelmas 2023

Acknowledgements

I am very grateful for the guidance of my supervisor, Marc Lackenby, who introduced me to a beautiful area of study. I have greatly benefited from his encyclopaedic knowledge of 3-manifold theory, geometric intuition, and clarity of thought. I would particularly like to express my thanks for his consistent support during a (globally) complicated and (personally) geographically tumultuous four years and eight house moves. I am grateful to the Clarendon Fund, Lincoln College and the Oxford-Australia Trust for funding my degree.

I have been fortunate to be part of the mathematical community on two continents during my doctorate. The topology group at Oxford has been a wonderful environment for me. I am grateful to all of its members during my time here, especially to Filippo Baroni, Misha Schmalian, Joe Scull, and Jan Steinebrunner, for excellent mathematical and non-mathematical conversations.

I would also like to thank the Australian geometric topology community for being so welcoming in the eighteen months I spent there in 2020-21. While I would have liked to meet more of them in person, the Australian Geometric Topology Webinar was a great source of mathematical community for me during an unsettled time. In particular, I am grateful to Ben Burton, Joan Licata, Jessica Purcell, and Stephan Tillmann, as well as the Australian National University, which hosted me for much of 2021.

I greatly appreciated the opportunity to attend several excellent conferences after my return from Australia and would like to express my thanks to their participants. The joint MATRIX/Oberwolfach workshop in September 2021 was an excellent start to a collaboration with Jack Brand, Ben Burton, Zsuzsi Dancso, Alex He, and Joan Licata. I was able to attend Oberwolfach in person for the workshop *Low-dimensional topology* in January 2023. I also greatly benefited from the conference *Computational problems in low-dimensional topology III* at Rutgers–Newark in April 2023. All of

these meetings were fantastic sources of ideas and collaboration after an extended period without much in-person mathematical discussion and were invaluable for developing the perspective of this thesis.

I am very lucky that my friends at Oxford are too many to name here, but I would particularly like to thank Luci Basualdo Bonatto, Darlan da Silva Cândido, Anjali Krishan, David Kohan Marzagão, Filippos Sytilidis, and Christoph Weis, and of course Jane Tan and Chris Williams, who made the jump from Canberra to Oxford with me. I am also grateful to Doreen Calder, Currydor, the Oxford cake seminar, and the generous and welcoming community at the Botley Meadow Allotments, who between them have kept me well-fed in heart and stomach.

Finally, I would like to thank my family, especially my Mum, my sister Carma, and my partner, Jean, for their support in my adventures on the other side of the world. None of this would have been possible without you. I wish my Dad could have seen this; I hope he would have been proud.

Abstract

In this thesis, we study 3-manifolds from an algorithmic point of view. Our motivation is the conjecture that 3-manifold homeomorphism is in **NP**. The hope is that if one can prove this conjecture in the special case of geometric 3-manifolds, then one can use geometrisation to prove the general statement. This thesis predominantly concerns itself with the Seifert fibered case.

The idea is to show that we can bound the size of characteristic topological structures in an arbitrary triangulation. To that end, we first study the size of minimal triangulations of Seifert fibered spaces. We show that we can realise certain important curves, the *singular fibres*, as subcomplexes by increasing the size of an arbitrary triangulation by a linear factor. We use this result to give bounds on the *triangulation complexity* of Seifert fibered spaces with non-empty boundary, which is the minimum number of tetrahedra in a triangulation of a manifold. We give a formula for the triangulation complexity which is correct up to a multiplicative factor. We also give an optimal bound for the triangulation complexity of a connect sum.

We then turn our attention to the 3-manifold homeomorphism problem. We first show that recognising (closed) Euclidean and Nil 3-manifolds, and deciding if a given set of Seifert data is correct for them, is in **NP**. We then consider the case of recognising Seifert fibered spaces with non-empty boundary. We prove that the recognition problem for Seifert fibered spaces with non-empty boundary is in **NP**, and that deciding whether a given Seifert fibered space with non-empty boundary admits certain Seifert data is in $\mathbf{NP} \cap \text{co-}\mathbf{NP}$. To do this, we give bounds on the size of a complete collection of normal annuli in such a Seifert fibered space. The technique we use for this bound, which is the theory of normal surfaces in *split handle structures*, comprises half this chapter and can be used generally to give a bound of the form $c^{|\mathcal{T}|^2}$ on the weight of a complete collection of (disjoint) incompressible normal surfaces in an arbitrary triangulation $\mathcal{T}$.

Statement of Originality

I declare that, to the best of my knowledge, the work in this thesis is original, unless indicated otherwise, and that no part of this material has been submitted for examination for any other degree. Section 3.3 is based on [Jac23a] and Chapter 5 is based on [Jac23b], which have both been submitted for publication.

Contents

1 Introduction **1**

2 Preliminaries **4**

2.1 Algorithms and computation 4

2.1.1 Complexity class bounds in 3-manifolds 5

2.2 Triangulations and handle structures 7

2.3 Normal surface theory 9

2.4 Algorithms for normal curves 11

2.5 Seifert fibered spaces 12

2.5.1 Classification of surfaces 13

2.5.2 Covers of Seifert fibered spaces 17

2.6 Layered solid tori and continued fractions 17

2.7 Cutting a 3-manifold along a surface 21

2.7.1 Parallelity bundles 21

2.7.2 Fast retriangulation 25

2.8 Working with handle structures 26

2.8.1 Nicely-embedded subnormal manifolds 26

2.8.2 Pre-tetrahedral handle structures 28

3 Triangulation complexity **32**

3.1 Introduction 32

3.2 Connect sums 35

3.3 Seifert fibered spaces 38

3.3.1 Simplicial singular fibres 38

3.3.2 Bounds on triangulation complexity 48

4 Recognising 3-manifolds **56**
4.1 Deciding n-manifold homeomorphism 56
4.1.1 Dimension 3 in theory and practice 57
4.1.2 3-MANIFOLD HOMEOMORPHISM: a geometric approach 58
4.2 The Euclidean and Nil cases 60
4.2.1 Reducible toroidal mapping tori 60
4.2.2 Recognising closed Euclidean manifolds 64
4.2.3 Recognising Nil manifolds 65

5 Recognition of Seifert fibered spaces with non-empty boundary **68**
5.1 Summary of results 68
5.2 A fundamental horizontal surface 69
5.3 Split handle structures 71
5.3.1 Normal surfaces 75
5.3.2 Normal surface sum 87
5.3.3 Fundamental normal surfaces 94
5.4 Bounding normal vertical annuli 97
5.4.1 Fundamental Möbius bands 97
5.4.2 A complete family of vertical annuli 99
5.5 The certificate 102

6 Further questions **110**

A Bounds on elementary disc types **114**

Bibliography **119**

List of Figures

2.1 An example of the intersection of a normal surface with a tetrahedron. 9
2.2 Gluing on a tetrahedron induces a 2-2- Pachner move. 18
2.3 A 2-2 Pachner move. 18
2.4 Part of the Farey triangulation. 19
2.5 The setup for an inessential annular simplification. 24
2.6 The possible boundary graphs in a pre-tetrahedral handle structure. . 29

3.1 A one-tetrahedron triangulation of S^3. 32
3.2 A one-vertex triangulation of the solid torus. 33
3.3 Performing a connect sum from the special spine point of view. 37
3.4 The triangulation to glue in, after identifying the labelled faces. . . . 37
3.5 An example of a maximal generalised parallelity bundle in Y_i. 41
3.6 The possible ways the boundary of a properly embedded disc in $\partial_v B$ could intersect $(\partial S_i, A_i)$. 43
3.7 Subdividing a triangle to add a new boundary component. 48
3.8 Applying three Pachner moves to the triangulation of a boundary torus of $\Sigma_n \times I$ to reduce it to a single vertex triangulation where two of the edges are the fibre and section of the product structure. 49

5.1 The pieces resulting from cutting along an elementary triangle in a tetrahedron, and the split handle structure point of view. 73
5.2 The induced split handle structure from cutting along an elementary disc in a split 0-handle that does not correspond to one in a tetrahedron. 76

A.1 The semitetrahedral split 0-handles. 114

List of Tables

3.1 The triangulation complexity $\Delta(M\#N)$ for M and N of triangulation complexity 1 or 2. 36

4.1 The Euclidean orbifolds. 60
4.2 Homology of Euclidean Seifert fibre structures. 65
4.3 Homology of Nil Seifert fibre structures. We require $e \neq 0$. 66

A.1 Homeomorphism types of elementary discs with boundary graph the complete graph on four vertices, and the number of normal isotopy types of each. 115
A.2 Homeomorphism types of elementary discs in a 0-handle with boundary graph the complete graph on four vertices with an edge removed, and the number of normal isotopy types of each. 116
A.3 Counts of elementary disc types in semitetrahedral 0-handles. 117

Notation

$\mathcal{T}^k$	when $\mathcal{T}$ is a triangulation, its k-skeleton; that is, the collection of k-cells in it
$\mathcal{T}^{(n)}$	the n^{th} barycentric subdivision of a triangulation $\mathcal{T}$
$\partial_v(\Sigma \times I)$	the vertical boundary $\partial\Sigma \times I$
$\partial_h(\Sigma \times I)$	the horizontal boundary $\Sigma \times \partial I$
$\lvert\mathcal{T}\rvert$	the number of tetrahedra in a triangulation $\mathcal{T}$
$\lvert\mathcal{H}\rvert$	the number of 0-handles in a handle structure $\mathcal{H}$
$\chi(S)$	the Euler characteristic of a manifold S
d_H	the maximum possible number of elementary disc types up to normal isotopy in a subtetrahedral (split) 0-handle
$e(M)$	when M is a Seifert fibered space, its Euler number
$s(F)$	the size of a normal surface F, which is the number of elementary discs in it
$w(F)$	the weight of a normal surface F
$\lVert q/p \rVert$	the continued fraction norm of q/p (see Definition 2.6.2)

Chapter 1

Introduction

In theory, theory and practice are the same.
In practice, they are different.

Attribution unknown

Given two mathematical objects, can we tell if they are the same? The answer, of course, depends on the class of object: given two group presentations, in general, we cannot decide if they are isomorphic; if we are told the two groups are δ-hyperbolic, we can. This thesis is motivated by the 3-manifold case.

Definition 4.1.1. The decision problem of 3-MANIFOLD HOMEOMORPHISM is the problem of, given two triangulations of compact, connected, orientable 3-manifolds with toroidal boundary, deciding if they are homeomorphic.

There is an algorithm to solve this problem as a consequence of Perelman's proof of Thurston's geometrisation conjecture, and as we discuss in Section 4.1 several proofs have been given of this fact. While known theoretical solutions could be extremely slow, in practice the question can be answered reasonably efficiently (see Section 4.1.1). We would like to better understand the difficulty of this problem from a theoretical point of view; that is, we wish to study its complexity class (as defined in Section 2.1) by asking the following well-known question.

Question A. *Is* 3-MANIFOLD HOMEOMORPHISM *in* ***NP****?*

To give a positive answer to this question, one approach is to first show that the homeomorphism problem is in **NP** for geometric 3-manifolds, which are hyperbolic, Sol, or Seifert fibered.

Theorem 5.1.1. *The problem* SEIFERT FIBERED SPACE WITH BOUNDARY RECOGNITION *is in* ***NP****.*

Theorem 5.1.2. *The problem* NAMING SEIFERT FIBERED WITH BOUNDARY *is in* $\boldsymbol{NP} \cap$ *co-*$\boldsymbol{NP}$.

To prove this result, we need to study how an arbitrary triangulation of a Seifert fibered space reflects its topology. This leads us to another question: bounding the triangulation complexity $\Delta(M)$ of a Seifert fibered space M, which is the number of tetrahedra in a minimal triangulation of M.

Theorem 3.3.1. *There exists $k > 0$ such that for any Seifert fibered manifold M with non-empty boundary other than $S^1 \times D^2$, whose Seifert data is $[\Sigma, (p_1, q_1), \ldots, (p_n, q_n)]$ with (without loss of generality) $0 < q_i < p_i$ for each i,*

$$\frac{1}{k}\left(|\chi(\Sigma)| + \sum_{i=1}^{n} \|q_i/p_i\| + 1\right) \leq \Delta(M) \leq k\left(|\chi(\Sigma)| + \sum_{i=1}^{n} \|q_i/p_i\| + 1\right).$$

Two main technical results underlie Theorems 3.3.1 and 5.1.1. First, we prove that we can find small representatives of almost all the singular fibres in any triangulation of a Seifert fibered space with non-empty boundary.

Theorem 3.3.2. *Let M be a Seifert fibered space with non-empty boundary and let $\mathcal{T}$ be a (material) triangulation of M. The collection of singular fibres of M that are not of multiplicity two have disjoint simplicial representatives in $\mathcal{T}^{(79)}$, the 79^{th} barycentric subdivision of $\mathcal{T}$. Furthermore, in $\mathcal{T}^{(82)}$, these simplicial singular fibres have disjoint simplicial solid torus neighbourhoods such that there is a simplicial meridian curve of length at most 48 for each such neighbourhood.*

Second, we consider sequences of normal surfaces in a triangulation such that each is fundamental in the complement of the previous surfaces. We use the theory of *split handle structures* to bound the complexity of the collection.

Corollary 5.3.35. *There exists a constant c_S, which we can take to be 10^{10^7}, such that the following holds. Let M be a manifold with a subtetrahedral split handle structure $\mathcal{H}$. Let $\{\Sigma_i\}$ be a collection of n disjoint normal surfaces in M such that, if we set $\mathcal{H}_0 := \mathcal{H}$ and $\mathcal{H}_i := \mathcal{H}_{i-1} \backslash\backslash \Sigma_i$, then Σ_{i+1} is a non-duplicate normal surface in $\mathcal{H}_i$ which is either fundamental or the double of a fundamental surface. There is a normal surface representative of the collection in $\mathcal{H}$ whose size is at most $c_S^{|\mathcal{H}|^2}$.*

In the context of Seifert fibered spaces, this lets us bound the complexity of a complete collection of *vertical annuli*.

Corollary 5.4.7. *There exists a constant c_T, which we can take to be $10^{10^{13}}$, such that the following holds. Let $\mathcal{T}$ be a triangulation of a Seifert fibered space M with non-empty boundary, other than the solid torus. There is a collection of disjoint normal essential annuli in $\mathcal{T}$, such that their complement is a collection of solid tori, of total weight at most $c_T^{|\mathcal{T}|^2}$. We can take these annuli to be vertical so long as M is not $T^2 \times I$ or $K \tilde{\times} I$.*

We use these results to give a "coordinate system" of distinguishing surfaces in an arbitrary Seifert fibered space with non-empty boundary, allowing us to give the polynomial time certificates needed for proving Theorem 5.1.1 and 5.1.2.

This thesis is organised as follows. We begin in Chapter 2 by surveying background material on complexity classes, normal surface theory in the algorithmic context, and some results about Seifert fibered spaces and manipulating different types of handle structures. We then turn to triangulation complexity in Chapter 3. We give bounds on the triangulation complexity of M and then study the Seifert fibered case, where we prove Theorems 3.3.1 and 3.3.2. In Chapter 4 we discuss the 3-manifold homeomorphism problem in general and prove that recognising Euclidean and Nil 3-manifolds is in **NP**. We restrict to the Seifert fibered case of the homeomorphism problem in Chapter 5. We spend much of this chapter developing the theory of split handle structures to prove Corollary 5.3.35. We hope that this work may be useful to other mathematicians in the field, as this corollary allows us to easily bound the size of a maximal collection of fundamental surfaces in a 3-manifold. We then apply the corollary to vertical annuli in a Seifert fibered space and use that result to prove Theorems 5.1.1 and 5.1.2. The final chapter briefly discusses some further questions. There is one appendix, which contains the case analysis necessary for the bound on elementary disc types in Lemma A.2. As we note in Remark 5.3.23, this optimal bound is not essential for our results; it only reduces the constants in statements such as Corollaries 5.3.35 and 5.4.7.

Chapter 2

Preliminaries

Most of this chapter consists of previously known material that we will need later. There are a few original results that are either routine or slight modifications of proofs by other authors: Lemmas 2.2.8, 2.4.2, 2.5.3, 2.5.7, 2.5.8, 2.6.3 and 2.6.4, and Corollary 2.7.10.

We assume a reasonable knowledge of 3-manifold theory and Seifert fibered spaces, such as that contained in Martelli's introductory book on the geometric topology of 3-manifolds [Mar16]. While we will state the results we need from normal surface theory, we will not rigorously establish this theory, and refer the reader to Chapters 3 and 4 of Matveev's book for proofs and precise definitions [Mat07].

2.1 Algorithms and computation

First, we give some background on algorithms and complexity; for a thorough introduction to this topic, see [AB09].

A *problem* is a function from the set of (finite) binary strings to itself; from its *input* to its *output*. The *size* of an input T, $|T|$, is the number of bits in the string. A *decision problem* is a problem whose range is the set {"yes", "no"}. An *algorithm* that *solves* a problem P is a Turing machine that, given an input to the problem, computes the corresponding output. We say a problem is *decidable* if there exists an algorithm to solve it. There exist undecidable decision problems, such as the *halting problem* (famously proven by Turing) and, given a group presentation, deciding whether it is a presentation of the trivial group (by work of Adian-Rabin).

A *complexity class* is (roughly speaking) a set of problems with solutions that satisfy certain restrictions. We will be interested in two complexity classes: **P** and **NP**. An algorithm runs in *polynomial time* if there is a polynomial p such that if T is an input, the running time of the algorithm is at most $p(|T|)$.

Definition 2.1.1. A decision problem lies in **P** if it has a polynomial time solution.

We might think of this as the class of problems that are "easy to evaluate".

Informally, a decision problem D lies in **NP** if, given an input T such that the output $D(T)$ is "yes", there is a proof of this output (a *certificate*) that can be verified in polynomial time in $|T|$. We might think of this as the class of problems that are "easy to verify".

Definition 2.1.2. A decision problem D lies in **NP** if there exists a polynomial time Turing machine V and a polynomial p such that for each input T, $D(T)$ is "yes" if and only if there exists some c, whose size is bounded by $p(|T|)$, such that $V(T, c) =$ "yes". We say that c is a *certificate* for T.

A decision problem D lies in co-**NP** if its complement problem, $\bar{D}$, which is defined by $\bar{D}(T) = \neg D(T)$, lies in **NP**.

A problem D is polynomial-time Karp reducible to another decision problem D' if there is a function f, computable in polynomial time, that takes each input to D to an input for D', such that $D(T) =$ "yes" if and only if $D'(f(T)) =$ "yes". A problem D is **NP**-hard if every problem in **NP** is polynomial-time Karp reducible to D. A problem is **NP**-complete if it is both in **NP** and is **NP**-hard.

The class **P** is a subset of **NP**. It is an open question (in fact, it is one of the Millennium Prize problems) whether **NP** is equal to **P**. If they are not equal, we expect that **NP** and co-**NP** should be distinct, and so no problem that is **NP**-complete should be in co-**NP**.

2.1.1 Complexity class bounds in 3-manifolds

We survey some decision problems in 3-manifolds that are known to be in **NP** or to be **NP**-hard. Welsh gave an interesting list of questions on the complexity class of knot theory decision problems [Wel93]. Among them was where unknot recognition sits in the complexity hierarchy. He noted that, as the unknot is characterised by being genus zero, one might hope to show that the problem is in **NP** by giving a spanning disc, but that examples were known where this spanning disc was of exponential size in the crossing number of the knot and hence in the input size [Sno90]. Hass, Lagarias and Pippenger used normal surface techniques to avoid this issue and prove that unknot recognition is in **NP** [HLP99], and Lackenby proved the problem is also in co-**NP** by using a combinatorial version of Gabai's sutured manifold hierarchies to show that giving the Thurston norm of a given homology class is in **NP** [Lac21]. As

it is generally believed that **NP** and co-**NP** are not equivalent, it seems unlikely that unknot recognition is **NP**-hard. The (open) natural next question is whether unknot recognition is in **P**.

If we generalise this problem to bounding the genus $g(K)$ of a knot K in a (triangulated orientable) 3-manifold M, Agol-Hass-Thurston showed deciding whether $g(K) < c$ where c, the triangulation of M, and K are the input is **NP**-complete [AHT06], while Lackenby and Yazdi proved that if we fix M, but allow K and c to vary, the problem is in co-**NP** [LY23].

Baldwin and Sivek showed, among other things, that recognising if a knot diagram is a torus knot is in $\mathbf{NP} \cap$ co-**NP** assuming the generalised Riemann hypothesis [BS19]. Two sets of authors independently proved that deciding if a given link diagram contains a k-component trivial unlink is **NP**-complete [KT21; MRST21], again among other results. Additionally, deciding if the crossing number of a link is at most k is **NP**-hard [MSS20]. This result implies that homeomorphism of 3-manifolds with boundary pattern is **NP**-hard. It is not known (as we discuss in Question D) whether the restriction of this problem to 3-manifold homeomorphism with no boundary pattern is **NP**-hard.

Moving away from knot theory, among irreducible orientable manifolds with non-empty boundary, deciding if a manifold admits a hyperbolic geometry is in co-**NP** [HH22]. Among Seifert fibered spaces, deciding if a manifold is a lens space is in co-**NP**[1] [HP22]. The problem of whether a manifold admits a genus g Heegaard splitting is known to be decidable by work of Johannson and Li [Li11], but the problem is also **NP**-hard [BDS17].

In Section 4.1.2 we will survey known results on when problems of the form "is M homeomorphic to N" or "does M admit a certain geometry" are in **NP**. To give an example, Ivanov and Schleimer independently proved that S^3 recognition is in **NP** [Iva08; Sch11] and Zentner showed that, assuming the Generalised Riemann Hypothesis, the problem is also in co-**NP** [Zen18].

[1]The result stated in this paper is that there is a polynomial time certificate that a Seifert fibered manifold with non-abelian fundamental group is not a lens space. The 3-manifolds with abelian fundamental groups are exactly the lens spaces, the 3-torus and $S^1 \times S^2$. Ivanov proved that recognition of $S^1 \times S^2$ is in **NP** [Iva08, Theorem 3], and in Theorem 4.2.2 we will see that recognition of the 3-torus is in **NP**, so we can remove this condition.

2.2 Triangulations and handle structures

We will use the definition of *triangulation* that is conventional in low-dimensional topology, and sometimes called a pseudo-triangulation: a (material) triangulation of a 3-manifold M is a collection of tetrahedra, with affine gluing maps between their faces, such that if we execute these gluing maps, the interior of the result is homeomorphic to M. We will also occasionally use *pseudo-ideal* triangulations. Such a triangulation is a collection of tetrahedra T with gluing maps and a set of marked vertices V, such that if we glue the tetrahedra by the gluing maps to get a space X, then $X - V$ is homeomorphic to M.

We will also later need to efficiently compute the homology of a (material) triangulation. This can be done in polynomial time for bounded dimensions as there is a polynomial time algorithm to compute the Smith normal form of an integer matrix [Ili89].

Proposition 2.2.1 (Question 33 of [KP03])**.** *There is a polynomial time algorithm to compute the homology of a finite simplicial complex of dimension at most 3 over the integers.*

We will often work with handle structures rather than triangulations, as if we cut a handle structure along an appropriately embedded transverse surface (a normal surface) we remain in the world of handle structures, which is not true for triangulations.

Convention 2.2.2. We take the definition of a handle structure to require the following:

1. each k-handle, with product structure $D^k \times D^{3-k}$, intersects the handles of lower index in exactly $\partial D^k \times D^{3-k}$, and is disjoint from the other k-handles;

2. 1-handles and 2-handles intersect in a manner compatible with their respective product structures; that is, a 1-handle $D^1 \times D^2$ intersects a 2-handle $D^2 \times D^1$ in segments of the form $D^1 \times \gamma$ in the 1-handle and $\lambda \times D^1$ in the 2-handle, where γ and λ are collections of arcs in ∂D^2 in the respective product structures.

Notation 2.2.3. Write $|\mathcal{H}|$ for the number of 0-handles in $\mathcal{H}$.

We will be particularly interested in handle structures that come from triangulations.

Definition 2.2.4. Let $\mathcal{T}$ be a triangulation (or cell structure) of a 3-manifold M. The *dual handle structure* $\mathcal{H}$ of M is formed by taking one $(3-k)$-handle for each k-simplex (or cell) of $\mathcal{T}$ that is not contained in the boundary and gluing them in the corresponding way.

We can characterise handle structures dual to triangulations using their *boundary graphs*.

Definition 2.2.5. The *boundary graph* of a 0-handle H is the thickened graph embedded in $\partial H \cong S^2$ whose thickened 0-cells are the intersections with the 1-handles and whose thickened 1-cells are the intersections with the 2-handles.

The boundary graph of a 0-handle that is dual to a tetrahedron that is disjoint from the boundary is the complete graph on four vertices.

Definition 2.2.6. A handle structure is *subtetrahedral* if its boundary graph is a subgraph of the complete graph on four vertices.

We can immediately observe the following lemma (see, for example, the discussion in §6.2 of [Lac21]).

Lemma 2.2.7. *The handle structure dual to a triangulation of a 3-manifold (possibly with boundary) is subtetrahedral.*

If our manifold has boundary, we can always remove any 3-handles.

Lemma 2.2.8. *Let $\mathcal{H}$ be a handle structure of a connected 3-manifold with non-empty boundary. There is a collection of 2-handles in $\mathcal{H}$ such that removing these 2-handles along with all the 3-handles in $\mathcal{H}$ does not change the resulting 3-manifold up to homeomorphism.*

Proof. Form a cell complex for the 3-manifold from $\mathcal{H}$ that has a $(3-k)$-cell for each k-handle, glued in the corresponding way. Also glue in a $(3-k-1)$-cell for each component of intersection of a k-handle and the boundary. The interior cells of this complex are dual to handles of $\mathcal{H}$. As this cell complex is connected, its 1-skeleton is connected, so each vertex has a path of edges from it to the boundary. The last edge in this path is dual to a 2-handle which is incident to a 3-handle. Remove this pair of handles to obtain a smaller handle structure. By induction, we can repeat this to remove all the 3-handles. □

2.3 Normal surface theory

A *normal curve* in a triangulated surface is a curve that intersects each triangle in arcs that run between *distinct* edges of the triangle. Normal surfaces, which are a higher-dimensional analogue of normal curves, were developed by Haken in the 1960s to give a combinatorial representation of interesting surfaces in a triangulated 3-manifold [Hak61b]. Let $\mathcal{T}$ be a triangulation of a 3-manifold M. A surface F in M is *normal* with respect to $\mathcal{T}$ if it intersects each tetrahedron in a collection of discs, each of whose boundary curves intersects each edge of the tetrahedron at most once. We call these *elementary discs*. One can prove that each of these discs must be a triangle or quadrilateral, and that there are seven possible disc types in each tetrahedron, some of which are shown in Figure 2.1.

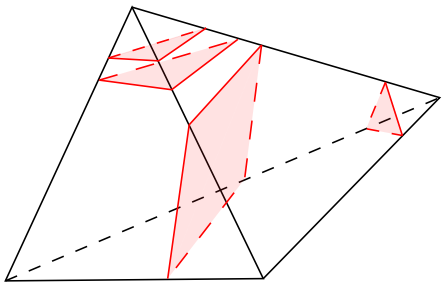


Figure 2.1: An example of the intersection of a normal surface with a tetrahedron.

Haken showed the following result when M is closed [Hak61a, Satz 2]; see Corollary 3.3.25 of [Mat07] for the general case.

Proposition 2.3.1. *Suppose that M is irreducible and ∂-irreducible, with a triangulation $\mathcal{T}$. Let F be an incompressible and ∂-incompressible surface in M. Then F is isotopic to a normal surface F' such that $|F' \cap \mathcal{T}^1|$ is at most $|F \cap \mathcal{T}^1|$.*

Definition 2.3.2. The *size* of a normal surface F, $s(F)$ is the number of elementary discs in it. The *weight* of a normal surface F, $w(F)$, is $|F \cap \mathcal{T}^1|$.

A normal surface is *minimal* if it is of minimal weight in its isotopy class in the 3-manifold.

We will state only the results of the general theory that we will need, and refer the unfamiliar reader to Chapters 3 and 4 of Matveev's book [Mat07]. We will, in fact, be using similar arguments to classical normal surface theory when we develop the theory of normal surfaces in *split handle structures* in Section 5.3. We can represent a normal surface as a unique vector in $\mathbb{Z}^{7|\mathcal{T}|}$. Conversely, a vector represents a normal

surface precisely when its entries are nonnegative, it contains only one quad type in each tetrahedron, and when two tetrahedra share a face, the count of elementary discs in each tetrahedron whose boundary contains a certain arc type is equal. Given two normal surfaces F and G in a 3-manifold, if their sum $F + G$ in $\mathbb{Z}^{7|\mathcal{T}|}$ happens to be a normal surface, then we can obtain this normal surface by resolving their curves of intersection in a canonical way.

If we have an incompressible, ∂-incompressible normal surface F, we can restrict the properties of its summands.

Proposition 2.3.3 (Theorem 6.5 [JT95], Theorem 4.1.36 [Mat07])**.** *Let F be a minimal, incompressible and ∂-incompressible normal surface in an irreducible, ∂-irreducible 3-manifold M. Suppose that F is either orientable or connected. If $F = G_1 + G_2$ then G_1 and G_2 are also incompressible and ∂-incompressible, and neither is of positive Euler characteristic.*

Definition 2.3.4. A normal surface F is *fundamental* if, whenever $F = G_1 + G_2$, one of G_1 and G_2 is empty.

A fundamental normal surface is a member of any basis over $\mathbb{Z}^+$ for the cone of normal surfaces in $\mathbb{Z}^{7|\mathcal{T}|}$: that is, it is in any *minimal Hilbert basis*. Linear integer programming techniques were used by Hass, Lagarias and Pippenger [HLP99, §6] to bound the sizes of these fundamental surfaces. The flavour of these linear programming results were summarised by Lackenby.

Proposition 2.3.5 (Theorem 8.1 [Lac21])**.** *Suppose that A is a $m \times n$ matrix. Consider the cone of solutions to $Ax = 0$, subject to the constraint that all entries of x are nonnegative. Suppose that each row of A has ℓ^2 norm at most k. If x is a fundamental solution (i.e. is in the integral Hilbert basis) to this system, then each coordinate of x is bounded by $n^{3/2}k^{n-1}$.*

Lemma 2.3.6 (Lemma 6.1 [HLP99])**.** *A fundamental normal surface F in a triangulation $\mathcal{T}$ has size at most $7|\mathcal{T}|^2 2^{7|\mathcal{T}|+2}$.*

Lemma 2.3.7 (Theorem 5.6 [HLP99], following from Corollary 6.4 [JT95])**.** *If a manifold M contains an essential disc, then in any triangulation $\mathcal{T}$ of M there is a fundamental essential disc.*

2.4 Algorithms for normal curves

We will need to quickly compute the slope of normal curves in the torus with respect to some coordinate curves. To do this, we need efficient algorithms for computing algebraic intersections. A normal curve is not equipped with an orientation, so this "algebraic" intersection is not signed; however, if we have three curves, the algebraic intersections of the pairs do have dependent signs. Thus we will arbitrarily assign orientations for the purpose of these computations. Erickson and Nayyeri dealt with this question by computing the *signed edge coordinates* of a (connected) curve from its normal vector, which is the list of its algebraic intersection number with each (oriented) edge of the triangulation. This format allows for an elegant, quick computation of the algebraic intersection number of two curves.

Proposition 2.4.1 (Corollary 6.11 [EN13])**.** *Let M be a compact 2-manifold with triangulation T. Let γ be a connected normal curve in T, represented by its unsigned normal coordinates. There is an algorithm to compute the signed normal coordinates of some orientation of γ in time polynomial in $|T|$ and $\log(s(\gamma))$, where $s(\gamma)$ is the number of normal arcs in γ.*

Lemma 2.4.2. *Let $\gamma_1, \ldots, \gamma_n$ be a collection of connected normal curves in a triangulation T of a compact 2-manifold, with an orientation of the surface that is given by an orientation of each triangle. We can compute the algebraic intersections of each pair of curves for some (fixed) choice of curve orientations in polynomial time in $|T|$ and $\sum_{i=1}^{n} \log(s(\gamma_i))$.*

Proof. This result is a modification of Corollary 6.12 of Erickson-Nayyeri [EN13] and the proof is along the same lines. Compute the signed normal coordinates of each curve and thus fix orientations on them. Fix a pair of curves, and for each edge of T, arbitrarily assign each curve to one end of the edge. Isotope the elementary arcs of the curve so that they intersect the edge only in that half, and draw each arc as a straight line arc per triangle. We can then compute the number of positive and negative crossings of the pair within each triangle in constant time by multiplying at most six pairs of signed normal coordinates. □

Schaefer, Sedgwick and Štefankovič used a similar idea to produce another polynomial time algorithm for algebraic intersection numbers using word equations [SSŠ02].

2.5 Seifert fibered spaces

A Seifert fibered space is an orientable 3-manifold which is a circle bundle over an orbifold. We can describe such a manifold by its Seifert data $[\Sigma, q_1/p_1, \ldots, q_n/p_n]$, along with, when it is closed, its Euler number e. The *singular fibres* of a Seifert fibered space are the circle fibres over the cone points of the base orbifold. All 3-manifolds with Euclidean, Nil, $S^2 \times \mathbb{R}$, spherical, $\mathbb{H}^2 \times \mathbb{R}$ and $\widetilde{SL_2}$ geometries and toroidal boundary are Seifert fibered. This is six of the eight 3-manifold geometries: the two remaining are the hyperbolic and Sol geometries.

When do two sets of Seifert data correspond to homeomorphic 3-manifolds? First, we describe Seifert data that correspond to the same Seifert fibration. A proof is given in Proposition 10.3.13 of [Mar16].

Proposition 2.5.1. *The Seifert fibrations associated to two sets of Seifert data* $[\Sigma_1, q_1/p_1, \ldots, q_m/p_m]$ *and* $[\Sigma_2, r_1/s_1, \ldots, r_n/s_n]$ *(where* p_i *and* s_j *are at least 2 for all* i *and* j*) are isomorphic (preserving orientation) if and only if* $m = n$*,* Σ_1 *and* Σ_2 *are homeomorphic, the Euler numbers* $\sum_{i=1}^m q_i/p_i$ *and* $\sum_{j=1}^m r_i/s_i$ *are equal, and (up to reordering)* $p_i = s_i$ *and* $q_i \equiv r_i \pmod{p_i}$*.*

Second, there are a few 3-manifolds that have more than one Seifert fibration, which follows from work of Waldhausen [Wal67b, Theorem 10.1]. A proof in English is as Theorem 10.4.19 of [Mar16].

Theorem 2.5.2. *Seifert fibered spaces admit unique Seifert fibrations, aside from:*

1. *the solid torus fibres as* $[D^2]$ *and* $[D^2, p/q]$*;*
2. $[D^2, 1/2, -1/2]$ *is also the circle bundle over the Möbius band;*
3. $[S^2, 1/2, -1/2, q/p] \cong [\mathbb{R}P^2, p/q]$*;*
4. $[S^2, 1/2, 1/2, -1/2, -1/2] \cong [K]$*;*
5. *lens spaces (including* $S^1 \times S^2$*) fibre in many ways.*

Lemma 2.5.3. *Given triangulated Seifert fibered manifolds* M_1 *and* M_2 *along with a set of Seifert data for each of them, where we are guaranteed that* M_1 *and* M_2 *are not lens spaces, there is a polynomial time algorithm to decide if* M_1 *and* M_2 *are homeomorphic.*

Proof. The algorithm is as follows. Replacing two singular fibres q_1/p_1 and q_2/p_2 with $(q_1 + p_1)/p_1$ and $(q_2 - p_2)/p_2$ or, if M has non-empty boundary, replacing q/p with $(q + p)/p$, does not change the Seifert fibration up to isomorphism. We can also remove $0/1$ singular fibres. Use these operations and their inverses to ensure that the multiplicity p is at least 2 for all singular fibres in M_1 and M_2. We can use the criteria in Proposition 2.5.1 to check (in polynomial time) if two sets of Seifert data with multiplicities at least 2 give orientation-preservingly isomorphic Seifert fibrations. To see if they are isomorphic disregarding orientation, we also compare the Seifert data for M_1 with that for M_2 with the signs of all of the singular fibre fractions reversed. If M_1 and M_2 have isomorphic Seifert fibrations, return "yes". Otherwise, we need to check if we are in one of the cases in Theorem 2.5.2.

If one Seifert fibration is $[D^2]$ or $[D^2, p/q]$, it is enough to check that the other is too. Similarly, if one manifold is a circle bundle over a Möbius band (case 2 of Theorem 2.5.2), check if the other is isomorphic to $[D^2, 1/2, -1/2]$. Cases 3 and 4 of Theorem 2.5.2 are analogous.

If none of these steps have concluded that the two manifolds are homeomorphic, then as they are not lens spaces, they are not. □

2.5.1 Classification of surfaces

We are interested in classifying incompressible and ∂-incompressible surfaces in Seifert fibered spaces. It is a standard result that if such a surface is orientable, it is either *horizontal* (transverse to the fibres) or *vertical* (consisting of regular fibres) (see Theorem 2.8 of [Wal67a] or for a proof in English, Proposition 10.4.9 of [Mar16]). We will state some results about non-orientable surfaces in some simple manifolds and extend this classification to the non-orientable case when the Seifert fibered space has non-empty boundary.

Proposition 2.5.4 ([Prz19; Rub78])**.** *An incompressible surface in $T^2 \times I$ is isotopic to one of the following:*

1. *a trivial sphere or disc;*
2. *an annulus $\gamma \times I$;*
3. *a ∂-parallel annulus or torus;*
4. *a non-orientable surface F, which is ∂-compressible and uniquely determined by two different slopes $\frac{p_0}{q_0} = F \cap (T^2 \times 0)$ and $\frac{p_1}{q_1} = F \cap (T^2 \times 1)$ where the curves representing these slopes intersect an even number of times.*

In the last case, F has non-orientable genus equal to the length of the minimal sequence of curves in the torus $(\gamma_1, \ldots, \gamma_n)$ from $\frac{p_0}{q_0}$ to $\frac{p_1}{q_1}$ where γ_i and γ_{i+1} intersect twice.

Proposition 2.5.5 ([Prz19; Ran96])**.** *The incompressible, non-∂-parallel, non-S^2 surfaces up to isotopy in a solid torus are classified by their intersections with the boundary, which are as single curves of slopes $\frac{p}{q}$ where q is even (that is, the slopes intersecting a meridian curve an even number of times). Any such surface with non-zero genus is non-orientable and ∂-compressible. Incompressible ∂-parallel surfaces are precisely the annuli which are unions of fibres in fibrations of the solid torus as an S^1-bundle, as well as ∂-parallel discs.*

We describe some of these incompressible non-orientable surfaces. Consider a thickened Möbius band, $S \tilde{\times} I$, which is homeomorphic to the solid torus. The Möbius band $S \tilde{\times} \{0\}$ in this solid torus is an incompressible, non-∂-parallel (but ∂-compressible) surface. We can isotope the Möbius band so it intersects a core curve of the solid torus in a single point. If we then remove a neighbourhood of the core curve from the solid torus, we obtain $T^2 \times I$ with an incompressible non-orientable surface Σ of genus one and two boundary components inside it. We can pick coordinates on the boundary of $T^2 \times I$ so that $\partial\Sigma$ is the curves of slope $1/2$ on one side (the boundary component from $\partial(\Sigma \tilde{\times} I)$) and $1/0$ on the other side. By layering this copy of $T^2 \times I$ on top of the solid torus, so that the boundaries of the interior surfaces agree, we can obtain incompressible non-orientable surfaces in $S^1 \times D^2$ of arbitrarily high genus.

Definition 2.5.6. Let M be a Seifert fibered space and let T be a collection of solid torus neighbourhoods of the singular fibres such that T is a union of fibres. A surface in a Seifert fibered space is *pseudo-vertical* if it is isotopic to a surface that is a union of fibres in $M - T$ and is incompressible in each solid torus component of T.

This notion generalises vertical surfaces.

Lemma 2.5.7. *Let M be an irreducible Seifert fibered space with non-empty boundary, and with (possibly empty) boundary pattern Γ consisting of a collection of vertical fibres. If S is an incompressible, ∂-incompressible surface in M disjoint from the boundary pattern then S is isotopic to a horizontal or pseudo-vertical surface, or is a ∂-parallel disc or trivial sphere.*

This result is a modification of the standard proof in the case when S is orientable and Mijatović's work in the case where there are no singular fibres [Mij04, Proposition 2.6]). If M is closed then S may additionally be *pseudo-horizontal*; this case is discussed in [Fro86, Theorem 2.5].

Proof. Suppose that S is not a ∂-parallel disc or a trivial sphere. If M is the solid torus, the result follows from Proposition 2.5.5, noting that S may be ∂-parallel if it is parallel to an annulus in the boundary containing curves of Γ. Otherwise M is ∂-irreducible. Take a collection of disjoint vertical annuli A disjoint from Γ: n that separate a neighbourhood of each singular fibre (that is disjoint from Γ) from M and then a minimal set of annuli that cut the remaining part of M into a solid torus. Isotope S so that it is transverse to A and minimises $|S \cap A|$. Consider $S \cap A$, which consists of arcs and closed curves. Note that $S \cap A$ does not contain any curves that are trivial in S or A, as then by the irreducibility of M and the incompressibility of S we could then reduce $|S \cap A|$. It also does not contain any arcs that are ∂-parallel in S or A as S is ∂-incompressible and M is irreducible and ∂-irreducible. As any arc in an annulus that starts and ends on the same boundary component is ∂-parallel, this means that $S \cap A$ contains none of these. Thus S intersects A in a collection of spanning arcs of A and vertical fibres. By the same reasoning, the same holds for the intersection of S with each annulus of $\partial M - \partial A$. As S is embedded it must intersect each annulus of A and $\partial M - \partial A$ in only one of these two types. If S intersects any annulus in a spanning arc, it intersects all neighbouring annuli to that one in a spanning arc, and hence (as M is connected) all annuli in the collection. Thus the two types are incompatible, so $S \cap A$ must consist of only one of the two.

Let M_0 be $M \backslash\backslash A$, and let S_0 be $S \backslash\backslash A$ in M_0. Let M_1 be the non-singular-fibre-neighbourhood component of M_0, and let S_1 be S_0 restricted to M_1. We claim that S_0 is incompressible: consider the boundary of some purported compressing disc. This curve bounds a disc in S as S is incompressible, and that disc intersects the annuli A in simple closed curves, which themselves bound discs in the annuli. We can use the irreducibility of M to isotope S through these discs, in the process reducing $|S \cap A|$.

Suppose that $S \cap A$ consists of vertical fibres. Recall the classification of incompressible surfaces in the solid torus from Proposition 2.5.5. Then S_1 is a collection of vertical annuli: it is incompressible, not a meridian disc, and cannot be an incompressible non-orientable surface as one of its boundary curves intersects the meridian once, which is odd. The remaining part of S, its intersection with the singular fibre neighbourhoods, can vary: if a given fibre has odd multiplicity, S cannot intersect the fibre so must be disjoint from the neighbourhood by the minimality of $|S \cap A|$. If the multiplicity is even, it may intersect the neighbourhood in a non-orientable incompressible surface. Gluing up, we can see that S is pseudo-vertical.

Suppose that $S \cap A$ consists of spanning arcs. (In this case Γ must be empty, as otherwise S would intersect it.) We claim that we can isotope S so that S_0 is

∂-incompressible: if it is not, let D be a non-trivial ∂-compression disc for S_0. Consider the arc $\alpha = \partial D \cap \partial M_0$. We will isotope α so that it is contained in the part of ∂M_0 coming from ∂M. In this situation we are done: as S is ∂-incompressible, we can use this boundary compression to reduce $|S \cap A|$. Note that α consists (up to isotopy) of a collection of arcs in annuli. These annuli are alternately from $\partial M_0 \cap \partial M$ and $\partial M_0 \cap A$. If α is contained in one component of $\partial M_0 \cap A$, we can use an isotopy in a collar of the boundary to push α into an adjacent component of $\partial M_0 \cap \partial M$. Otherwise, starting at one end of α, use an isotopy in the collar of the boundary to push this arc α into the adjacent annulus. We can continue this until α is contained in a single annulus. Thus S_0 is ∂-incompressible so is a collection of meridian discs in each of the solid tori, and hence is horizontal. □

We completely describe one situation that will be a pathological case for the recognition algorithm in Chapter 5: Seifert fibered spaces with a horizontal annulus.

Lemma 2.5.8. *The only Seifert fibered spaces containing horizontal Möbius bands are those with the Seifert data* $[D^2, 1/2, 1/2]$ *or that of a circle bundle over a Möbius band. The only additional Seifert fibered space containing a horizontal annulus has the Seifert structure of a circle bundle over an annulus.*

In the first two cases M is homeomorphic to $K \tilde{\times} I$.

Proof. Let M be such a Seifert fibered space. As the Euler characteristic of the Möbius band is zero, the base orbifold G has $\chi(G) = 0$. If we write $\chi(G) = 2-a-b-\sum_i(1-\frac{1}{p_i})$ where b is the number of boundary components, a is twice its genus (if orientable) or its non-orientable genus (otherwise), and $\{p_i\}$ are the multiplicities of the singular fibres, then we can see that $a + (b-1) + \sum_i(1-\frac{1}{p_i}) = 1$. As $b \geq 1$, each of these terms is nonnegative.

That M contains a horizontal Möbius band is equivalent to the statement that the base orbifold G of M is covered by a Möbius band. We can work by cases: if $a > 0$, we have a Möbius band with no singular fibres. If $b > 1$, then b must be 2 so we have an annulus, which the Möbius band does not cover but the annulus (trivially) does. Otherwise, $a = 0$ and $b = 1$, so G is a disc with some singular fibres with multiplicities such that $\sum_i(1-\frac{1}{p_i}) = 1$. For each $p > 1$, we have $\frac{1}{2} \leq 1 - \frac{1}{p} < 1$, so the only solution is two singular fibres, each with multiplicity two. □

2.5.2 Covers of Seifert fibered spaces

We will also need a good understanding of covers of Seifert fibered spaces.

A *horizontal cover* $\tilde{M} \to M$ between Seifert fibered spaces is a covering map induced by a finite-degree cover on the level of orbifolds. Each finite-degree orbifold cover induces a unique covering map on the Seifert fibrations. A *vertical cover* $\tilde{M} \to M$ is a cover that induces a homeomorphism between the base orbifolds.

Proposition 2.5.9 (Prop. 10.3.21 and 10.3.22 [Mar16])**.** *Every finite degree cover of Seifert fibered spaces $p : \tilde{M} \to M$ is a composition $p = p_h \circ p_v$ of a vertical cover p_v and a horizontal cover p_h. If we write d_h for the degree of p_h, d_v for the degree of p_v, and $\tilde{\Sigma}$ and Σ for the base orbifolds of $\tilde{M}$ and M respectively, then:*

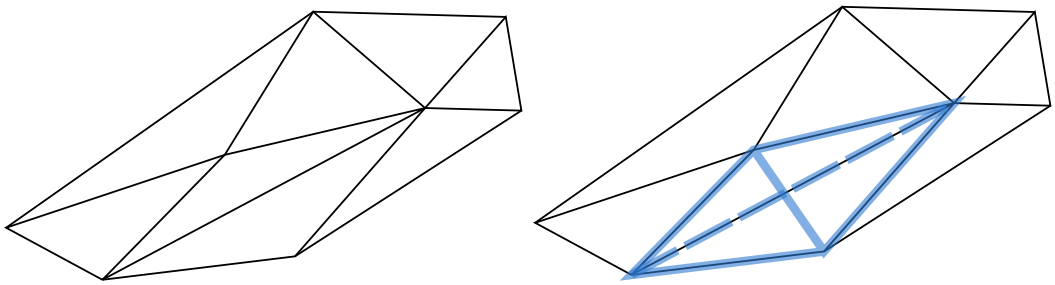
1. $\chi(\tilde{\Sigma}) = d_h\chi(\Sigma)$, *and*
2. $e(\tilde{M}) = \frac{d_h}{d_v}e(M)$.

Theorem 2.5.10 ([GH75; Sco83b])**.** *If M is a closed, orientable, irreducible 3-manifold that is finitely covered by a Seifert fibered space, then M is also Seifert fibered.*

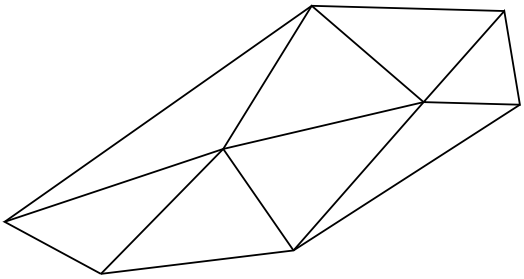
The cited papers proved this result with the additional assumption that M had infinite fundamental group. With the resolution of the geometrisation conjecture, we now know that all 3-manifolds with finite fundamental group have the elliptic geometry and thus are Seifert fibered.

2.6 Layered solid tori and continued fractions

To bound triangulations of Seifert fibered spaces, and give efficient triangulations of Dehn fillings more generally, we will use *layered triangulations* of solid tori. These were introduced by Jaco and Rubinstein [JR06]. Given a single vertex triangulation of the torus, which consists of two triangles, we can attach a tetrahedron to it by identifying two of the faces of the tetrahedron (which are joined by a single edge) with the two triangles. The effect of such a gluing, as shown in Figure 2.2 is to execute a *2-2 Pachner move* (depicted in Figure 2.3) on the triangulation. Now, if we start with the single tetrahedron triangulation of the solid torus(see [Bak] for a depiction), layering tetrahedra in this manner induces another single vertex triangulation on its boundary. Jaco and Rubinstein conjectured that this method produces minimal triangulations of the solid torus with proscribed boundary triangulation slopes.



(a) Gluing the blue tetrahedron to the surface.



(b) The induced surface triangulation.

Figure 2.2: Given a surface with an induced triangulation, if we glue on the blue tetrahedron, the new induced triangulation is the result of a 2-2 Pachner move.

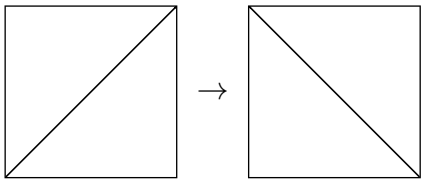


Figure 2.3: A 2-2 Pachner move.

There is a beautiful connection between layered solid tori, the Farey tree, and continued fractions. Consider the Poincaré disc model of the hyperbolic plane, $\mathbb{H}^2$, and label its boundary with $\mathbb{R} \cup \{\infty\}$ as in the upper half space model. The *Farey triangulation* of $\mathbb{H}^2$ has an edge between each pair of fractions p/q and r/s in the circle at infinity such that $|ps - qr| = 1$; a small part of it is shown in Figure 2.4. The *Farey tree* is the trivalent tree in $\mathbb{H}^2$ that is dual to this triangulation.

The Farey tree is isomorphic to the single-vertex triangulation graph of the torus, whose vertices are single-vertex triangulations of T^2 and edges are 2-2 Pachner moves. First, we identify its vertices with triangulations of T^2. As $\pi_1(T^2)$ is the abelian group $\mathbb{Z}^2$, once we have fixed two coordinate curves in T^2 intersecting once, we can identify essential simple closed curves in T^2 with $\mathbb{Q} \cup \{\infty\}$ by taking the ratio of the two coordinates. Three such curves in T^2 form a one-vertex triangulation of T^2 exactly if each pair of them intersects once. This corresponds to there being edges between each

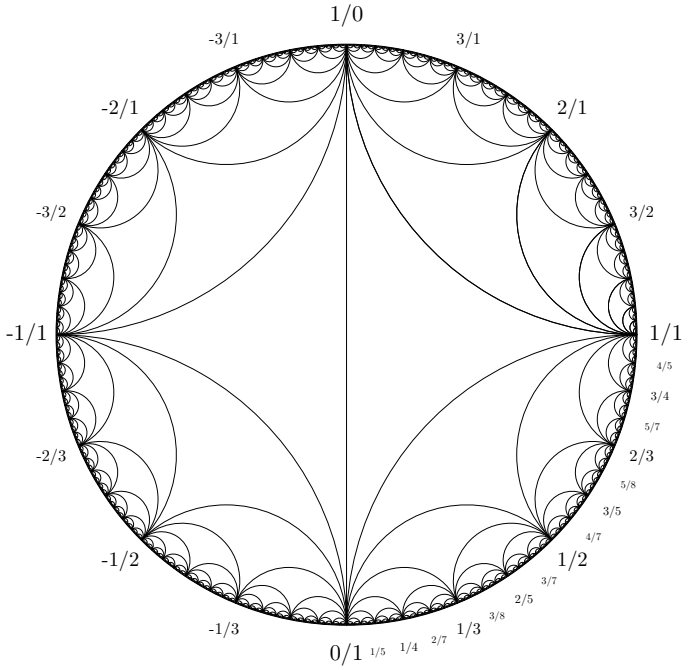


Figure 2.4: Part of the Farey triangulation.

pair of the corresponding points in $\mathbb{R} \cup \{\infty\} \cong \partial\mathbb{H}^2$ in the Farey triangulation, which bound a triangle that is dual to a vertex of the Farey tree.

For the edges, fix coordinates on the boundary of the single-tetrahedron solid torus triangulation as follows: let m be the (unique) meridian and let k and l be two of the edges of the triangulation, chosen so that the coordinate $l - k$ is the meridian m, and l is a longitude. The three edges of this triangulation of T^2 are then l, $l - m$ and $2l - m$; that is, in $\mathbb{Q} \cup \{\infty\}$ they are $0/1$, $-1/1$ and $-1/2$. Say we perform a 2-2 Pachner move over the $-1/2$ edge. The resulting triangulation will be the other single vertex triangulation of T^2 containing $0/1$ and $-1/1$, which has third edge $1/0$. In general, each pair of edges a and b that intersect once has two possible completions into a triangulation, $a + b$ and $a - b$, and the two triangulations differ by a Pachner move.

We will use continued fractions to bound the size of a minimal layered triangulation containing some edge of proscribed slope. First we need the following result of Lackenby and Purcell. The line $L(q/p)$ in the Farey tree is the union of the edges of the Farey tree dual to edges in the Farey triangulation that have one end at q/p; the vertices in $L(q/p)$ correspond to single vertex triangulations of T^2 containing q/p.

Lemma 2.6.1 (Lemma 5.4, [LP22a])**.** *Let $0 < q < p$. The distance between $L(q/p)$ and $L(\infty)$ in the Farey tree is $\|q/p\| - 1$.*

The term $\|q/p\|$ is defined as follows.

Definition 2.6.2. The positive continued fraction of a reduced fraction q/p with $q, p > 0$ is a sequence of integers $[a_0, \ldots, a_n]$ with $a_i > 0$ for all $i > 0$ such that

$$\frac{q}{p} = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\cdots + \cfrac{1}{a_n}}}}.$$

The continued fraction norm $\left\|\frac{q}{p}\right\|$ is $\sum_{i=0}^{n} a_i$.

The requirement that $a_i > 0$ for $i > 0$ means that this continued fraction is unique, so this norm is well-defined.

Lemma 2.6.3. *Let T be a one-vertex triangulation of the torus, and let λ and μ be two of its edges. Let U be the manifold that results from gluing a solid torus to T by the map taking its meridian to $p\mu + q\lambda$, where $0 < |q| < p$. There is a triangulation of U, with induced boundary triangulation T, by at most $\left\|\frac{|q|}{p}\right\| + 2$ tetrahedra.*

Proof. When we glue in a layered solid torus, we start with a single tetrahedron triangulation of the solid torus, which induces a one-vertex triangulation of its boundary. As on the previous page, let k and l be two of its edges, chosen so that the coordinate $m = l - k$ is a meridian and l is a longitude. The three edges of this triangulation of T^2 are l, $l - m$ and $2l - m$. We can thus glue on one tetrahedron so that the boundary triangulation is $(l, l - m, m)$.

The gluing map we wish to execute is $m \mapsto p\mu + q\lambda$ and $l \mapsto r\mu + s\lambda$, where r and s are integers so that $ps - qr = 1$. As our only requirement was the meridian gluing, any such r and s will do.

It is thus sufficient to determine how many tetrahedra we need to glue on to our original triangulation $(\mu, \lambda, \lambda \pm \mu)$ such that the new boundary triangulation has an edge of slope q/p. This is equivalent to asking for the distance in the Farey tree from $(0, \infty, \pm 1)$ to the line of vertices that are adjacent to the q/p slope. Up to the choice of ± 1, this distance is the same for q/p and $-q/p$, so this quantity is equal to the distance from the line of vertices adjacent to the $|q|/p$ slope to the closer of $(0, \infty, 1)$ and $(0, \infty, -1)$. As we have required that $0 \leq |q| \leq p$, this is equal to the distance from the *line* of faces adjacent to ∞ to the line of faces adjacent to $|q|/p$, which is $\||q|/p\| - 1$ (Lemma 2.6.1). We may need to add an extra tetrahedron to take us from $(0, \infty, 1)$ to $(0, \infty, -1)$ or vice versa. Thus this layered triangulation possibly first has a tetrahedron to achieve the closest choice of the ± 1 parameter, and then has $\|q/p\| - 1$ tetrahedra to change the coordinates to have a q/p slope edge, and finally

two tetrahedra to have the effect of gluing in a solid torus with the correct meridian slope: a total of $\|\,|q|/p\| + 2$ tetrahedra. The configuration $(0, \infty, \pm 1)$ is invariant under exchanging the coordinates. Doing this takes q/p to p/q, and we do indeed see that $\|\,|q|/p\| = \|p/|q|\|$. □

The following lemma gives another symmetry in the continued fraction sum.

Lemma 2.6.4. *Suppose* $0 < q < p$. *Then* $\left\|\frac{q}{p}\right\| = \left\|\frac{p-q}{p}\right\|$.

Proof. Note that $\left\|\frac{q}{p}\right\| - 1$ is the distance from the line about $\frac{q}{p}$ to the line about infinity. By reflecting in the vertical line through $1/2$, which takes $\frac{q}{p}$ to $\frac{p-q}{p}$, we see that $\left\|\frac{q}{p}\right\| - 1$ is equal to the distance from the line about $\frac{p-q}{p}$ to infinity, which is $\left\|\frac{p-q}{p}\right\| - 1$. □

2.7 Cutting a 3-manifold along a surface

We often find ourselves in the situation where we have an interesting normal surface F in a triangulated manifold M, where the number of elementary discs in F is bounded by an exponential in the number of tetrahedra in M. We may wish to cut along F to form a new manifold, $M\backslash\backslash F$, and use it to better understand M. We can do this by modifying the dual handle structure (Definition 2.2.4) to the triangulation to produce a handle structure for $M\backslash\backslash F$. Naïvely, however, cutting along F produces exponentially more handles. Most of these new handles are quite boring: they lay between parallel normal discs of F and have an induced I-bundle structure. We call the collection of these handles the *parallelity bundle*. We can take two approaches to understanding it: first, if we are analysing it theoretically, Lackenby gives us restrictions on how it sits in F [Lac09]. Alternatively, if we wish to work with $M\backslash\backslash F$ algorithmically, we have little choice other than to quickly retriangulate it (in time polynomial in the number of tetrahedra and the logarithm of the weight of the normal surface) using Lackenby's results together with work of Agol, Hass and Thurston [AHT06].

Definition 2.7.1. If $B \cong \Sigma \overset{(\sim)}{\times} I$ is an I-bundle over a surface, its *horizontal boundary*, $\partial_h B$, is $\Sigma \overset{(\sim)}{\times} \partial I$, and its *vertical boundary*, $\partial_v B$, is $\partial\Sigma \overset{(\sim)}{\times} I$.

2.7.1 Parallelity bundles

Definition 2.7.2. Let S be a surface equipped with a handle structure. A 1-manifold properly embedded S is *standard* if it is disjoint from the 2-handles, intersects the

0-handles in arcs, and intersects each 1-handle $D^1 \times D^1$ in a collection of arcs, each of which is $D^1 \times \{*\}$ for some point $*$ in D^1.

If S is a subsurface of the boundary of a 3-manifold M, we say that a handle structure $\mathcal{H}$ is a *handle structure for the pair* (M, S) if ∂S is standard with respect to the induced handle structure on ∂M.

When S is a normal surface in a handle structure for a 3-manifold N, the induced handle structure on $N \backslash\backslash S$ is a handle structure for $(N \backslash\backslash S, S)$. The parallelity handles then are as follows:

Definition 2.7.3. Let $\mathcal{H}$ be a handle structure for (M, S). A handle H of $\mathcal{H}$ is a *parallelity handle* with respect to S if it has a product structure $D^2 \times D^1$ such that $\partial_h H = H \cap S$ and such that if H' is another handle of $\mathcal{H}$, then $H \cap H'$ has the form $\gamma \times D^1$, where γ is a subset of ∂D^2.

To obtain good control over the parallelity bundle, we consider *generalised* parallelity bundles. The idea here is that some components of the complement of the parallelity bundle may have I-bundle structures that are compatible with the fibered structure. We extend the parallelity bundle over these regions.

Definition 2.7.4 (Definition 5.2 [Lac09])**.** A *generalised parallelity bundle* $\mathcal{B}$ in a handle structure $\mathcal{H}$ for (M, S) is a (possibly disconnected) subset of the handles of $\mathcal{H}$ such that:

1. $\mathcal{B}$ is endowed with an I-bundle structure over a surface;
2. any handle of $\mathcal{B}$ that intersects $\partial_v \mathcal{B}$ is a parallelity handle whose I-bundle structure agrees with that of $\mathcal{B}$;
3. $\partial_h \mathcal{B}$ is $\mathcal{B} \cap S$; and
4. if H is a handle in $\mathcal{B}$, and H' is adjacent to H and of higher index, then H' is also in $\mathcal{B}$.

A *maximal* generalised parallelity bundle is a generalised parallelity bundle that is not contained in any other generalised parallelity bundle.

The fourth condition here implies that $\mathcal{H} - \mathcal{B}$ is itself a handle structure. It is a straightforward check to show that the parallelity bundle is indeed a generalised parallelity bundle [Lac09, Lemma 5.3]. To give an example, if M is $\Sigma \times S^1$ for Σ a closed orientable surface and S is a fibre $\Sigma \times \{*\}$, with a handle structure that is

dual to a triangulation, then the maximal generalised parallelity bundle of $(M \backslash\backslash S, S)$ will be the whole manifold. The parallelity bundle of $(M \backslash\backslash S, S)$ will be some proper subset of this, as there always are some non-parallelity handles.

There is an important way in which a maximal generalised parallelity bundle of (M, S) may be unnecessarily complicated. This is when there is a component of the vertical boundary that bounds (together with ∂M) the complement of a (knotted) cylinder in M. We may not be able to incorporate this cylinder into the parallelity bundle. Figure 2.5 gives an example; as the tube passing through the center of P is knotted, we cannot extend the parallelity bundle on the complement of P to P itself. However, we can simply remove the knot complement bounded by the vertical boundary from M. (This is a special case of annular simplifications as defined by Lackenby in Definition 5.4 of [Lac09].)

Definition 2.7.5. Let $\mathcal{H}$ be a handle structure for (M, S). Suppose that there are two annuli, G and G', in M, with $\partial G = \partial G' \subset S$, such that G is properly embedded and G' is contained in the boundary. Suppose that $G \cup G'$ bounds a region P in M, which is a union of handles, such that:

1. G is a vertical boundary component of a generalised parallelity bundle contained in P;
2. P is contained in a 3-ball;
3. if H is a handle in P and H' is a handle of higher index that is adjacent to H, then H' is also in P;
4. if H is a parallelity handle of (M, S) intersecting P, then H is in P;
5. $G' \cap (\partial M - S)$ is either empty or a regular neighbourhood of a core curve of G'.

Removing the interiors of P and G' is an *inessential annular simplification.*

In an inessential annular simplification, as G is a properly embedded annulus in a 3-ball, it is compressible. When we remove P and the interior of G', we are taking a 3-ball, cutting it along an annulus, and keeping the component containing a compressing disc for the annulus. The boundary of this component is two discs and a copy of the annulus, and as the annulus is compressible in it and it is a subset of the 3-ball which is irreducible, by compressing the annulus we can see that this component is itself a 3-ball. Thus an inessential annular simplification produces a manifold M' that is homeomorphic to M.

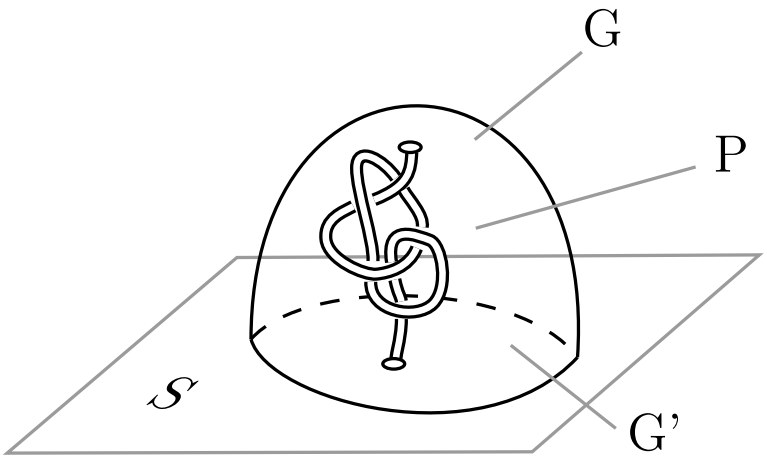


Figure 2.5: The setup for an inessential annular simplification. The annulus G consists of the knotted tube and the surface of the dome.

The new manifold M' also inherits a subsurface S' of its boundary, such that (M', S') is homeomorphic to (M, S), as follows. In the part of $\partial M'$ that comes from ∂M, S' contains the regions that were in S. If $G' \cap (\partial M - S)$ is empty, the part of $\partial M'$ that does not come from ∂M is also in S'. Otherwise, it is not.

So long as there are no possible inessential annular simplifications, the components of a maximal generalised parallelity bundle are in some sense essential.

Theorem 2.7.6 (Proposition 5.6 [Lac09])**.** *Suppose M is a compact orientable irreducible 3-manifold with a handle structure $\mathcal{H}$. Let S be an incompressible surface in ∂M with standard boundary and that is not a 2-sphere. Suppose that M does not admit any inessential annular simplifications. Let $\mathcal{B}$ be a maximal generalised parallelity bundle for (M, S). Then each component of $\mathcal{B}$ has incompressible horizontal boundary.*

Although we have defined only inessential annular simplifications, rather than Lackenby's more general annular simplifications, our case is the only one used in the proof of the proposition.

If F is an incompressible, ∂-incompressible minimal normal surface in a triangulated 3-manifold M, then if we cut along F to obtain $M\backslash\backslash F$, with a copy of the double of F, $\tilde{F}$, in its boundary, then $(M\backslash\backslash F, \tilde{F})$ admits no annular simplifications. If it did, then we could isotope F (in effect, take it through the annular simplification) to obtain a surface F' which would be transverse to the triangulation and of lower weight than F. Applying the normalisation procedure to F' which (since the surface is incompressible and ∂-incompressible) does not change its isotopy class, we would obtain a normal representative of lower weight than F.

2.7.2 Fast retriangulation

While the maximal generalised parallelity bundle lets us restrict how the horizontal boundary of this I-bundle intersects the subsurface in question, the original notion of the parallelity bundle (from Definition 2.7.3) can be used to quickly find a small triangulation of $M\backslash\backslash S$.

Proposition 2.7.7 (Proposition 13 [HH22] following Theorem 9.3 [Lac21])**.** *There is an algorithm whose input is a compact orientable triangulation $\mathcal{T}$ and a connected normal surface S in $\mathcal{T}$ given as a vector, and outputs a triangulation of $\mathcal{T}\backslash\backslash S$, whose running time is bounded by a polynomial in $|\mathcal{T}|$, $\chi(S)$, and $\log(w(S))$.*

We briefly describe the idea of the proof, as we will use related techniques. Let M be the underlying manifold of $\mathcal{T}$ and let $\mathcal{H}$ be the (subtetrahedral) handle structure on $M\backslash\backslash S$ induced by cutting the handle structure dual to $\mathcal{T}$ along S. If $\mathcal{B}$ is the parallelity bundle of $\mathcal{H}$, there are linearly many handles in $\mathcal{H} - \mathcal{B}$, of finitely many types, and so $\mathcal{H} - \mathcal{B}$ can be triangulated in time linear in $|\mathcal{T}|$.

The key idea is that as S is connected, the only way that a component of $M\backslash\backslash S$ can be entirely composed of parallelity handles is if either S is boundary parallel or S bounds an I-bundle over a non-orientable surface. Otherwise, components of $\mathcal{B}$ each border at least one non-parallelity handle, and each non-parallelity handle borders at most four parallelity handles, so we can bound the total number of components of $\mathcal{B}$ by a linear function in $|\mathcal{T}|$ and a constant, which is altogether bounded by a linear function in $|\mathcal{T}|$. We then apply the following important theorem.

Theorem 2.7.8 (Theorem 16 [AHT06])**.** *Suppose that we have some set of integers $\{1, \ldots, N\}$ and a set of k* pairings*: bijections from one subinterval $[a_i, b_i]$ of $\{1, \ldots, N\}$ to another, which are each strictly increasing or decreasing. An* orbit *of the pairing is a maximal collection of integers in $\{1, \ldots, N\}$ such that each pair in the orbit is identified by some concatenation of pairings (and their inverses). Suppose, in addition, that we are given a weight function $w : \{1, \ldots, N\} \to \mathbb{Z}^d$ such that there are m integers j whose weight differs from their neighbour $j + 1$ and such that the sum of the entries of $w(i)$ is bounded by D for any i. There is an algorithm to determine the number of orbits, a representative of each orbit, and the total weight of each orbit, where the runtime of the algorithm is polynomial in $kmd \log D \log N$.*

This theorem is very useful in algorithms that use normal surfaces. We can view a normal surface as a set of pairings between intervals of integers as follows: enumerate the elementary discs of a normal surface F by, for each tetrahedron, enumerating each

disc type in order along an edge that it intersects. Use this enumeration to label the elementary discs by the integers $\{1, \ldots, s(F)\}$. Each face induces at most five pairings between intervals of integers.

One application of the Agol-Hass-Thurston algorithm is to quickly compute the topologies of the components of a normal surface.

Proposition 2.7.9 (Corollary 17 [AHT06])**.** *Let M be a 3-manifold with a triangulation $\mathcal{T}$ and let F be a normal surface in M. There is a procedure for counting the number of components of F and determining the topology of each component that runs in time polynomial in $|\mathcal{T}| \log w(F)$.*

For retriangulation, one can use the Agol-Hass-Thurston algorithm to determine the topology of each component of $\partial_h \mathcal{B}$. One then needs to, for each component, determine whether it is in the boundary of a product I-bundle or a twisted product, but this can be done as described in the proof of Theorem 9.2 of [Lac21].

We give a slight generalisation of Proposition 2.7.7.

Corollary 2.7.10. *There is an algorithm that takes as its input both a compact connected orientable triangulation $\mathcal{T}$ of a 3-manifold M and a (possibly disconnected) normal surface S in $\mathcal{T}$ given as a vector, such that no two components of S are normally isotopic, and provides as its output a triangulation of $M \backslash\backslash S$ whose size is bounded by a polynomial in $|\mathcal{T}|$, $\log w(S)$, and the minimal Euler characteristic of the components of S, and runs in time polynomial in those same three parameters. It also produces normal surface vectors for each component of the boundary of the copy of the double of S in $\partial(M \backslash\backslash S)$.*

Proof. By Kneser-Haken finiteness there are at most $61|\mathcal{T}|$ components of S. Cut along S, applying Proposition 2.7.7 at each step. The construction of the triangulation in Proposition 2.7.7 follows Lackenby's argument in Theorem 9.2 of [Lac21], which is in the setting of sutured manifolds, and keeping track of the boundary of the surface we cut along rather than of some new sutures does not significantly modify the argument. We perform $O(|\mathcal{T}|)$ steps, each of which takes time that is at most polynomial in $|\mathcal{T}|$, $\log w(S)$ and the minimal Euler characteristic of the components of S. □

2.8 Working with handle structures

2.8.1 Nicely-embedded subnormal manifolds

In Chapter 3 we will need to consider subsets of triangulations that come from cutting along elementary discs and removing parts of tetrahedra, where no subset lies between

two parallel elementary discs. We can obtain such a subset by taking a complementary component of a normal surface in a manifold, and then removing some cells from it. Bounds on the combinatorics of such subsets were studied by Lackenby and Schleimer [LS22, §7]. We will use their results – which were obtained using elementary combinatorial techniques – in §3.3, to find simplicial embeddings of such subsets in bounded barycentric subdivisions of the original triangulation.

Definition 2.8.1. A *subnormal subset* of a handle in a handle structure $\mathcal{K}$ is a subset of a handle K of $\mathcal{K}$ that is a component of the complement of some (possibly empty) collection of elementary discs in K. A *subnormal handle structure* $\mathcal{H}$ in a handle structure $\mathcal{K}$ is a handle structure with an embedding $\mathcal{H} \hookrightarrow \mathcal{K}$ such that the image of each handle H of $\mathcal{H}$ is a subnormal subset of a handle of $\mathcal{K}$.

Definition 2.8.2. Consider a handlebody M' with handle structure $\mathcal{H}'$, consisting of only 0- and 1-handles, and marked disjoint annuli A in $\partial M'$. We say that $(\mathcal{H}', A)$ is $(0,0)$-*nicely embedded* in a handle structure $\mathcal{H}$ if:

1. each k-handle of $\mathcal{H}'$ is a subnormal subset of a k-handle of $\mathcal{H}$;
2. no handle of $\mathcal{H}'$ lies between parallel elementary discs in $\mathcal{H}$;
3. each component of intersection of an annulus in A and a handle of $\mathcal{H}'$ is a component of intersection of the handle and a handle of $\mathcal{H}$.

A 3-manifold with handle structure $\mathcal{H}''$ is $(0,0)$-*nicely embedded* in $\mathcal{H}$ if $(\mathcal{H}', A)$ is $(0,0)$-nicely embedded where $\mathcal{H}'$ is the collection of 0- and 1-handles of $\mathcal{H}''$ and $A \subseteq \partial\mathcal{H}'$ is the attaching annuli of the 2-handles.

If we take a normal surface where the only parallelity handles of its complement come from 2-handles, then any component of its complement is $(0,0)$-nicely embedded in the original manifold.

Definition 2.8.3. A subset of a manifold M with triangulation $\mathcal{T}$ is *simplicial* if it is a subcomplex of $\mathcal{T}$.

Theorem 2.8.4 (Theorem 7.6 [LS22])**.** *Suppose that M is a 3-manifold with a triangulation $\mathcal{T}$. Let $\mathcal{H}$ be the dual handle structure. Let $\mathcal{H}'$ be a handle structure for a handlebody, with A a union of disjoint annuli in its boundary, such that $(\mathcal{H}', A)$ is $(0,0)$-nicely embedded in $\mathcal{H}$. We can arrange that the following are simplicial in $\mathcal{T}^{(30)}$: the zero handles of $\mathcal{H}'$, the 1-handles of $\mathcal{H}'$ (vertically collapsed onto their co-cores), and the annuli A.*

Theorem 2.8.5 (Theorem 7.7 [LS22]). *Suppose that M is a 3-manifold with a triangulation $\mathcal{T}$. Let $\mathcal{H}$ be the dual handle structure. Let $\mathcal{H}'$ be a handle structure for a solid torus that is $(0,0)$-nicely embedded in $\mathcal{H}$. There is a core curve of the solid torus that is simplicial in $\mathcal{H}^{(79)}$.*

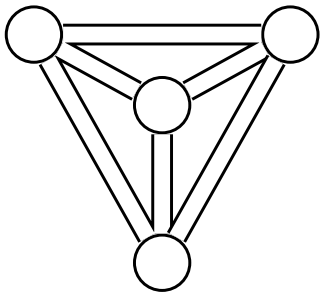
The idea of the proof of these theorems is to put an affine structure on $\mathcal{H}$ (which can be made simplicial in $\mathcal{T}^{(2)}$). One can then construct an induced affine structure on $\mathcal{H}'$. Lackenby and Schleimer then take successive barycentric subdivisions of $\mathcal{H}$: five suffices to make the boundaries of the 1-handles of $\mathcal{H}'$ simplicial, then three more to deal with the intersections of 2-handles and the normal discs cutting out 0-handles, and 16 to make these normal discs simplicial; this completes the sketch of Theorem 2.8.4. To obtain Theorem 2.8.5, they use Theorem 2.8.4 to realise the solid torus handle structure $\mathcal{H}'$ as a simplicial subset of $\mathcal{T}^{(30)}$. They then apply Theorem 4.2 of [Lac14], which gives a core curve in an affine handle structure of a solid torus, running through only the 0- and 1-handles, that intersects each handle in a bounded number of straight arcs. Lackenby and Schleimer again use barycentric subdivisions to make this core curve simplicial. The penultimate step realises this core curve as simplicial arcs in the boundaries of the original tetrahedra. The last step is to take two additional subdivisions to realise the core curve as simplicial in the interior of the original tetrahedra to ensure that it is in the interior of the manifold; as a result, the core curve contains only edges created in $\mathcal{T}^{(79)}$.

2.8.2 Pre-tetrahedral handle structures

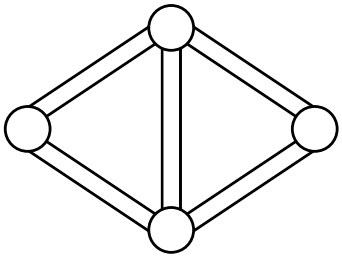
We have already discussed subtetrahedral handle structures in Section 2.2, which are handle structures where all 0-handle boundary graphs are subgraphs of the complete graph on four vertices. We will now consider a more restrictive flavour of handle structure: the *pre-tetrahedral* ones.

Definition 2.8.6. A handle structure is *pre-tetrahedral* if the boundary graph of each 0-handle is tetrahedral, semi-tetrahedral, product of length three or parallelity of length four, as illustrated in Figure 2.6.

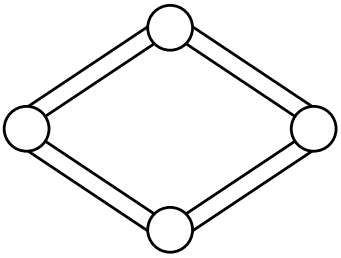
Let $\mathcal{H}$ be a pre-tetrahedral handle structure. For each 0-handle H, let α be its number of components of intersection with the 3-handles, and let β be $\frac{1}{2}$ if H is tetrahedral, $\frac{1}{4}$ if H is semi-tetrahedral, and otherwise zero. The *tetrahedral complexity* $\Delta(\mathcal{H})$ of $\mathcal{H}$ is the sum of $\frac{\alpha}{8} + \beta$ over the 0-handles.



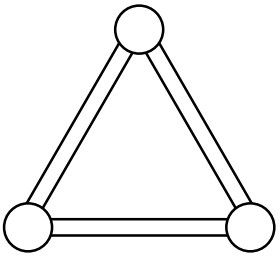
(a) Tetrahedral boundary graph. Any component of the complement may arise from intersection with ∂M.



(b) Semi-tetrahedral boundary graph. One of the components of the complement, which is adjacent to four 1-cells, must arise from ∂M.



(c) Parallelity of length four boundary graph. Both of the components of the complement must arise from intersection with ∂M.



(d) Product of length three boundary graph. At least one of the components of the complement must arise from intersection with ∂M.

Figure 2.6: The possible boundary graphs in a pre-tetrahedral handle structure.

This definition may seem rather unmotivated; its origin is that if F is a normal surface in a triangulation $\mathcal{T}$ of a closed manifold M, and $\mathcal{H}$ is the induced dual handle structure on $M\backslash\backslash F$, then the tetrahedral complexity of $\mathcal{H}$ is equal to the number of tetrahedra in T [LP22b, Lemma 4.12]. We can lower bound the tetrahedral complexity of handle structures for $T^2 \times [0,1]$ with certain short curves. For this we will need the *spine graph*.

Definition 2.8.7. A *spine* in a closed surface is an embedded graph, with no vertices of degree 1 or 2, whose complement is a disc. A *cellular spine* in a surface equipped with a cell structure is a spine that is a subset of the 1-skeleton.

An *edge contraction* is the operation of collapsing an edge that joins two distinct vertices, while an *edge expansion* is its inverse. If Γ is a spine in a surface S, e is an arc properly embedded in the complement of Γ, and f is an edge of the graph $\Gamma \cup e$ with two different components of $S\backslash\backslash(\Gamma \cup e)$ on each side of it, then an *edge swap* is the operation of removing f from Γ and adding e.

Definition 2.8.8. The *spine graph* of S, denoted $Sp(S)$, is the graph whose vertices are spines of S up to isotopy and whose edges are edge contractions and expansions.

Lemma 2.8.9. *There is a constant $k > 0$ such that the following holds. Suppose that $\mathcal{H}$ is a pre-tetrahedral handle structure of $T^2 \times [0,1]$ that admits no annular simplification. Let $\mathcal{C}$ be the (thickened) cell structure of $T^2 \times [0,1]$ where each handle becomes a 3-cell, and each component of intersection of j handles or of $(j-1)$ handles and the boundary becomes a $(4-j)$-cell. Let γ_0 and γ_1 be distinguished essential simple closed curves in the 1-skeleton of $\mathcal{C}$ that are in $T^2 \times \{0\}$ and $T^2 \times \{1\}$ respectively. Put coordinates on $T^2 \times \{0\}$ such that γ_0 has slope $1/0$. Let q/p be the slope of γ_1 in the coordinates induced from those on $T^2 \times \{0\}$ by the product structure. Suppose that each γ_i contains at most ℓ_i 1-cells, and that $0 < q < p$. Then the tetrahedral complexity $\Delta(\mathcal{H})$ is at least $\frac{1}{k}\left(\|q/p\| - 578 - 96\min(\ell_0, \ell_1)\right)$.*

The proof of this lemma is a modification of part of the proof of Theorem 1.1 in [LP22a] and is also related to the proof of Theorem 2 of [BF23].

Proof. By relabelling we can assume that $\ell_0 \geq \ell_1$. Consider the cell structure that $\mathcal{C}$ induces on the boundary. We can use the product structure to compare the cell structures on the two boundary components by projecting them onto the same torus. Pick a cellular spine Γ_0 for $T^2 \times \{0\}$ containing γ_0. Such a spine exists as γ_0 is not separating, so we can take Γ to be the union of γ_0 and some arc in the 1-skeleton that starts and ends on γ_0. There is a universal constant k' such that there is a sequence of at most $k'\Delta(\mathcal{H})$ edge contractions and expansions taking Γ_0 to some cellular spine Γ_1 in $T^2 \times \{1\}$ [LP22a, Theorem 9.1]. Now there is a sequence of at most $24 + 4\ell_1$ edge swaps taking Γ_1 to some cellular spine Γ_2 for $T^2 \times \{1\}$ containing γ_1 [LP22a, Lemma 4.15]. As we can perform any edge swap in a spine on a surface S using at most $24g(S)$ edge expansions and contractions, where $g(S)$ is the genus of S [LP22b, Lemma 8.3], we can move from Γ_0 to Γ_2 using a total of at most $k'\Delta(\mathcal{H}) + 24(24 + 4\ell_1)$ edge expansions and contractions.

Consider two complexes associated to the torus: the graph of one-vertex triangulations, $Tr(T^2)$, where the edges are 2-2 Pachner moves, and the spine graph $Sp(S)$. Each one-vertex triangulation has a spine dual to it, and edges between one-vertex triangulations, which are 2-2-Pachner moves, can be sent to a composition of an edge contraction and an expansion, so there is a quasi-embedding $Tr(T^2) \to Sp(T^2)$ that multiplies all distances by 2. Any vertex of $Sp(T^2)$ is at most distance one from one of these dual spines, so this map is in fact a quasi-isometry. (In general, for an orientable surface, this argument holds but the distance bound from an arbitrary spine to one dual to a triangulation depends on the genus.)

By this argument, Γ_0 and Γ_2 are within distance one of spines Γ_0' and Γ_2' with only valence three vertices. Let T_0 and T_2 be the dual triangulations to these spines, which contain edges of the same slopes as γ_0 and γ_1 respectively. As we can identify the Farey tree and $Tr(T^2)$, their distance in $Tr(T^2)$ is at least the distance between the lines in the Farey tree whose vertices are the triangulations containing edges of these slopes, which is $\|q/p\| - 1$ by Lemma 2.6.1.

From the quasi-isometry $\frac{1}{A}d_{Tr(T^2)}(T_0, T_2) - B \leq d_{Sp(T^2)}(\Gamma_0', \Gamma_2')$ for some universal constants A and B, so

$$\frac{1}{A}(\|q/p\| - 1) - B \leq d_{Sp(T^2)}(\Gamma_0, \Gamma_2) + 2 \leq k'\Delta(\mathcal{H}) + 24(24 + 4\ell_1) + 2$$

and if we pick k to be at least $k'A$ and also at least $2B$, we have the result. □

The constant k in this result relies on the constant in Theorem 6.1 [MMS12], which as far as the author knows has not been bounded.

Chapter 3

Triangulation complexity

3.1 Introduction

The *triangulation complexity* of a compact 3-manifold M, $\Delta(M)$, is the minimum number of tetrahedra in a triangulation of M. As with the crossing number of a knot, this invariant is challenging to compute. Some examples of minimal triangulations are shown in Figures 3.1 and 3.2, and the minimal triangulations for all closed manifolds up to 11 tetrahedra are available from Regina [BBP+22; Bur11].

The triangulation complexity is known precisely for only a few families of manifolds, including:

1. the 3-manifolds with triangulation complexity one are S^3, the lens spaces $L(4,1)$ and $L(5,2)$, the 3-ball and the solid torus,

2. when Σ is a closed orientable surface of genus g, $\Delta(\Sigma \times I) = 10g - 4$ [BFP15; JJST16];

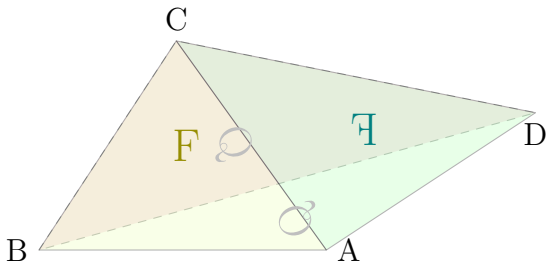


Figure 3.1: A one-tetrahedron triangulation of S^3. We identify the face ABC with the face ACD by the vertex identifications $C \to C$, $A \to A$, $B \to D$. We also identify the face ABD with the face BCD by the identifications $B \to B$, $D \to D$, $A \to C$.

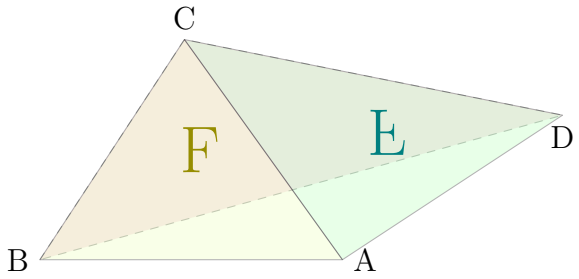


Figure 3.2: A one-vertex triangulation of the solid torus. As in the S^3 case, we identify a pair of adjacent faces ABC and ACD. However, we identify by a twist: $A \to D$, $B \to C$, and $C \to A$. See [Bak] for more images of this.

3. if M is a hyperbolic once-punctured torus bundle, its monodromy ideal triangulation (which is the same as its veering triangulation in the sense of Agol) is minimal [JRST19];

4. the triangulation complexity of the lens spaces $L(2n, 1)$ and $L(4n, 2n-1)$ is achieved by their minimal layered triangulations [JRT09; JRT11];

5. the triangulation complexities of orientable irreducible closed 3-manifolds of complexity at most 13 are known by explicit enumeration [MT20].

A few additional small families are known precisely [FV11; JRST20; JRST22; RST21], and the triangulation complexity is known to within (non-constructive) multiplicative bounds for closed hyperbolic mapping tori and for closed manifolds with elliptic or sol geometries [LP22a; LP22b].

We will consider triangulation complexity to be defined over material triangulations rather than ideal ones, but up to a factor the material triangulation is a lower bound for the ideal one. The following well-known lemma makes this precise.

Lemma 3.1.1. *Let $\Delta(M)$ be the minimal number of tetrahedra in a material triangulation of a compact 3-manifold M with boundary, and let $\Delta'(M)$ be the minimal number of tetrahedra in an ideal triangulation of M. Then $\Delta(M) \leq 14\Delta'(M)$.*

Proof. Consider an ideal triangulation of M. Truncate the tetrahedra to obtain a cell structure homeomorphic to M. Cone each non-triangle face, each of which has six sides, to a vertex of the face. Each of these six-sided faces is now triangulated by four triangles, so the 3-cells now have 20 faces each. In each truncated tetrahedron, pick a vertex of highest valence: it will be of valence at least six, and so neighbours at least six triangles. Cone the tetrahedron to this vertex, giving a triangulation of this cell

with one tetrahedron for each triangle that did not neighbour the vertex, the count of which is at most 14. Thus we have realised a material triangulation of M using $14\Delta'(M)$ tetrahedra. □

It is well-known that triangulation complexity gives a natural upper bound for simplicial volume, which was first defined by Gromov [Gro82]. We give a proof for the sake of exposition. Let N be a closed orientable n-manifold. Consider a singular n-cycle over $\mathbb{R}$ in N, $\alpha = \sum_{i=1}^{k} \gamma_i \alpha_i$ where $\gamma_i \in \mathbb{R}$ and α_i is a singular n-simplex. Define $|\alpha| = \sum_{i=1}^{k} |\gamma_i|$. Recall that, as N is closed and orientable, $H_n(N;\mathbb{R}) \cong \mathbb{R}$ and we write $[N]$ for a distinguished generator of this group, the *fundamental class*. There are many singular n-cycles representing $[N]$.

Definition 3.1.2. The *simplicial volume* or *Gromov norm* of a closed orientable manifold N is

$$\|M\| = \inf\{|\alpha| \,\big|\, [\alpha] = [N] \in H_n(N;\mathbb{R})\}.$$

Restrict to considering the case where N is a 3-manifold and fix an orientation on it. If T is a triangulation of N, by considering each of the tetrahedra t_i in T as a singular n-simplex s_i, we can build a singular 3-chain corresponding to the sum of all the tetrahedra in T with their preferred orientations, $\sum s_i$. We might hope that this is a singular n-simplex, but it is not quite: there were many different parameterisations of the t_i that we could have picked, and these differ in the ordering they give to the vertices. If we take the differential of s_i and s_j where t_i and t_j intersect at a face, while we will get a singular 2-simplex from each corresponding to the shared face, we may get different permutations of the vertices. We can correct for this as follows: consider all permutations of the 4 vertices of a tetrahedron t_i. We have a fixed 3-simplex s_i mapping Δ^3 to t_i. Consider averaging this map over all the permutations by setting

$$s_i' = \frac{1}{24} \sum_{\sigma \in S_4} (-1)^{sign(\sigma)} s_i \circ \sigma$$

where S_4 is the symmetric group on the vertices. If we consider

$$\sum s_i' = \sum_{i=1}^{|T|} \frac{1}{24} \sum_{\sigma \in S_4} (-1)^{sign(\sigma)} s_i \circ \sigma$$

we no longer encounter the issue with different permutations on shared faces. The norm of this 3-cycle, which represents the fundamental class of N, is $|T| \cdot \frac{1}{24} \cdot |S_4| = |T|$. As a result we have the following lemma.

Lemma 3.1.3. *The triangulation complexity of a closed orientable 3-manifold is at least its simplicial volume.*

In the next chapter we will consider the triangulation complexity of Seifert fibered spaces, so it is worth noting that this bound does not give us new information in that setting.

Proposition 3.1.4 (Proposition 13.2.10, [Mar16])**.** *If M is a closed Seifert-fibered 3-manifold, $\|M\| = 0$.*

(There is an extension of the definition of simplicial volume to the case with boundary, with much the same properties, which we do not discuss here. The proof given in Martelli's book shows that the simplicial volume is also zero in this case.)

We can define a related manifold invariant, the *integral simplicial volume*, by

$$\|M\|_{\mathbb{Z}} = \inf\{|\alpha| \,\big|\, [\alpha] = [N] \in H_n(N;\mathbb{Z})\}.$$

For 1-manifolds and 2-manifolds, simplicial volume, integral simplicial volume, and triangulation complexity agree [Löh18, Prop. 4.3]. Simplicial volume and triangulation complexity cannot be equal for 3-manifolds as triangulation complexity is always positive. A natural question for 3-manifolds, then, is whether integral simplicial volume and triangulation complexity are equal, or at least whether they are Lipshitz equivalent. Bertolotti and Frigerio recently showed that they are not [BF23].

Triangulation complexity is also related to Matveev complexity, which is the number of vertices in the minimal *almost simple spine* for the manifold (the full definition is given in [Mat90]).

Theorem 3.1.5 (Theorem 5, [Mat90])**.** *Suppose M is a compact, orientable, irreducible and ∂-irreducible 3-manifold, with no essential annuli. Then its triangulation complexity and Matveev complexity agree, so long as the Matveev complexity of M is not 0 (that is, $M = S^3$, $\mathbb{R}P^3$, or $L(3,1)$, possibly with a ball removed).*

3.2 Connect sums

We would like to understand how triangulation complexity behaves when we modify a 3-manifold. Both triangulation complexity and Matveev complexity are submultiplicative under finite covers: if $\tilde{M} \to M$ is a cover of degree d, $\Delta(\tilde{M}) \leq d\Delta(M)$ and similarly for Matveev complexity. While Matveev complexity is also subadditive (and in fact additive [Mat90, Theorem B]) under connect sums, triangulation complexity is more

pathological. The size of *simplicial* triangulations of 3-manifolds is subadditive under connect sum, but as in a general triangulation the boundary of each tetrahedron may not be embedded, one cannot simply remove a tetrahedron from each 3-manifold and identify their boundaries.

Proposition 3.2.1. *Let M and N be 3-manifolds, possibly with boundary. The triangulation complexity of their connect sum is bounded by*

$$\Delta(M\#N) \leq \Delta(M) + \Delta(N) + 4$$

and there exist M and N for which this is an equality.

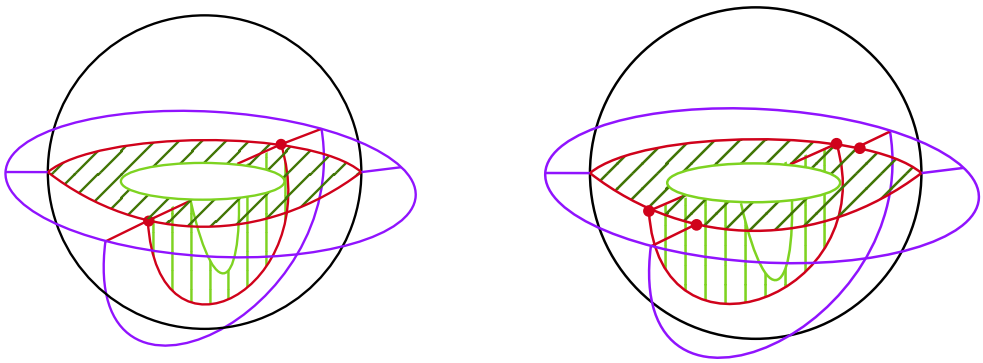
Finding M and N such that this upper bound is achieved is not difficult: we can use Regina to compute the triangulation complexities of connect sums where M and N have triangulation complexity 1 or 2, as shown in Table 3.1. Recall that the manifolds of triangulation complexity 1 are S^3, $L(4,1)$ and $L(5,2)$. Now $L(4,1)\#L(4,1)$, for example, has $\Delta(L(4,1)\#L(4,1)) = 6$.

	$L(4,1)$	$L(5,2)$	$S^2 \times S^1$	$\mathbb{R}P^3$	$L(3,1)$	$L(5,1)$	$L(7,2)$	$L(8,3)$	Y
$L(4,1)$	6	6	6	5	5	7	7	7	7
$L(5,2)$		6	6	5	5	7	7	7	7
$S^2 \times S^1$			6	5	5	7	7	7	7
$\mathbb{R}P^3$				4	4	6	6	6	6
$L(3,1)$					4	6	6	6	6
$L(5,1)$						7	7	7	7
$L(7,2)$							7	7	7
$L(8,3)$								7	7
Y									7

Table 3.1: The triangulation complexity $\Delta(M\#N)$ for M and N of triangulation complexity 1 or 2. The manifold Y in this table is the Seifert fibered space $[S^2, (2,1), (2,1), (2,-1)]$. All triangulation complexities were computed using Regina [BBP+22].

Barchechat proved Proposition 3.2.1 in the closed case with a more complicated construction than we will now give. Since the December 2021 update, Regina uses our construction to build connect sums [BBP+22, Version 7.0].

An (ideal) triangulation of a 3-manifold is dual to a 2-complex called a *special spine* where a tetrahedron of the triangulation corresponds to a vertex of the spine (see [Mat07, §1.1.4] for relevant definitions). Special spines were defined by Casler [Cas65] (under the name "standard spine") and developed in part by Matveev [Mat90]. Let M and N be 3-manifolds. The construction for Proposition 3.2.1 is as follows: from the special spine point of view, remove a ball centred



(a) (b)

Figure 3.3: Performing a connect sum from the special spine point of view.

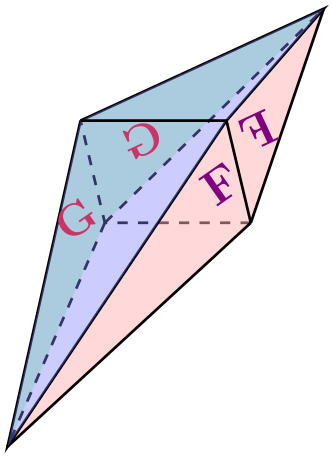


Figure 3.4: The triangulation to glue in, after identifying the labelled faces.

on an edge from each manifold. Identify the resulting boundary components, which produces the local picture shown in Figure 3.3a. To resolve this to a special spine, modify this picture to produce Figure 3.3b. We have added four vertices, so in the dual triangulation we have added four tetrahedra.

From the triangulation point of view, the procedure is to take a face pairing in each manifold and unpair the faces. (You can visualise this as puffing up the face like a crisp packet to make an empty pocket.) This manifold may not be homeomorphic to M with a 3-ball removed, as if the triangulation is ideal, the vertices of this face may not be in M. Consider the triangulation in Figure 3.4 (identifying the labelled faces). Its boundary is two copies of S^2, joined at a point. Glue them in to these unpaired face regions in the 3-manifolds. This configuration is dual to the diagram from the special spine perspective.

3.3 Seifert fibered spaces

We will determine the triangulation complexity of Seifert fibered spaces with non-empty boundary, up to multiplicative error.

Theorem 3.3.1. *There exists $k > 0$ such that for any Seifert fibered manifold M with non-empty boundary other than $S^1 \times D^2$, whose Seifert data is $[\Sigma, (p_1, q_1), \ldots, (p_n, q_n)]$ with (without loss of generality) $0 < q_i < p_i$ for each i,*

$$\frac{1}{k}\left(|\chi(\Sigma)| + \sum_{i=1}^{n} \|q_i/p_i\| + 1\right) \leq \Delta(M) \leq k\left(|\chi(\Sigma)| + \sum_{i=1}^{n} \|q_i/p_i\| + 1\right).$$

The solid torus, with $\Delta(D^2 \times S^1) = 1$, is an exception as it has infinitely many non-isomorphic Seifert fibrations (it is $[D^2, (p, q)]$ for all (p, q)), and so no such lower bound can hold for it. This triangulation complexity result is very different to the Matveev complexity situation: there, it is known that the Matveev complexity of M is at most $\sum_{i=1}^{n} \max(\|p_i/q_i\| - 3, 0)$ [FW13, Theorem 1.2], and, in particular, is zero if M has no singular fibres.

We will prove the upper bound by constructing a suitable triangulation. The lower bound is more challenging: the part from the Euler characteristic of Σ follows from the rank of the first homology, but the singular fibres are quite subtle. As we have noted, the simplicial volume of a Seifert fibered space is zero, so that lower bound does not help us here. Instead, we will make most of the singular fibres simplicial in a suitable barycentric subdivision.

Theorem 3.3.2. *Let M be a Seifert fibered space with non-empty boundary and let $\mathcal{T}$ be a (material) triangulation of M. The collection of singular fibres of M that are not of multiplicity two have disjoint simplicial representatives in $\mathcal{T}^{(79)}$, the 79^{th} barycentric subdivision of $\mathcal{T}$. Furthermore, in $\mathcal{T}^{(82)}$, these simplicial singular fibres have disjoint simplicial solid torus neighbourhoods such that there is a simplicial meridian curve of length at most 48 for each such neighbourhood.*

3.3.1 Simplicial singular fibres

We will first prove Theorem 3.3.2. A related result was proved by Mijatović [Mij04, Theorem 5.1]: he showed that a neighbourhood of each singular fibre (including those of multiplicity two) is simplicial in a suitable subdivision of $\mathcal{T}$. His bound was that this subdivision could be obtained by making at most $2^{2^{400|\mathcal{T}|^2}}$ Pachner moves. If our

bound were phrased in this language, the number of Pachner moves would be linear in the size of the triangulation.

Theorem 3.3.2 is trivial if M has no singular fibres other than those of multiplicity two. It was already known for the solid torus:

Theorem 3.3.3 (Theorem 6.14 [LS22], using Theorem 4.2 [Lac14])**.** *Let $\mathcal{T}$ be a triangulation of the solid torus. There is a simplicial core curve in $\mathcal{T}^{(51)}$.*

Lackenby and Schleimer proved this using *affine handle structures* dual to triangulations.

Definition 3.3.4 (Affine handle structure)**.** An *affine handle structure* for a 3-manifold M is a handle structure $\mathcal{H}$ such that each 0-handle and 1-handle of $\mathcal{H}$ is identified with a compact polyhedron in $\mathbb{R}^3$, satisfying the following conditions:

1. the polyhedron for each 1-handle is the product of a convex polygon and an interval;
2. the faces of the polyhedra are convex;
3. when a 0-handle and 1-handle intersect, their intersection is identified with a face of each of the corresponding polyhedra such that the map between the two faces is affine;
4. each component of intersection of a 0-handle with a 2- or 3-handle or with ∂M is a union of faces of the corresponding polyhedron.

They used the following theorem.

Theorem 3.3.5 (Theorem 4.2 [Lac14])**.** *Let $\mathcal{H}$ be an affine handle structure of the solid torus. Suppose that each 0-handle has at most four components of intersection with the 1-handles, and that each 1-handle has at most three components of intersection with the 2-handles. Then M has a core curve that intersects only the 0- and 1-handles, that respects the product structure on the 1-handles, that intersects each 1-handle in at most 24 straight arcs, and that intersects each 0-handle in at most 48 arcs. Moreover, the arcs in each 0-handle are simultaneously parallel to a collection of arcs in the boundary of the corresponding polyhedron, and each component of this collection intersects each face of the polyhedron in at most six straight arcs.*

This notion of "straight" is induced by the affine structure on the face of the polyhedron. Up to affine isotopy, there are only finitely many configurations of n straight arcs in a given polygon.

For the Seifert fibered spaces other than the solid torus, we will find *subnormal neighbourhoods* of the singular fibres and then apply Theorem 3.3.5 to them. Recall the definition of a subnormal handle structure (Definition 2.8.1).

Definition 3.3.6. Let γ be a curve in a handle structure $\mathcal{K}$. A *subnormal neighbourhood* of γ is a subnormal handle structure $\mathcal{H}$ in $\mathcal{K}$ containing γ that is homeomorphic to a solid torus for which γ is a core curve.

Let Y and Y' be 3-dimensional submanifolds of some 3-manifold M, where Y and Y' may have some boundary decoration. We say that Y and Y' are *isotopic in* M to mean that there is a (not necessarily ∂M-fixing) isotopy, preserving the boundary decoration, of the embedding map of Y in M to that of Y'.

Let $\mathcal{T}$ be a triangulation of M, and let $\mathcal{H}$ be the dual handle structure (as in Definition 2.2.4). Our approach is to take a collection of normal annuli that cut out a neighbourhood of each singular fibre. We will then prove that, for each of these singular fibres, in the induced handle structure of their solid torus neighbourhood there is a core curve consisting of a bounded number of arcs in each 0- and 1-handle that avoids the parallelity bundle altogether.

Let A be a properly-embedded vertical annulus in M that, together with an annulus A' in the boundary with $\partial A = \partial A'$, separates a neighbourhood of a singular fibre from M. Take a disjoint family of one such A_i for each non-multiplicity-two singular fibre γ_i. Since these annuli are essential (as M is not a solid torus), we can isotope them to be simultaneously normal. Write $\{Y_i\}$ for the collection of subnormal neighbourhoods of singular fibres that they separate from M, and write $\{\mathcal{H}_i\}$ for the induced handle structures of these subnormal neighbourhoods. As annular simplifications do not change $(Y_i, \partial A_i)$ up to isotopy in M, and reduce the number of handles, perform them until no more are possible.

Removing the parallelity bundle

Fix our attention on one singular fibre neighbourhood Y_i, recalling that the singular fibre which is its core curve is not of multiplicity two. Let $\mathcal{B}$ be a maximal generalised parallelity bundle for $\mathcal{H}_i$ with respect to A_i. As A_i is essential and normal in M, so is incompressible in ∂Y_i and has standard boundary, by Theorem 2.7.6 the horizontal boundary of $\mathcal{B}$ is an incompressible subsurface of ∂Y_i, which is contained in A_i by

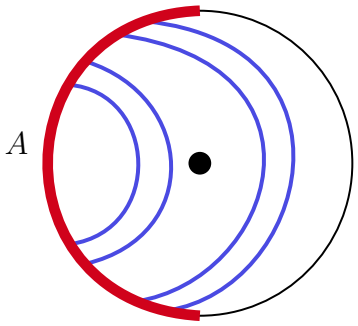


Figure 3.5: An example of a maximal generalised parallelity bundle in Y_i. We show the projection of the parallelity bundle onto the underlying orbifold: a disk with one cone point (the black dot, which is the image of the singular fibre). Here, $\mathcal{B}$ has two components, which are both I-bundles over annuli. The vertical boundary of one intersects the annulus A.

definition. Thus each component of $\mathcal{B}$ is an I-bundle over a disc, annulus, or Möbius band; in the annulus and Möbius band bundle cases, their horizontal boundaries form a collection of annuli in A_i which have core curves that are isotopic to the core curve of A_i. We show an example of what $\mathcal{B}$ might look like in Figure 3.5. We will prove that Möbius band bundles do not occur when the singular fibre γ_i is not of multiplicity two, and then will remove annulus bundles and any disc bundles that intersect ∂Y_i in their vertical boundary without changing (Y_i, A_i) up to isotopy in M.

Lemma 3.3.7. *Let B be an annulus bundle or Möbius band bundle component of $\mathcal{B}$. Each component of the boundary of $\partial_v B$ intersects at most one boundary component of A_i.*

Proof. As $\partial_h B$ is a collection of incompressible annuli and is a subset of A_i, the core curve of each component of $\partial_h B$ is isotopic to the core curve of A_i. If a component of the boundary of $\partial_v B$ intersects both boundary components, the arc that intersects both components cuts A_i into a disc, so intersects the core curve of some component of $\partial_h B$, which is a contradiction. □

We will use this lemma to show that, if we remove such a parallelity bundle component and we know that this does not change Y_i up to isotopy in M, then it preserves the pair (Y_i, A_i) up to isotopy in M.

Lemma 3.3.8. *If B is a component of the generalised parallelity bundle that is a Möbius bundle or annulus bundle, and some component of $\partial_v B$ intersects the boundary of Y_i in its interior but is not a subset of ∂Y_i, then we can remove handles from the*

handle structure of (Y_i, A_i) *such that this component is a subset of* ∂Y_i *without altering the resulting submanifold up to isotopy in* M*.*

Proof. If a component α of $\partial_v B$ is not properly embedded but is also not contained in ∂Y_i, we can decompose this annulus as follows. Since $\partial_h B$ is incompressible and not a disc, $\partial_v B$ is incompressible so as α is in a solid torus it is boundary-parallel. Consider the components of its intersection with each handle. These are discs with four sides: two that are its intersection with A_i, and the other two that may be in the interior of the solid torus or may be in $\partial S_i - A_i$. Colour each of these squares by whether their interior is in ∂Y_i or not. Then consider the tiling of α where we merge adjacent squares of the same colour (noting that, as α is neither properly embedded nor contained in the boundary, the whole annulus is not of only one colour). Let D be one of these tiles of $\partial_v B$ whose interior is in the interior of the solid torus. Note that D is properly embedded: two of its edges are contained in A_i and two are in $\partial Y_i - A_i$. This disc cannot be a meridian disc for Y_i as then the singular fibre would have multiplicity zero, so it must be boundary-parallel – that is, ∂D bounds a disc in ∂Y_i. We thus have three possibilities for how ∂D can lie with respect to (Y_i, A_i) up to homeomorphism, shown in Figure 3.6(a)-(c). As the two arcs of ∂D in A_i are parts of $\partial_v B$ and so do not span A_i, they each start and end on one component of ∂A_i. They may intersect the same component or two different components, and in the case when they intersect the same component, they may be nested or not. Up to homeomorphism, the two arcs in $\partial Y_i - A_i$ are determined by the fact that ∂D bounds a disc in ∂Y_i.

Consider the 3-ball R that D bounds. This 3-ball intersects B only in D: if they intersected elsewhere, then since B is connected and D is separating in M, the entirety of B would be contained in a 3-ball, but then its horizontal boundary would be compressible. We can thus remove this 3-ball without affecting (Y_i, A_i) up to isotopy in M, by replacing (Y_i, A_i) with $(Y_i - R, A_i - R)$ in the case of Figure 3.6(a) and (b), or with $(S_i - R, (A_i - R) \cup D)$ in the case of Figure 3.6(c). We have reduced the number of squares in the tiling of α, so by induction we can repeat this process until α is a subset of ∂S_i. □

The cell structure of (Y_i, A_i) after this process may not satisfy Convention 2.2.2 and so may not be a handle structure. We will correct this when we remove the parallelity bundle component B.

Lemma 3.3.9. *The maximal generalised parallelity bundle has no Möbius band bundle components.*

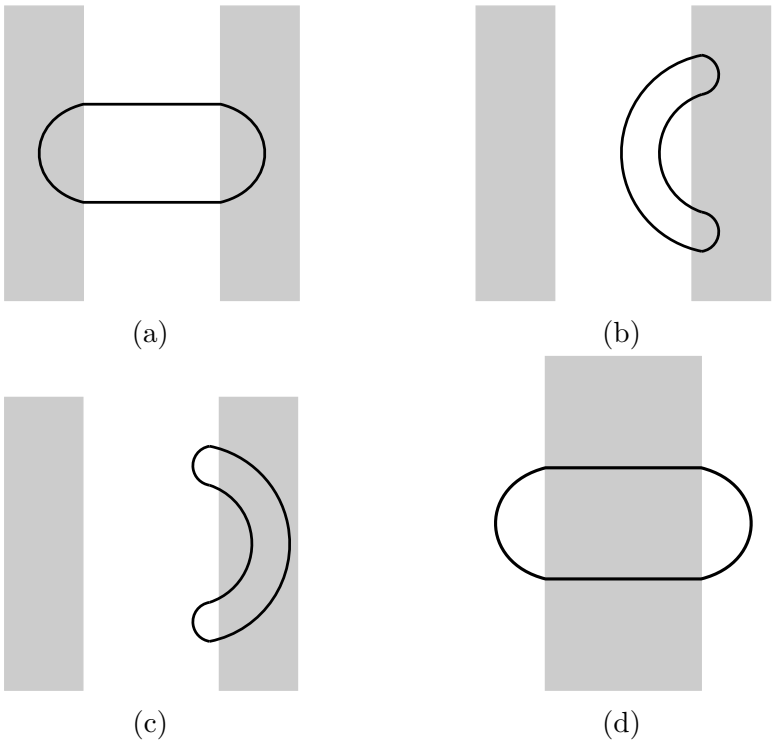


Figure 3.6: The possible ways the boundary of a properly embedded disc in $\partial_v B$ could intersect $(\partial S_i, A_i)$. These diagrams show subsets of ∂S_i. The shaded section is A_i. The black curve is ∂D.

Proof. Suppose that B were such a component. Consider the incompressible and hence boundary-parallel annulus that constitutes its vertical boundary, $\partial_v B$.

If $\partial_v B$ intersects ∂Y_i, then by Lemma 3.3.8 we can assume it is a subset of ∂Y_i. Then, as its interior is disjoint from A_i and its remaining boundary components are subsets of A_i, it must be the entirety of $\partial Y_i - A_i$. Now, as the boundary of B is the entirety of ∂Y_i, Y_i is the neighbourhood of a singular fibre of multiplicity 2, so we have a contradiction.

If $\partial_v B$ is properly embedded, we can isotope it to be vertical with respect to the Seifert fibration of Y_i that is induced from M. Let D_p be the orbifold that is a disc with one cone point of angle $\frac{2\pi}{p}$, recalling that p is the multiplicity of this singular fibre. The Seifert fibration of M gives us an orbifold fibration $Y_i \to D_p$ with S^1-fibres. Now $\partial_v B$ is the preimage of a properly-embedded arc under this map and A_i is the preimage of a boundary arc. The arc from $\partial_v B$ has its endpoints on the arc from A_i, and the union of the arc from $\partial_v B$ and the connected segment of the arc from A_i cut out by these endpoints is the boundary of a sub-disc of D_p. If this sub-disc does not contain the cone point, then the 3-manifold covering it is a solid torus whose meridian

disc is bounded by one arc from A_i and one from $\partial_v B$: this is not a Möbius band bundle. If this sub-disc *does* contain the cone point, then a meridian disc of the solid torus covering it is a $2p$-gon whose boundary alternates arcs from $\partial_v B$ and A_i. To be a Möbius band bundle, p must be 2, which is a contradiction. □

Lemma 3.3.10. *Let $\mathcal{H}$ be a handle structure for (Y, A) where Y is a 3-manifold and A is a subsurface of its boundary. Let B be a component of a maximal generalised parallelity bundle with respect to A. Then $\mathcal{H} - B$ satisfies Convention 2.2.2 and so is a handle structure.*

Proof. As the intersections of handles in $\mathcal{H} - B$ are a subset of those in $\mathcal{H}$, it is enough to show that each k-handle H (with product structure $D^k \times D^{3-k}$) intersects the handles of lower index in exactly $\partial D^k \times D^{3-k}$. Certainly H intersects the handles of lower index in a subset of this set. If a handle adjacent to H in $\mathcal{H}$ was in B, then this had to have been of higher index than H, as otherwise H would also have been in B. Thus we did not remove any intersections between H and lower index handles. □

Lemma 3.3.11. *There is a subnormal neighbourhood Y_i' of γ_i with a distinguished annulus A_i' in $\partial Y_i'$, whose maximal generalised parallelity bundle has no annulus bundle or Möbius band bundle components, where the induced handle structure for Y_i' from $\mathcal{H}$ is a handle structure for the pair (Y_i', A_i'), and such that (Y_i', A_i') and (Y_i, A_i) are isotopic in M.*

Proof. By Lemma 3.3.9 we know that any subnormal handle structure which is isotopic to (Y_i, A_i) in M has no Möbius band bundle components. Suppose that B is some annulus bundle component of the parallelity bundle of (Y_i, A_i). If the interior of $\partial_v B$ intersects ∂Y_i, then by Lemma 3.3.8 we can remove handles from (Y_i, A_i) without changing the resulting submanifold up to isotopy in M. By induction, we can assume that if the interior of $\partial_v B$ intersects ∂Y_i, then one of the components α of $\partial_v B$ is contained in ∂Y_i. As in the Möbius bundle case, this component is the entirety of $\partial Y_i - A_i$. The other component β of $\partial_v B$ must be properly embedded, as it is disjoint from α and its interior is disjoint from A_i. We thus have one properly embedded component of $\partial_v B$, and the other is either properly embedded or a subset of the boundary.

Again as in the Möbius bundle case, we can isotope both A_i and each component of $\partial_v B$ to be vertical in the Seifert fibration; that is, so that they are preimages of arcs in D_p under the projection $Y_i \to D_p$. The two components of $\partial_v B$ project to parallel arcs. The annulus A_i projects to an arc in the boundary of D_p.

If a component of $\partial_v B$ is contained in the boundary, cut along the other component. We have two resulting pieces: the parallelity bundle, which we discard, and a piece Y_i' where $(Y_i', A_i \cap Y_i')$ is isotopic to (Y_i, A_i) in M.

Otherwise, suppose both components are properly embedded. In D_p, one of the two parallel arcs from $\partial_v B$ bounds a disc together with a segment of A_i that does not contain the other arc of $\partial_v B$. This disc may or may not contain the cone point. (The area between the two arcs cannot contain the cone point, as then the multiplicity of the cone point would have to be 2, which it is not.) If this disc disc does contain the cone point, cut along the innermost arc (with respect to the cone point), which has the same effect as in the previous case, except now one piece of the complement is the parallelity bundle glued to a solid torus along an annulus. The other piece, Y_i', has the same core curve (when included into M) as Y_i, and $(Y_i', A_i \cap Y_i')$ is isotopic to (Y_i, A_i) in M. If the disc does not contain the cone point, we cut along the innermost arc with respect to the cone point, removing the parallelity bundle glued to a solid torus, and once again can see that the other piece, Y_i', has the same core curve as Y_i, and $(Y_i', (A_i \cup \partial_v B) \cap Y_i')$ is isotopic to (Y_i, A_i) in M, so we set the distinguished annulus in Y_i' to be $(A_i \cup \partial_v B) \cap Y_i'$.

As the number of handles is reduced by this operation, we can recompute the parallelity bundle and repeat it until there are no annulus bundle components of $\mathcal{B}$ remaining. After performing these modifications we have a handle structure by Lemma 3.3.10. □

The remaining components of the generalised parallelity bundle $\mathcal{B}$ in (Y_i, A_i) are disc bundles.

Lemma 3.3.12. *If B is a disc bundle component of the maximal generalised parallelity bundle such that its vertical boundary intersects ∂Y_i in its interior, then there is a properly embedded disc, which is a subset of $\partial_v B$, such that cutting along this disc separates a 3-ball from Y_i that contains B and does not change (Y_i, A_i) up to isotopy in M.*

Proof. Suppose $\partial_v B$ intersects ∂Y_i in its interior. The annulus $\partial_v B$ cannot be a subsurface of ∂Y_i: if it were, then as it intersects A_i in its boundary only, it would have to have the same boundary curves as A_i, but boundary curves of A_i do not bound discs (i.e. candidates for $\partial_h B$) in A_i. Thus, as in the proof of Lemma 3.3.8, we can divide $\partial_v B$ into squares according to whether a region of its interior is in ∂Y_i or in the interior of Y_i. Again as in the previous case, since p is not zero or two, each of these squares is boundary-parallel.

Unlike in the previous case, the boundaries of the properly embedded squares may be nested in ∂Y_i. Consider a square D whose boundary is to an outermost such curve in ∂Y_i. Suppose that B is not contained in the ball D bounds. In this case, B intersects this ball only in the square D, and as ∂D was an outermost curve, B intersects the ball bounded by each properly-embedded disc in $\partial_v B$ in only the properly-embedded disc itself. We can thus see that $\partial_v B$ is isotopic in Y_i to an annulus in ∂Y_i that is disjoint from the rest of B, and in particular from $\partial_h B$ (though possibly not from A_i). But then both boundary curves of $\partial_h B$ bound discs in ∂Y_i, which is impossible.

Cutting along D removes a 3-ball R from Y_i, which contains the parallelity bundle component B. There are four possible ways, up to homeomorphism, for ∂D to lie in $(\partial Y_i, A_i)$, which are shown in Figure 3.6. Replace (Y_i, A_i) with $(Y_i - R, A_i - R)$ in the case of Figure 3.6(a) or (b), and with $(Y_i - R, (A_i - R) \cup D)$ in the case of Figure 3.6(c) or (d). This does not change (Y_i, A_i) up to isotopy in M. □

There may be parallelity bundle components remaining that we cannot remove without changing (Y_i, A_i) up to isotopy in M. We will replace them with 2-handles to obtain a handle structure for (Y_i, A_i) that is $(0, 0)$-nicely embedded in M; in particular, that has no parallelity bundle. We will then be able to find a core curve for Y_i that runs a bounded number of times through each 0- or 1-handle of M and thus find one that is simplicial in a suitable barycentric subdivision, using the combinatorial bounds of Theorem 2.8.4.

Proposition 3.3.13. *By removing handles of $\mathcal{H}_i$ and replacing the maximal parallelity bundle with 2-handles, we can obtain a handle structure $\mathcal{G}$ for the pair (Y_i', A_i') that is $(0, 0)$-nicely embedded subnormal in $\mathcal{H}$ and is isotopic to (Y_i, A_i) in M. In particular, it has no parallelity 0- or 1-handles.*

Proof. By Lemma 3.3.9 the maximal parallelity bundle has no Möbius band bundle components. If there is an annulus bundle, apply Lemma 3.3.11 to remove it; if there is a disc bundle whose vertical boundary intersects ∂Y_i, apply Lemma 3.3.12 to remove annulus bundles and disc bundles whose vertical boundaries intersect ∂Y_i. The resulting handle structure satisfies Convention 2.2.2 by Lemma 3.3.10. The remaining generalised parallelity bundle is a collection of disc bundles with properly-embedded vertical boundary.

Replace these parallelity bundle components with 2-handles and let $\mathcal{H}_i'$ be the resulting collection of handles. So long as this is a handle structure, it is immediate that it satisfies the definition of being $(0, 0)$-nicely embedded subnormal in $\mathcal{H}$. As any handle that is adjacent to the maximal parallelity bundle $\mathcal{B}$ via a parallelity handle of

lower index is itself in $\mathcal{B}$, and the vertical boundary of $\mathcal{B}$ is properly embedded, these new 2-handles $D^2 \times D^1$ intersect handles of lower index in all of $\partial D^2 \times D^1$. They are also disjoint from other 2-handles, and any components of their intersection with 1-handles come from components of intersection of 2-handles and 1-handles in $\mathcal{H}_i'$, which is a handle structure, so $\mathcal{H}_i'$ is a handle structure satisfying Convention 2.2.2. □

An appropriate core curve

Proposition 3.3.14. *For each subnormal neighbourhood Y_i of a non-multiplicity-two singular fibre with handle structure $\mathcal{H}_i$, we can find a core curve satisfying the conditions in Theorem 3.3.5 that is disjoint from the parallelity bundle of Y_i with respect to the annulus A_i.*

Proof. Let $\mathcal{G}_i$ be the handle structure from Proposition 3.3.13. We know that $\mathcal{G}_i$ has no parallelity handles, its 0- and 1-handles are parts of 0- and 1-handles of $\mathcal{H}_i$, and $\mathcal{G}_i$ is a solid torus neighbourhood of γ_i. A core curve of $\mathcal{G}_i$ thus includes into $\mathcal{H}_i$ as a core curve. The intersection count conditions in Theorem 3.3.5 were satisfied by $\mathcal{H}_i$ as it came from cutting along a normal surface in a handle structure dual to a triangulation. These conditions are also satisfied in $\mathcal{G}_i$: its 0- and 1-handles are parts of 0- and 1-handles of $\mathcal{H}_i$ respectively, so it suffices to check the intersection conditions on the 2-handles. After replacing generalised parallelity bundle components with 2-handles, all intersections between 1-handles and these new 2-handles include into $\mathcal{H}_i$ as components of intersection between the 1-handles and 2-handles. Thus Theorem 3.3.5 applies; that is, there is a core curve γ for Y_i that, in $\mathcal{G}_i$, has the properties from Theorem 3.3.5, and so in particular is disjoint from the 2-handles in $\mathcal{G}_i$ and hence from the parallelity bundle of $\mathcal{H}_i$. □

Proof of Theorem 3.3.2. As we have already noted, the solid torus case is Theorem 6.14 of [LS22]. Otherwise, recall that $\mathcal{H}$ is dual to the triangulation $\mathcal{T}$ of M. Consider the collection of core curves $\gamma = \{\gamma_i\}$ from Proposition 3.3.14: this collection consists of one curve isotopic to each non-multiplicity-two singular fibre. These core curves are contained in the 0- and 1-handles in the handle structures $\mathcal{G}_i$, which are simultaneously $(0,0)$-nicely embedded subnormal in $\mathcal{H}$. By Theorem 2.8.5, the curves $\{\gamma_i\}$ are simultaneously simplicial in $\mathcal{T}^{(79)}$. As we discussed after the statement of Theorem 2.8.5, these simplicial curves contain only edges of $\mathcal{T}^{(79)}$ that are not edges of $\mathcal{T}^{(78)}$. In the solid torus case we had a simplicial core curve in $\mathcal{T}^{(51)}$, so we can take 27 further barycentric subdivisions, each time, for each edge of the simplicial core curve, replacing it with two edges which were created in the barycentric subdivision.

In a barycentric subdivision, the link of an edge e that was created in the subdivision is a circle of four or six edges. When we take an additional barycentric subdivision, each of these edges is subdivided once, so we have a circle of eight or twelve edges. After two further barycentric subdivisions, the simplicial neighbourhoods of these simplicial singular fibres are solid tori and do not intersect in their interior; one more suffices to ensure that they are disjoint. Thus, if we take three barycentric subdivisions to obtain $\mathcal{T}^{(82)}$, we ensure that the simplicial neighbourhoods of the singular fibres are disjoint solid tori and that transverse to each edge of each singular fibre there is a meridian curve of length at most $6 \times 2^3 = 48$. □

3.3.2 Bounds on triangulation complexity

We will now prove the triangulation complexity bounds. First, we give the upper bound.

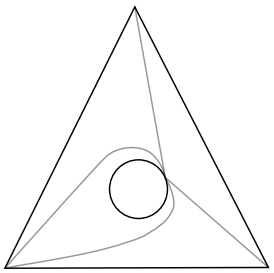


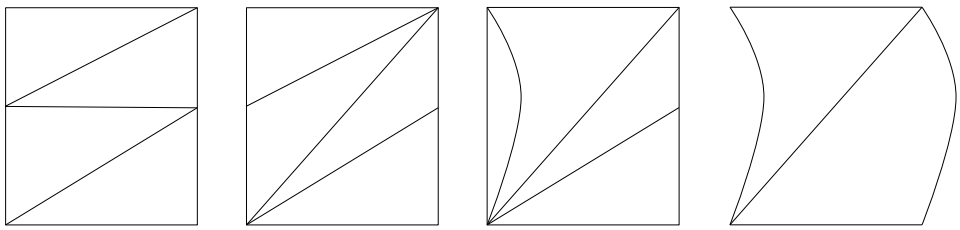
Figure 3.7: Subdividing a triangle to add a new boundary component to a surface triangulation. The edges in grey are new internal edges; the central circle is the new boundary component.

Proposition 3.3.15. *Let M be a manifold with Seifert data $[\Sigma, (p_1, q_1), \ldots, (p_n, q_n)]$ with non-empty boundary. Its triangulation complexity is bounded by*

$$\Delta(M) \leq 96|\chi(\Sigma)| + 176 + 70 \sum_{i=1}^{n} \|q_i/p_i\| .$$

Proof. Let b be the number of boundary components of Σ, and let a be, if Σ is orientable, twice its genus, and otherwise let it be its non-orientable genus. We have $\chi(\Sigma) = 2 - a - b$. Write Σ_n for the surface Σ with n additional boundary components added.

Suppose that Σ_n is orientable. When Σ_n is an annulus or disc, we can construct a triangulation of size three – in particular, of size at most $6a - 1 + 4(b + n)$ – with



(a) (b) 2-2 move (c) 2-2 move (d) 3-1 move

Figure 3.8: Applying three Pachner moves to the triangulation of a boundary torus of $\Sigma_n \times I$ to reduce it to a single vertex triangulation where two of the edges are the fibre and section of the product structure.

one vertex on each boundary component. Otherwise, we can triangulate the closed orientable surface of the correct genus with at most $2a - 2$ faces. We can then add a boundary component in the middle of a face at the cost of 3 additional faces, as shown in Figure 3.7, again with one vertex on each boundary component, to give the same bound. Take the product of this triangulation with an interval to obtain a subdivision of $\Sigma_n \times I$ into triangular prisms, where the same triangulation is induced on the top and bottom faces. Glue top to bottom to obtain $\Sigma_n \times S^1$, divided into at most $6a - 1 + 4(b + n)$ triangular prisms.

Otherwise Σ_n is non-orientable. There is a triangulation of the Möbius band using two triangles with two vertices on its boundary, and by adding one more triangle we can ensure that the boundary has a single vertex. We can add a boundary component to such a surface using three additional triangles, as in the orientable case. We can use this triangulation of the non-orientable surface Γ_p with non-orientable genus one and p boundary components, by a total of $3p$ triangles, to construct a cell structure on $\Gamma_p \tilde{\times} I$ using $3p$ triangular prisms: one over each of these triangles. Now, we can glue two of these cell structures, for $\Gamma_{p'} \tilde{\times} I$ and $\Gamma_{p''} \times I$, along two annuli $\gamma \times I$ that sit over one each of their boundary components. Thus we obtain a cell structure for $\Sigma_n \tilde{\times} I$, and we can glue two of these by the identity map on their boundaries to get a cell structure for $\Sigma_n \tilde{\times} S^1$ composed of at most $2(6a - 6 + 4(b + n))$ triangular prisms.

Pick a subdivision of each quadrilateral face into two triangles, then cone each triangular prism to an interior point, subdividing it into eight tetrahedra: one for each of the faces of this triangulation of the boundary of the prism. This procedure gives a triangulation of $\Sigma_n \overset{(\sim)}{\times} S^1$ by at most $16(6a - 1 + 4(b + n))$ tetrahedra. Note that we were free to choose the triangulation of the quadrilaterals in the boundary.

Some of these boundary tori correspond to boundary components of M, while others will be Dehn filled to create singular fibres. We restrict our attention to

the latter class. The induced triangulation on each boundary torus depends on the orientability of Σ_n. We wish it to be a single vertex triangulation, where the fibre and section curves of the product structure are two of the edges. If Σ_n is orientable, this is automatic. If it is non-orientable, the initial triangulation of each boundary component is by taking a strip of two rectangles and identifying opposite sides, then subdividing each rectangle into two triangles in a way we can choose. As shown in Figure 3.8, we can apply three Pachner moves (that is, add three additional tetrahedra) to obtain the required properties. We have used a total of at most $16(6a - 6 + 4(b + n)) + 3n$ tetrahedra.

Consider a toroidal boundary component that we wish to Dehn fill by gluing the meridian of a solid torus to the slope q_i/p_i curve. By Lemma 2.6.3, we can do this using at most $\|q_i/p_i\| + 2$ tetrahedra. Thus $\Delta(M)$ satisfies

$$\Delta(M) \leq 16(6a - 1 + 4(b + n)) + 3n + \sum_{i=1}^{n}(\|q_i/p_i\| + 2)$$

so, noting that $\|q_i/p_i\|$ is at least 1 so $\sum_i \|q_i/p_i\|$ is at least n and that $|\chi(\Sigma)| = |a+b-2|$, we have that

$$\begin{aligned}\Delta(M) &\leq 96a - 16 + 64b + 69n + \sum_{i=1}^{n} \|q_i/p_i\| \\ &\leq |96a + 96b - 192| + 176 + 70 \sum_{i=1}^{n} \|q_i/p_i\|\end{aligned}$$

which gives us the bound. □

We next give the lower bound from the Euler characteristic of the underlying surface, which follows from the rank of first homology.

Lemma 3.3.16. *Let M be a manifold with Seifert data $[\Sigma, (p_1, q_1), \ldots, (p_n, q_n)]$ and non-empty boundary. The triangulation complexity $\Delta(M)$ is at least $\frac{1}{6}(|\chi(\Sigma)| + 1)$.*

Proof. As $\Delta(M) \geq 1$, this bound holds when M is a solid torus, so assume that $\chi(\Sigma)$ is at most 0. Let $\mathcal{T}$ be a minimal triangulation of M. Note that $\mathcal{T}$ has at most $6|\mathcal{T}|$ edges so its first homology group is of rank at most $6|\mathcal{T}|$. We can use van Kampen's theorem to construct $\pi_1(M)$, in the same order as in the proof of Proposition 3.3.15: first, we take $\Sigma_g^{b+n} \overset{(\sim)}{\times} S^1$ and then glue in n solid tori with their meridians glued to the specified slopes. This presentation is given explicitly, for example, in Theorem 2.2.2 of [Bri93]. Abelianising, we find that the first homology has a free summand of rank $|2 - \chi(\Sigma)|$ if Σ is orientable, and $|1 - \chi(\Sigma)|$ otherwise. As $\chi(\Sigma)$ is nonpositive, these are each at least $|\chi(\Sigma)| + 1$. □

We now need to show that $\Delta(M)$ is bounded below in terms of the sum of the continued fraction expansions. The idea is as follows: in $\mathcal{T}^{(82)}$, for each non-multiplicity-two singular fibre we have a simplicial neighbourhood with a meridian curve of length at most 48. Considering the $(0,0)$-nicely embedded subnormal neighbourhood $\mathcal{G}_i$ of this singular fibre with distinguished annulus A_i in its boundary from Proposition 3.3.13, the boundary components of A_i give us a normal curve in the boundary of $\mathcal{G}_i$ of slope q_i/p_i (with respect to some (meridian, longitude) coordinate system on the boundary of the solid torus) which runs through each handle in the boundary at most once. We thus bring our attention to the complexity of $T^2 \times I$ where we have a controlled length curve of a specified slope in each boundary component, which we studied in Lemma 2.8.9. We need two further elementary results.

Lemma 3.3.17. *Let A be a simplicial triangulation of an annulus, composed of $|A|$ triangles. There is a triangulation of $D^2 \times I$, comprising at most $3|A|$ tetrahedra, that induces the triangulation A on the boundary $\partial D^2 \times I$.*

Proof. Consider the cone CA of A. This is homeomorphic to $D^2 \times I$ pinched at a point in the centre (the cone point) and has $|A|$ top-dimensional cells. We have an induced triangulation X of $D^2 \times \partial I$ from the cone $C\partial A$ over ∂A. Take the cone over each component of X, which is homeomorphic to two balls, and glue it to CA by identifying the inclusions of $C\partial A$ into CX and CA. The resulting cell complex is homeomorphic to $D^2 \times I$ and the triangulation it induces on $\partial D^2 \times I$ is A; as it was built by gluing together cones over triangulated surfaces, each 3-cell is a tetrahedron. This cell complex contains $|A| + |\partial A|$ top-dimensional cells, where $|\partial A|$ is the length of the boundary of A, and that $|\partial A|$ is at most $2|A|$ since no triangle of A has all three edges in the boundary. □

Lemma 3.3.18. *Let $\mathcal{T}$ be a material triangulation of a compact 3-manifold M. There is a pre-tetrahedral handle structure $\mathcal{H}$ for M, containing at most $5|\mathcal{T}|$ 0-handles, that does not contain any parallelity handles with respect to ∂M, and such that the induced cell structure on the boundary is the same as the boundary triangulation induced by $\mathcal{T}$.*

Proof. The idea here is to take the dual handle structure to $\mathcal{T}$ and rule out the possibility of any parallelity handles by ensuring that each handle has at most one component of intersection with the boundary. A (2-dimensional) cell structure A on a closed surface S induces a (3-dimensional) cell structure on $S \times I$, which we denote as $A \times I$, whose induced boundary cell structure on each boundary component is

a copy of A, and where the intersection of each cell with a boundary component is connected. We can consider our triangulation $\mathcal{T}$ to be a cell structure. Let A be the induced cell structure on ∂M from $\mathcal{T}$. Let $\mathcal{C}$ be constructed from the cell structure $\mathcal{T}$ by attaching a copy of $A \times I$ to $\partial\mathcal{T}$. Let $\mathcal{H}$ be the dual handle structure to $\mathcal{C}$, as in Definition 2.2.4. The only handles that intersect the boundary are those added by attaching $A \times I$, each of whose cells intersect each boundary component of $A \times I$ in a connected subsurface, and hence intersect ∂M in at most one component. As the triangulation of ∂M from $\mathcal{T}$ contained at most $4|\mathcal{T}|$ faces, $\mathcal{C}$ has at most $5|\mathcal{T}|$ 3-cells, so $\mathcal{H}$ has at most $5|\mathcal{T}|$ 0-handles.

Now, $\mathcal{H}$ is pre-tetrahedral. The interior 0-handles are dual to the interior cells of $\mathcal{C}$ which are tetrahedra, so these 0-handles have tetrahedral boundary graphs. Those that intersect the boundary are semi-tetrahedral: they are dual to cells that are the product of a triangle and an interval, where one of the triangular faces is on the boundary. □

We can now prove the lower bound when there are no multiplicity two fibres.

Proposition 3.3.19. *Let M be a manifold with Seifert data $[\Sigma, (p_1, q_1), \ldots, (p_n, q_n)]$ other than the solid torus where the first m singular fibres in this list are not of multiplicity two. Then $\Delta(M) \geq \frac{1}{6}m$, and there is a universal constant $k > 0$ such that*

$$\Delta(M) \geq \frac{1}{k}\left(\sum_{i=1}^{m} \|q_i/p_i\| - 9794m\right).$$

Proof. Let $\mathcal{T}$ be a triangulation of M, and let $\mathcal{H}$ be the dual handle structure. Consider, for each singular fibre γ_i that is not of multiplicity two, the $(0,0)$-nicely embedded subnormal handle structure $\mathcal{G}_i$ in M for the subnormal neighbourhood (Y_i, A_i) of γ from Proposition 3.3.13, which exists as M is not a solid torus. Let $\mathcal{C}$ be the disjoint union of these handle structures.

Each subnormal neighbourhood $\mathcal{G}_i$ contains at least one non-parallelity 0-handle, and each 0-handle of $\mathcal{H}$ contains at most six such possible subnormal 0-handles. Thus, as each 0-handle in $\mathcal{H}$ is dual to a tetrahedron in $\mathcal{T}$, $m \leq 6|\mathcal{T}|$.

As the subnormal neighbourhoods are simultaneously $(0,0)$-nicely embedded, in $\mathcal{T}^{(30)}$ we can arrange that the following parts of $\mathcal{C}$ are simplicial using Theorem 2.8.4: the 0-handles, the cocores of the 1-handles, and the attaching annuli of the 2-handles. We already know from Theorem 3.3.2 that after an additions 82 barycentric subdivisions, a core curve of each of these solid tori is simplicial and that these simplicial

curves have disjoint solid tori as their simplicial neighbourhoods, with a simplicial meridian curve of length at most 48 for each such neighbourhood.

Subdivide one more time, giving $\mathcal{T}^{(113)}$, so that the simplicial neighbourhood of each core curve is disjoint from the attaching annuli of the 2-handles of $\mathcal{C}$. Doing this doubles the degree of each edge, so increases the length of the simplicial meridian curve to at most 96, and ensures that the simplicial neighbourhoods are disjoint from $\partial\mathcal{G}_i$ for each i. The induced triangulations of the attaching annuli in $\mathcal{T}^{(113)}$ consist of subdivided triangles from $\mathcal{T}^{(30)}$. Each of these faces is subdivided into six triangles in each barycentric subdivision, so the attaching annuli triangulations contain a total of at most $6^{83} \times 4 \times 24^{30}|\mathcal{T}|$ triangles. By Lemma 3.3.17 we can triangulate the 2-handles (although not necessarily as a subset of $\mathcal{T}^{(113)}$) using at most $3 \times 6^{83} \times 4 \times 24^{30}|\mathcal{T}|$ additional tetrahedra, so we can triangulate $\mathcal{C}$ using at most $2 \times 24^{113}|\mathcal{T}|$ tetrahedra. Consider a subnormal neighbourhood (Y_i, A_i) of a core curve, which we have triangulated. The annulus A_i is simplicial, as it was a subset of the cell structure of the boundary, and so its boundary curves are simplicial. Take one of these. This curve has slope (q_i, p_i), the relevant Seifert data for that singular fibre for some choice of (meridian, longitude) coordinates on the boundary of the solid torus.

Drill out a simplicial neighbourhood of each of the core curves, which we recall is disjoint by construction from the attaching annuli of the 2-handles as well as from $\partial\mathcal{G}_i = \partial S_i$ and has a meridian of length at most 96. We now have a copy of $T^2 \times I$ for each singular fibre, triangulated using a total of at most $2 \times 24^{113}|\mathcal{T}|$ tetrahedra, each with coordinates such that they have a $(1, 0)$ curve of length at most 96 on one side and, on the other, a simplicial (q_i, p_i) curve. By Lemma 3.3.18 there is a pre-tetrahedral handle structure for this collection of copies of $T^2 \times I$, inducing the same triangulation of the boundary, that has no parallelity handles and contains at most $10 \times 24^{113}|\mathcal{T}|$ 0-handles. Each 0-handle contributes at most 1 to the tetrahedral complexity of the handle structure so the tetrahedral complexity of the collection is at most $10 \times 24^{113}|\mathcal{T}|$. By Lemma 2.8.9 we thus have

$$|\mathcal{T}| \geq \frac{1}{10 \cdot 24^{113}k'} \left(\sum_{i=1}^{m} \|q_i/p_i\| - 9794m \right)$$

for some $k' > 0$. □

Corollary 3.3.20. *Let M be a manifold with Seifert data $[\Sigma_g^b, (p_1, q_1), \ldots, (p_n, q_n)]$ where the first m singular fibres in this list are not of multiplicity two. There is a*

universal constant k such that

$$\Delta(M) \geq \left(\frac{1}{24} - \frac{9794}{k}\right) m + \frac{1}{24}|\chi(\Sigma)| + \frac{1}{k}\sum_{i=1}^{m} \|q_i/p_i\| + \frac{7}{24}.$$

Proof. Let $\mathcal{T}$ be a triangulation of M. From Lemma 3.3.16 and Proposition 3.3.19 we have four different lower bounds for $|\mathcal{T}|$ including the trivial one: that it is at least one. The maximum of these is certainly at least the average of the four, so as we can set k to also be at least one,

$$|\mathcal{T}| \geq \frac{1}{4}\left(\frac{m}{6} + \frac{1}{6}(|\chi(\Sigma)| + 1) + \frac{1}{k}\sum_{i=1}^{m} \|q_i/p_i\| - \frac{9794}{k}m + 1\right)$$

which gives the bound. □

Proof of Theorem 3.3.1. The following upper bound is given in Proposition 3.3.15:

$$\Delta(M) \leq 96|\chi(\Sigma)| + 176 + 70\sum_{i=1}^{n} \|q_i/p_i\|.$$

The lower bound when there are no singular fibres of multiplicity two follows from Proposition 3.3.19. To extend our bound to multiplicity two singular fibres, we take a double cover to remove all but at most one of them.

Consider the degree two cover of M constructed as follows: M is a circle bundle over an orbifold S. The multiplicity two singular fibres are the preimages of degree two cone points. Pair these cone points (with at most one left over) by picking some disjoint family of arcs that begin and end on degree two points. Let s be the number of arcs. There is a homomorphism $\pi_1(S) \to \mathbb{Z}/2$ by taking a loop to its intersection count with this family of arcs mod 2. The degree two cover $\widetilde{S} \to S$ associated to the index 2 kernel of this homomorphism induces a horizontal degree two cover $\widetilde{M} \to M$ on the level of 3-manifolds. (Equivalently, consider the collection of Klein bottles in M that is the preimage of the arcs in S and take the corresponding degree two cover.)

Which Seifert fibred manifold is $\widetilde{M}$? Using the Riemann-Hurwitz formula, the base orbifold $\widetilde{S}$ of $\widetilde{M}$ has Euler characteristic $2\chi(S) - \sum_{i=1}^{2s} 1 = 4 - 2a - 2b - 2s$, where a is the non-orientable genus or twice the orientable genus, according to whether S is non-orientable or not, and b is the number of boundary components of S. It has $2b$ boundary components, two singular fibres of index $\pm q_i/p_i$ for each non-multiplicity-two singular fibre of $\widetilde{M}$, as well as, if there was a multiplicity two singular fibre left over, two multiplicity two singular fibres. A $-q_i/p_i$ singular fibre, which is a $(p_i - q_i)/p_i$

singular fibre after a change of coordinates, contributes the same continued fraction sum by Lemma 2.6.4. We deduce that if $\widetilde{S}$ is orientable, it has genus $a+s-1$, or if it is non-orientable, it has $2a+2s-2$ crosscaps. Sort the list of singular fibres of M was so that those with multiplicity two are at the end, and let m be the number of non-multiplicity-two singular fibres.

When $a+b-2 \geq 0$, $|s+a+b-2| = |s|+|a+b-2|$. Otherwise, in which case Σ is a disc, $|s+a+b-2| = |s|-1$ so either way $|s+a+b-2|+2 \geq |s|+|a+b-2|$. We can take k to be large enough that $\frac{1}{24} - \frac{9794}{k}$ is at least $\frac{1}{48}$. Using these two bounds, we can modify the lower bound on $\Delta(\widetilde{M})$ from Corollary 3.3.20 in the following way:

$$\begin{aligned}
\Delta(\widetilde{M}) &\geq \left(\frac{1}{24} - \frac{9794}{k}\right) 2m + \frac{1}{24}|4-2a-2b-2s| + \frac{1}{24}\cdot 4 + \frac{1}{k}\sum_{i=1}^{m} 2\,\|q_i/p_i\| + \frac{3}{24} \\
&\geq 2\cdot\frac{1}{48}m + \frac{1}{12}(|s|+|a+b-2|) + \frac{2}{k}\sum_{i=1}^{m}\|q_i/p_i\| + \frac{3}{24} \\
&\geq \frac{1}{24}m + \frac{1}{12}|\chi(S)| + \frac{1}{12}s + \frac{2}{k}\sum_{i=1}^{m}\|q_i/p_i\| + \frac{3}{24}
\end{aligned}$$

but then each multiplicity two fibre contributes 2 to the sum of $\|q_i/p_i\|$ over all the singular fibres, so $\sum_{i=1}^{n}\|q_i/p_i\| \leq \sum_{i=1}^{m}\|q_i/p_i\| + 4s + 2$. By taking k to be at least 96, we have that $\frac{4}{k}$ is at most $\frac{1}{24}$ and $\frac{8}{k}s$ is at most $\frac{1}{12}s$. We can thus obtain the bound that

$$\Delta(\widetilde{M}) \geq \frac{1}{24}m + \frac{1}{12}|\chi(S)| + \frac{2}{k}\sum_{i=1}^{n}\|q_i/p_i\| + \frac{1}{12}.$$

By lifting a minimal triangulation of M through the double cover, we see that $\Delta(\widetilde{M}) \leq 2\Delta(M)$. Thus we have the result for a suitable choice of k. □

Chapter 4

Recognising 3-manifolds

That works very well in practice, but how does it work out in theory?

Attribution unknown

4.1 Deciding n-manifold homeomorphism

Given two (compact, orientable) n-manifolds, can we tell if they are homeomorphic? That is, is there an algorithm to solve "n-manifold homeomorphism"?

The answer to this question depends on the dimension. For $n = 1$, there is only one compact n-manifold, so the algorithm that always outputs "yes" suffices. For $n = 2$, by the classification of surfaces, a 2-manifold is determined by its homology and its number of boundary components [Bra21]. Considering higher dimensions, for $n = 4$ the problem is undecidable, due to Markov [Mar58] – the simplest 4-manifold that is known to be unrecognisable is $\#_9 S^2 \times S^2$ [Tan23]. For $n \geq 5$, even the n-sphere is known to be unrecognisable (while the recognition problem for S^4 is open) [VKF74, appendix by Novikov].

For $n = 3$ there is an algorithm to recognise when two 3-manifolds are homeomorphic, but its construction required Perelman's resolution of the geometrisation conjecture [Per02; Per03a; Per03b]. Proofs of this fact when the manifolds are closed have been given by Sela, Scott-Short, Matveev, Bessières-Besson-Boileau-Maillot-Porti and Kuperberg, among others [Sel95; SS14; Mat07; Bes+10; Kup19]. Several authors state that Jaco gave a proof in the case with boundary in a set of lectures given at the Peking Summer School in July 2005, but unfortunately his lecture notes are no longer available[1].

[1] Personal correspondence with W. Jaco.

How hard is 3-manifold homeomorphism? To make this question precise, we need to pin down our input format. As the input to the homeomorphism problem, we use triangulations (see Section 2.2). As we are in 3-dimensions, and for 3-manifolds the piecewise-linear, smooth, and topological categories are isomorphic, we can describe every (compact) 3-manifold this way.

Definition 4.1.1. The decision problem of 3-MANIFOLD HOMEOMORPHISM is the problem of, given two triangulations of compact, connected, orientable 3-manifolds with toroidal boundary, deciding if they are homeomorphic.

Perhaps the more natural question is to specialise to closed 3-manifolds, but as geometrisation also applies to 3-manifolds with toroidal boundary, we will see in Section 4.1.2 that the closed case approach works in this setting too.

Question A. *Is* 3-MANIFOLD HOMEOMORPHISM *in* ***NP****?*

Whether we speak of triangulated 3-manifolds or of triangulations in general makes no difference, as the following well-known result shows (see Theorem 3.4 of [Sch11] for one proof):

Proposition 4.1.2. *Deciding whether, given a triangulation, its underlying topological space is a 3-manifold, has a polynomial time solution.*

The idea is that, to show that a triangulation $\mathcal{T}$ gives a closed 3-manifold, we need to show that all points in T have 3-ball neighbourhoods. This is automatic for points in the interior of tetrahedra or faces (as no face is glued to itself) and for points on edges that are not at the midpoint or end of the edge. A neighbourhood of a vertex or the midpoint of an edge is determined by its link, so it suffices to check that these links (of which there are a linear number in $|\mathcal{T}|$) are 2-spheres. Checking whether a (2-dimensional) triangulation is a 2-sphere can be done in polynomial time. For 3-manifolds with boundary, we additionally allow points to have neighbourhoods that are 3-ball hemispheres: that is, their neighbourhoods may also be discs. This does not make the algorithm substantially more complicated.

4.1.1 Dimension 3 in theory and practice

Question A is one precise way of asking: how hard is 3-manifold homeomorphism? From a theoretical perspective, the best known bound was given by Kuperberg [Kup19]. He showed that the problem is elementary recursive, which implies that there is an algorithm for 3-MANIFOLD HOMEOMORPHISM that has running time bounded by

a bounded tower of exponentials in the total number of tetrahedra t in the two triangulations – that is, that it is at most of the order of

$$2^{2^{\cdots^t}}$$

for some fixed height [Kup19] – where the height is not known for the general case. For hyperbolic manifolds, it is bounded by $2^{2^{t^{O(t)}}}$ [Scu21]. While these bounds were not easy to find, they also grow very quickly: with eight tetrahedra, for example, they give no assurance that an algorithm to decide 3-manifold homeomorphism should run in any reasonable amount of time.

In practice, however, we are quite good at deciding 3-MANIFOLD HOMEOMORPHISM. Matveev and Tarkaev produced a census of all 103041 (orientable, closed, irreducible) 3-manifolds up to 13 tetrahedra in 2020 [MT20]. In the same year, Ben Burton tabulated the more than 300 million prime knots of up to 19 crossings, which involves showing that an associated list of 300 million 3-manifolds does not contain duplicates [Bur20]. Both groups work by enumerating all possible manifolds in the class they are considering, then for each pair, either show they are homeomorphic using ad hoc methods or distinguish them using invariants such as homology groups, Turaev-Viro invariants, and the number of subgroups of the fundamental group of a given index (see [Bur11] and [MT20] for further discussion on the challenges of this approach). If the 3-manifold homeomorphism problem took anything like the theoretical running time in practice, these results would have been impossible.

This observation motivates Question A: can we improve our understanding of the complexity class of the problem? A more ambitious question would be whether 3-MANIFOLD HOMEOMORPHISM is in **P**, but even showing that the recognition problem is in **P** for the solid torus would be a significant advancement, as that would prove that unknot recognition is in **P**.

4.1.2 3-manifold homeomorphism: a geometric approach

One algorithm for solving 3-MANIFOLD HOMEOMORPHISM that directly uses geometrisation is, given triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ of two 3-manifolds M_1 and M_2, (1) find the geometric decompositions of $\mathcal{T}_1$ and $\mathcal{T}_2$, (2) check if the geometric pieces are homeomorphic, and (3) check if the gluing maps are equivalent. This is the approach used by Kuperberg, for example [Kup19], and suggests a possible proof that 3-MANIFOLD HOMEOMORPHISM is in **NP**. If M_1 and M_2 are homeomorphic, we want to, in polynomial time in $|\mathcal{T}_1| + |\mathcal{T}_2|$,

1. give the geometric decompositions of M_1 and M_2 and certify that they are in fact decompositions of these manifolds;

2. certify that the geometric pieces are homeomorphic; and

3. certify that the gluing maps are equivalent.

Giving the prime decomposition of the manifolds is possible by Jaco and Tollefson's proof that there is a complete set of reducing spheres consisting of fundamental normal surfaces [JT95, Theorem 5.2]; to give the geometric decomposition using a bounded amount of data is more challenging. We will concern ourselves with the second step: whether the homeomorphism problem is in **NP** for *geometric* 3-manifolds, and whether deciding if a manifold is geometric is in **NP**. The geometrisation theorem tells us that the pieces of the geometric decomposition of a 3-manifold each have one of eight geometries: $\mathbb{H}^3$, Sol, $\mathbb{S}^3$, $\mathbb{S}^2 \times \mathbb{R}$, $\mathbb{E}^3$, Nil, $\mathbb{H}^2 \times \mathbb{R}$, and $\widetilde{SL_2}$, where the latter six of these are Seifert fibered spaces.

As we mentioned in Section 2.1, the recognition problem is already known to be in **NP** for a few classes of manifolds, some of which we will need later.

Theorem 4.1.3 ([Iva08; Sch11])**.** *Recognition of S^3 is in **NP**.*

Theorem 4.1.4 (Theorem 3 [Iva08])**.** *Recognition of B^3, $S^1 \times S^2$, $\mathbb{R}P^3$ and $S^1 \times D^2$ is in **NP**, as is recognition of any fixed connect sum or ∂-connect sum of them.*

Since $S^2 \times S^1$ and $\mathbb{R}P^3 \# \mathbb{R}P^3$ are the only manifolds with the $\mathbb{S}^2 \times \mathbb{R}$ geometry, recognition of this geometry is in **NP**.

Theorem 4.1.5 (Theorem 23 [HH22])**.** *Recognising $\Sigma \times I$ and $\Sigma \tilde{\times} I$ where Σ is a compact surface is in **NP**.*

Theorem 4.1.6 (Theorems 12.1 and 12.2 [LS22])**.** *Recognising if a triangulated 3-manifold is elliptic (the $\mathbb{S}^3$ geometry) is in **NP**, and, given a triangulated elliptic 3-manifold, at least one set of valid Seifert data for it can be certified in polynomial time in the size of the triangulation.*

We will see in Section 4.2 that both the problem of recognising if a manifold is Euclidean or Nil and, that of, given two such manifolds, deciding if they are homeomorphic, are in **NP**. In Chapter 5 we will prove that, similarly, recognising if a manifold with non-empty boundary is Seifert fibered is in **NP**, and that deciding if a given such manifold admits a certain set of Seifert data is in $\mathbf{NP} \cap \text{co-}\mathbf{NP}$. The remaining cases, then, are the closed manifolds with $\mathbb{H}^2 \times \mathbb{R}$ or $\widetilde{SL_2}$ geometries, the hyperbolic manifolds, and the Sol geometry.

4.2 The Euclidean and Nil cases

When considering closed manifolds, two geometries – the Euclidean and Nil ones – are determined purely by their homology and some covering data. This is surely known to experts in the field. We will use this to show that the recognition and homeomorphism problems for closed manifolds with these geometries is in **NP** among all 3-manifolds.

Theorem 4.2.1. *The problem* CLOSED EUCLIDEAN AND NIL RECOGNITION *is in* ***NP****, and deciding if a given closed Euclidean or Nil manifold admits a given set of Seifert data is in* ***NP*** $\cap$ *co-****NP****.*

All closed Euclidean and Nil manifolds are Seifert fibered, and can be described as circle bundles over Euclidean orbifolds (see Scott [Sco83a] or Martelli [Mar16, Prop. 12.3.3] for an exposition). There are only seven such (closed) orbifolds (with isolated cone points), as one can prove by considering the possible lists of integers such that the Euler characteristic of the underlying surface plus $\sum(\frac{1}{p_i} - 1)$ is equal to zero. We describe them by listing the underlying surface and the angles of the cone points:

Surface	Cone points	Surface	Cone points
T^2	None	S^2	$2\pi/2, 2\pi/2, 2\pi/2, 2\pi/2$
K	None	S^2	$2\pi/3, 2\pi/3, 2\pi/3$
S^2	$2\pi/2, 2\pi/4, 2\pi/4$	S^2	$2\pi/2, 2\pi/3, 2\pi/6$
$\mathbb{R}P^2$	$2\pi/2, 2\pi/2$		

Table 4.1: The Euclidean orbifolds.

4.2.1 Reducible toroidal mapping tori

Some Euclidean and Nil manifolds can be viewed as mapping tori of the torus by a reducible mapping class: one that fixes an essential curve. Identify the mapping class group of the torus with $SL(2, \mathbb{Z})$. A reducible mapping class fixes an essential curve, and the mapping class group acts transitively on the set of curves of the torus. Up to conjugacy there are therefore two reducible mapping classes: $\phi_e = \begin{pmatrix} 1 & e \\ 0 & 1 \end{pmatrix}$ and $\bar{\phi}_e = \begin{pmatrix} -1 & e \\ 0 & -1 \end{pmatrix}$. The mapping torus of the torus with monodromy ϕ_e is homeomorphic to the Seifert fibered space $[T^2, e]$, and the mapping torus of $\bar{\phi}_e$ is homeomorphic to the Seifert fibered space $[K, e]$. We can assume that e is nonnegative as reversing the orientation of these circle bundles flips the sign of e.

Let REDUCIBLE TOROIDAL MAPPING TORUS RECOGNITION be the problem to determine if a given 3-manifold is a torus bundle over the circle with reducible monodromy. Let NAMING REDUCIBLE TOROIDAL MAPPING TORI be the problem of, given a triangulated reducible toroidal mapping torus and the Seifert data of another reducible toroidal mapping torus, deciding if the manifold admits that Seifert data.

Theorem 4.2.2. *The problem* REDUCIBLE TOROIDAL MAPPING TORUS RECOGNITION *is in* ***NP****. The problem* NAMING REDUCIBLE TOROIDAL MAPPING TORI *is in* ***NP*** $\cap$ *co-****NP****.*

Our main tool to prove Theorem 4.2.2 will be the homology of torus bundles over the circle.

Proposition 4.2.3. *The homology of a torus bundle over the circle M with monodromy ϕ is the following. We have $H_3(M;\mathbb{Z}) \cong \mathbb{Z}$. The second homology is*

$$H_2(M;\mathbb{Z}) = \begin{cases} \mathbb{Z}^3 & \phi = \mathrm{id} \\ \mathbb{Z}^2 & \phi \textit{ conjugate to } \phi_e,\ e \neq 0 \\ \mathbb{Z} & \textit{else.} \end{cases}$$

If ϕ is not Anosov, then the first homology is

$$H_1(M;\mathbb{Z}) = \begin{cases} \mathbb{Z}^3 & \phi = \mathrm{id} \\ \mathbb{Z}^2 \oplus \mathbb{Z}_e & \phi \textit{ conjugate to } \phi_e,\ e \neq 0 \\ \mathbb{Z} \oplus \mathbb{Z}_2^2 & \phi \textit{ conjugate to } \bar{\phi}_e,\ e \textit{ even} \\ \mathbb{Z} \oplus \mathbb{Z}_4 & \phi \textit{ conjugate to } \bar{\phi}_e,\ e \textit{ odd} \\ \mathbb{Z} \oplus \mathbb{Z}_3 & \phi \textit{ of order } 3 \\ \mathbb{Z} \oplus \mathbb{Z}_2 & \phi \textit{ of order } 4 \\ \mathbb{Z} & \phi \textit{ of order } 6. \end{cases}$$

Proof. Up to conjugacy and taking inverses, the elements of $SL(2,\mathbb{Z})$ are $\mathrm{id} = \phi_0$, ϕ_e, $\bar{\phi}_e$, Anosov maps, and finite order maps of order 2, 3, 4 and 6, which have respective representatives $\bar{\phi}_0 = \begin{pmatrix} -1 & 0 \\ 0 & -1 \end{pmatrix}$, $\begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix}$, $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ and $\begin{pmatrix} 1 & -1 \\ 1 & 0 \end{pmatrix}$.

Let A and B be two copies of $T^2 \times I$, that agree with the fibration of M as a torus bundle over the circle, such that $A \cup B$ covers M and $A \cap B$ is two copies of $T^2 \times I$. Apply the Mayer-Vietoris sequence to compute the homology of $M = A \cup B$. Doing this, we see that the map from $H_1(A \cap B) \cong \mathbb{Z}^2 \oplus \mathbb{Z}^2$ to $H_1(A) \oplus H_1(B) \cong \mathbb{Z}^2 \oplus \mathbb{Z}^2$ is the matrix

$$\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & \phi_{11} & \phi_{12} \\ 0 & 1 & \phi_{21} & \phi_{22} \end{pmatrix}$$

which reduces to

$$\begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & \phi_{11} - 1 & \phi_{12} \\ 0 & 0 & \phi_{21} & \phi_{22} - 1 \end{pmatrix}$$

so has a kernel of the same dimension as the eigenspace of ϕ with eigenvalue 1. This dimension is two if $\phi = \mathrm{id}$, one if $\phi = \phi_e$ for $e \neq 0$, and zero else, which is enough to give us the second homology. For the first homology, we compute the image of this matrix for the specific mapping classes in question. □

The manifold M is irreducible and the only (orientable) essential surfaces in it are vertical tori, unless $e = 0$, in which case there are also horizontal tori.

Lemma 4.2.4. *Suppose M is an irreducible closed 3-manifold that contains an essential torus. Let $\mathcal{T}$ be a triangulation of M. Then there is an essential torus that is either a fundamental normal surface or the double of one.*

Proof. Let F be a normal incompressible torus that is of minimal weight among all normal surfaces that are incompressible tori. Suppose that $F = G_1 + G_2$ (that is, that F is not fundamental). By Theorem 2.3.3, both of G_1 and G_2 is incompressible and of Euler characteristic at most 0. As $\chi(G_1) + \chi(G_2) = \chi(F) = 0$, we have either two incompressible Klein bottles or two incompressible tori. We cannot have tori, as then F was not of minimal weight among the class of incompressible tori. Thus we have two incompressible Klein bottles. As $2F$ is also an incompressible minimal normal surface and $2F = 2G_1 + 2G_2$, the surfaces $2G_1$ and $2G_2$ are also incompressible tori. Now, we must have the weight of $2G_1$ equal to the weight of $2G_2$ equal to that of F, so consider G_1, which is a Klein bottle in M with incompressible double. By the same argument, if $G_1 = H_1 + H_2$, then H_1 and H_2 must be Klein bottles or tori, whose doubles are incompressible. At least one must be a Klein bottle since G_1 is non-orientable, so we obtain a Klein bottle of lower weight than G_1 and thus its double is an incompressible torus of lower weight than that of F, a contradiction. □

Proof of Theorem 4.2.2. Given a triangulation $\mathcal{T}$ of some 3-manifold M, if it is a reducible toroidal mapping torus with Seifert data $[T^2, e]$, we give the following certificate:

1. the Seifert data $[T^2, e]$,
2. a fundamental normal surface (or a double of one) F that is an incompressible torus, given as a vector $v \in \mathbb{Z}^{7|\mathcal{T}|}$,

3. the data from following the algorithm in Proposition 2.7.7 to construct a triangulation $\mathcal{T}'$ of $M\backslash\backslash F$,

4. and a certificate from Theorem 4.1.5 that $\mathcal{T}'$ is homeomorphic to $T^2 \times I$.

The algorithm to verify that $\mathcal{T}$ is $[T^2, e]$ from this certificate is the following. Check that $\mathcal{T}$ is closed. Use Proposition 2.7.9 to check that F is indeed a torus. Use the algorithm of Proposition 2.7.7 to construct a triangulation of $M\backslash\backslash F$ and verify that it is $\mathcal{T}'$. Verify the certificate of Theorem 4.1.5 to prove that $\mathcal{T}'$ is $T^2 \times I$. At this point we know that M is a torus bundle over the circle, and we wish to prove that the monodromy is (up to conjugacy) ϕ_e for some value of e that we will verify. Fortunately for us, Proposition 4.2.3 tells us that we can recognise $[T^2, e]$ among the class of torus bundles over the circle by its homology: for $e = 0$, it is uniquely determined by $H_2(M; \mathbb{Z}) = \mathbb{Z}^3$, and for $e \neq 0$, it is picked out that we are in this class by computing $H_2(M; \mathbb{Z}) = \mathbb{Z}^2$, and then we can detect e (noting that the sign of e is irrelevant) from the degree of the torsion in $H_1(M; \mathbb{Z})$. So the remainder of the certificate is to compute the homology of M, which (as the dimension of M is fixed) can be done in polynomial time by Proposition 2.2.1.

If it is a reducible toroidal mapping torus with Seifert data $[K, e]$ we take a slightly different approach. As K is double covered by T^2, $[K, e]$ is double covered by $[T^2, 2e]$. (Another way of viewing this is that $\bar{\phi}_e^2 = \phi_{2e}$.) The certificate is:

1. the Seifert data $[K, e]$,

2. the second barycentric subdivision $\mathcal{T}'$ of $\mathcal{T}$,

3. a triangulation $\tilde{\mathcal{T}}$ of the double cover $[T^2, 2e]$ of M, along with a simplicial projection map to $\mathcal{T}'$,

4. a certificate that $\tilde{\mathcal{T}}$ is in fact $[T^2, 2e]$ (from the first part of this proof),

5. a homology computation for $H_k(M; \mathbb{Z})$, for $k = 1, 2$.

Check that the projection map of $\tilde{\mathcal{T}}$ is in fact a degree two covering map by checking this on a neighbourhood of each tetrahedron of $\mathcal{T}'$. As the homology computation of $H_1(M; \mathbb{Z})$ shows that it is rank one, we can certify that $\pi_1(M)$ is infinite. As M is finitely covered by a Seifert fibered space, by Theorem 2.5.10 M is Seifert fibered and so in particular ϕ is not Anosov. Among the non-Anosov cases, $H_1(M; \mathbb{Z})$ certifies that ϕ is $\bar{\phi}_e$ for some (unknown) e; that is, that M is $[K, e]$.

As discussed in Proposition 2.5.9 any covering map $\tilde{M} \to M$ between two Seifert fibered spaces is a composition of a horizontal map (which is injective when restricted to a regular fibre of $\tilde{M}$ and is induced by a covering map of the orbifolds) and a vertical map (which induces a homeomorphism between orbifolds). As the second homology computation certifies that M is not a circle bundle over the torus, the covering map from $[T^2, 2e]$ to M must have a non-trivial horizontal component. As it is degree two, and the degree of a covering map is the product of the horizontal and vertical degrees, it must just be a horizontal map. Thus the Euler number of M is half that of $\tilde{M}$, which we have already certified. We have now proved the first part of Theorem 4.2.2.

To prove that NAMING REDUCIBLE TOROIDAL MAPPING TORI is in $\mathbf{NP} \cap$ co-$\mathbf{NP}$, given a triangulation of a manifold $\mathcal{T}$ and a candidate set of Seifert data $[T^2, e]$ or $[K, e]$, take the certificate we have just described for REDUCIBLE TOROIDAL MAPPING TORUS RECOGNITION. This certificate includes the Seifert data of $\mathcal{T}$ and we have just described how to certify it. Now by Lemma 2.5.3, as $[T^2, e]$ and $[K, e]$ are not lens spaces, we can decide whether $\mathcal{T}$ is homeomorphic to the manifold described by the given set of Seifert data or not in polynomial time. □

4.2.2 Recognising closed Euclidean manifolds

The closed Euclidean manifolds are the closed Seifert fibered spaces whose Euler class and orbifold Euler characteristic are both zero, which are all covered by T^3. We consider the closed Euclidean 3-manifolds other than T^3 and $K \tilde{\times} S^1$. There are only five of these, and they are easy to certify. Let CLOSED EUCLIDEAN 3-MANIFOLD RECOGNITION be the problem to determine if a given closed 3-manifold is Euclidean. Let NAMING CLOSED EUCLIDEAN 3-MANIFOLDS be the problem of, given a triangulated closed Euclidean 3-manifold and the Seifert data of another closed Euclidean 3-manifold, deciding if the manifold admits that Seifert data.

Proposition 4.2.5. *The problem* CLOSED EUCLIDEAN 3-MANIFOLD RECOGNITION *is in* ***NP****. The problem* NAMING CLOSED EUCLIDEAN 3-MANIFOLDS *is in* ***NP*** $\cap$ *co-****NP****.*

Proof. Given a triangulation $\mathcal{T}$ of a closed 3-manifold M, if M is Euclidean, the certificate we give (for either problem) is the following:

1. the Seifert data of M,
2. a compatible set of orientations of the tetrahedra of $\mathcal{T}$,
3. the 2^{nd} barycentric subdivision $\mathcal{T}'$ of $\mathcal{T}$,

4. a triangulation $\tilde{\mathcal{T}}$ of T^3, containing at most $6 \times 24^2|\mathcal{T}|$ tetrahedra, a covering map from it to $\mathcal{T}'$, and

5. a certificate that $\tilde{\mathcal{T}}$ is indeed T^3.

First, we show that such a triangulation of T^3 does exist. There are only seven closed Euclidean Seifert fibered structures, corresponding to the seven closed Euclidean orbifolds with no reflections. Table 4.2 gives their Seifert data, first homology group, and the degree of the cover by T^3. From the table, we can see that every closed Euclidean 3-manifold is covered by T^3 with degree at most six.

The certificate that $\tilde{\mathcal{T}}$ is indeed T^3 exists by Theorem 4.2.2. We show that we can verify this certificate in polynomial time in $|\mathcal{T}|$. Verify that the given set of orientations of $\mathcal{T}$ is indeed compatible. By Theorem 4.2.2 we can verify the certificate of T^3. Certify that the covering map is indeed a covering map, and thus that M is an orientable Euclidean manifold. Compute the first homology of M over $\mathbb{Z}$ in polynomial time as described in Proposition 2.2.1. Note that $K \tilde{\times} I$ is homeomorphic to the Seifert fibered spaces with data $[K]$ and $[S^2, 1/2, 1/2, -1/2, -1/2]$. Observe from Table 4.2 that the first homology determines the homeomorphism type of M among closed Euclidean 3-manifolds. We can thus certify the Seifert data. Now to decide NAMING CLOSED EUCLIDEAN 3-MANIFOLDS, by Lemma 2.5.3 we can decide whether this Seifert data gives a homeomorphic manifold to the Seifert data we were originally given in polynomial time. □

Seifert data	$H_1(M;\mathbb{Z})$	degree of cover by T^3
$[T^2]$	$\mathbb{Z}^3$	1
$[K]$	$\mathbb{Z} \oplus \mathbb{Z}_2^2$	2
$[S^2, 1/2, 1/2, -1/2, -1/2]$	$\mathbb{Z} \oplus \mathbb{Z}_2^2$	2
$[S^2, 1/3, 1/3, -2/3]$	$\mathbb{Z} \oplus \mathbb{Z}_3$	3
$[S^2, 1/2, -1/4, -1/4]$	$\mathbb{Z} \oplus \mathbb{Z}_2$	4
$[S^2, 1/2, -1/3, -1/6]$	$\mathbb{Z}$	6
$[\mathbb{R}P^2, 1/2, -1/2]$	$\mathbb{Z}_4$	4

Table 4.2: Homology of Euclidean Seifert fibre structures.

4.2.3 Recognising Nil manifolds

The Nil 3-manifolds have the same base orbifolds as the Euclidean orbifolds, and we will use similar techniques to recognise them. While Euclidean Seifert fibered spaces all have Euler number zero, with Nil manifolds we must also detect this invariant.

Fortunately we can detect it using covers of small degree by reducible toroidal mapping tori. Table 4.3 gives the first homology of the Nil 3-manifolds. The homology groups were computed by abelianising the standard presentation of the fundamental group of a Seifert fibered space and computing its Smith normal form. Let NIL 3-MANIFOLD RECOGNITION be the problem to determine if a given (closed) 3-manifold is Nil. Let NAMING NIL 3-MANIFOLDS be the problem of, given a triangulated (closed) Nil 3-manifold and the Seifert data of another Nil 3-manifold, deciding if the manifold admits that Seifert data.

Seifert data	$H_1(M;\mathbb{Z})$	degree of cover by $[T^2, e']$ for suitable e'
$[T^2, e]$	$\mathbb{Z}^2 \oplus \mathbb{Z}_e$	1
$[K, e]$, e even	$\mathbb{Z} \oplus \mathbb{Z}_2^2$	2
$[K, e]$, e odd	$\mathbb{Z} \oplus \mathbb{Z}_4$	2
$[S^2, 1/2, 1/2, -1/2, -1/2, e]$	$\mathbb{Z}_2^2 \oplus \mathbb{Z}_{4e}$	2
$[S^2, 1/3, 1/3, -2/3, e]$	$\mathbb{Z}_3 \oplus \mathbb{Z}_{9e}$	3
$[S^2, 1/2, -1/4, -1/4, e]$	$\mathbb{Z}_2 \oplus \mathbb{Z}_{16e}$	4
$[S^2, 1/2, -1/3, -1/6, e]$	$\mathbb{Z}_{36e}$	6
$[\mathbb{R}P^2, 1/2, -1/2, e]$	$\mathbb{Z}_4^2$	4

Table 4.3: Homology of Nil Seifert fibre structures. We require $e \neq 0$.

Proposition 4.2.6. *The problem* NIL 3-MANIFOLD RECOGNITION *is in* ***NP****. The problem* NAMING NIL 3-MANIFOLDS *is in* $\boldsymbol{NP} \cap$ *co-*$\boldsymbol{NP}$*.*

Proof. Let M be a (closed) Nil 3-manifold with a triangulation $\mathcal{T}$. We have already seen the result for NIL 3-MANIFOLD RECOGNITION in Theorem 4.2.2 when M is $[T^2, e]$ or $[K, e]$, so suppose M is not one of these. The certificate we give is the following:

1. the Seifert data of M,
2. a compatible set of orientations of the tetrahedra of $\mathcal{T}$,
3. the 2^{nd} barycentric subdivision $\mathcal{T}'$ of $\mathcal{T}$,
4. an oriented triangulation $\tilde{\mathcal{T}}$ of a manifold $\tilde{M}$, containing at most $3 \times 24^2|\mathcal{T}|$ tetrahedra and a covering map from it to $\mathcal{T}'$, such that the homeomorphism type of $\tilde{\mathcal{M}}$ is $[T^2, e']$ where $e' = ke$ for $2 \leq k \leq 6$,
5. a certificate of the homeomorphism type of $\tilde{\mathcal{T}}$.

The required triangulation of $\tilde{M}$ exists as it is a cover of M of degree at most 6. By Theorem 4.2.2, there is a polynomial time certificate that $\tilde{M}$ admits the Seifert data $[T^2, e']$. Given this cover, M must also be Seifert fibered (Theorem 2.5.10). From Table 4.3, a Nil 3-manifold is identified up to its Euler class by its first homology, and, for the manifolds other than $[K, e]$ or $[\mathbb{R}P^2, 1/2, -1/2, e]$, its Euler class is also determined. We are assuming we are not in the $[K, e]$ case. Compute the first homology of M using Proposition 2.2.1, and thus deduce its Seifert data except for the case of $[\mathbb{R}P^2, 1/2, -1/2, e]$. In this final case, we have certified all the Seifert data of the manifold except for the Euler number e. We know that there is a degree four cover by $[T^2, e']$ where we have certified the value of e' and deduced that the cover is horizontal. By Proposition 2.5.9, $e' = e(\tilde{M}) = d_h e(M) = 4e(M)$. We can thus deduce from the certificate that $e(M) = \frac{1}{4}e'$.

Now for the NAMING NIL 3-MANIFOLDS part, take this certificate or, if M is $[T^2, e]$ or $[K, e]$, the certificate of its Seifert data from Theorem 4.2.2. Then by Lemma 2.5.3 we can decide if M admits the given set of Seifert data in polynomial time. □

Proof of Theorem 4.2.1. Given triangulations $\mathcal{T}$ of a closed 3-manifold M, we can certify that M is Euclidean and Nil, and if so, certify its Seifert data using Theorem 4.2.2 and Propositions 4.2.5 and 4.2.6. Then we can decide if M admits the given set of Seifert data in polynomial time (given this certificate) using Lemma 2.5.3. □

Chapter 5

Recognition of Seifert fibered spaces with non-empty boundary

5.1 Summary of results

In this chapter, we will prove that recognising Seifert fibered spaces with non-empty boundary is in **NP** and that the homeomorphism problem for this class of manifolds is in **NP** $\cap$ co-**NP**. As we described in Section 4.1.2, this is a step towards a proof that the 3-manifold homeomorphism problem is in **NP**. The second step of this program is to show that the homeomorphism problem is in **NP** for geometric 3-manifolds. One would then hope to certify whether or not two 3-manifolds are homeomorphic by giving appropriate certificates for their geometric pieces and for the gluing maps. As we have discussed, the recognition problem for closed elliptic, Euclidean and Nil 3-manifolds is in **NP**, as is (given such a manifold) deciding if a given set of Seifert data gives a homeomorphic manifold. In this chapter we give the result for Seifert fibered spaces with non-empty boundary, which leaves closed Seifert fibered spaces, hyperbolic manifolds and Sol manifolds as the remaining cases.

We consider two algorithmic problems. First, the SEIFERT FIBERED SPACE WITH BOUNDARY RECOGNITION decision problem: given a triangulation of some 3-manifold, decide whether it admits a Seifert fibered structure and has non-empty boundary.

Theorem 5.1.1. *The problem* SEIFERT FIBERED SPACE WITH BOUNDARY RECOGNITION *is in* ***NP****.*

Second, the NAMING SEIFERT FIBERED WITH BOUNDARY problem: given a triangulation of a Seifert fibered 3-manifold with non-empty boundary and a set of Seifert data, decide if the manifold is homeomorphic to that Seifert fibration.

Theorem 5.1.2. *The problem* NAMING SEIFERT FIBERED WITH BOUNDARY *is in* $\boldsymbol{NP} \cap$ *co-*$\boldsymbol{NP}$.

From Theorem 5.1.2, we can obtain the following result.

Corollary 5.1.3. *Recognition of torus knot complements is in* $\boldsymbol{NP}$*, as is deciding if a torus knot complement is homeomorphic to the complement of a given torus knot* $T(p, q)$.

The complement of the torus knot $T(p, q)$ in the 3-sphere is the Seifert fibered space $[D^2, -s/q, -r/p]$ where r/s is a fraction such that $ps - qr = 1$. Baldwin and Sivek showed previously that, given a knot diagram, deciding if it is a diagram of a torus knot is in **NP** [BS19]. They compute the Alexander polynomial and signature of the knot in order to determine which torus knot it is. By applying our Seifert fibered space naming algorithm, we obtain the same result in the more general 3-manifold setting. However as we do not keep track of orientations, we cannot distinguish the complements of $T(p, q)$ and $T(-p, q)$.

We will prove Theorems 5.1.1 and 5.1.2 by giving a set of surfaces that carve the manifold M into solid tori. First, we prove in Proposition 5.2.2 that (aside from a few exceptions) M contains a fundamental horizontal surface. Our second goal is to provide a collection of normal vertical annuli in M, of total weight at most an exponential of a polynomial of the size of the triangulation, whose complement is a collection of solid tori. To do this we develop the theory of *split handle structures* in Section 5.3: these give us a variant of normal surface theory that is well-adapted to finding collections of interesting disjoint surfaces in a 3-manifold. We use these normal surfaces to build the certificate, first for recognising circle bundles over surfaces with boundary, and second for Seifert fibered spaces with only multiplicity two singular fibres. We can reduce the general case to these situations by using Theorem 3.3.2 to move to a triangulation where the non-multiplicity-two singular fibres are simplicial, which allows us to drill them out and record their meridian.

One difficulty with extending to the closed case is that, as well as the fact that Theorem 3.3.2 applies only to the case with boundary, in the closed case we must also certify the Euler number of the manifold. This is a great deal more subtle.

5.2 A fundamental horizontal surface

Let M be a Seifert fibered space whose boundary is non-empty, equipped with a triangulation $\mathcal{T}$. In this section we will show that, aside from a short list of exceptional

manifolds, there is a fundamental normal horizontal surface in M whose induced covering of the base orbifold of the Seifert fibration is of minimal degree. To do this, we will use normal surface theory as discussed in Section 2.3 and our understanding of essential surfaces in Seifert fibered spaces from Section 2.5.1.

Lemma 5.2.1. *Suppose that M is a Seifert fibered space with non-empty boundary other than $S^1 \times D^2$, $T^2 \times I$, or $K \tilde{\times} I$. If F is a horizontal surface in M, then $\chi(F)$ is negative.*

Proof. It is enough to show that the Euler characteristic of the base orbifold O of M is negative, as $\chi(F)$ is a multiple of it. We classified the orbifolds with boundary and Euler characteristic zero in Lemma 2.5.8, which have corresponding Seifert fibered spaces $T^2 \times I$ and $K \tilde{\times} I$. For $\chi(O) > 0$, the underlying surface of O must be a disc and O has at most one cone point, so M is the solid torus. □

Proposition 5.2.2. *Suppose that M is a Seifert fibered space with non-empty boundary other than $S^1 \times D^2$, $T^2 \times I$, or $K \tilde{\times} I$. Let $\mathcal{T}$ be a triangulation of M. Let p be the lowest common multiple of the multiplicities of the singular fibres. There is a fundamental horizontal normal surface in M that is a degree p cover of the underlying orbifold. If M is $S^1 \times D^2$, then there is a fundamental normal meridian disc.*

Proof. Lemma 2.3.7 gives the solid torus case, so we can assume that M is irreducible and ∂-irreducible. Consider M to be constructed by taking a circle bundle M' over a surface Σ with non-empty boundary, then gluing n solid tori (that is, neighbourhoods of singular fibres) on along vertical annuli with respect to the fibration of M', on a single boundary component of M'. The meridian of one of these solid tori, containing a (p_i, q_i) singular fibre, intersects the gluing annulus in p_i spanning arcs.

Take a degree p horizontal surface in M': if Σ is orientable, this will be p copies of Σ; otherwise it will be $\lfloor \frac{p}{2} \rfloor$ copies of the double of Σ, in addition to one copy of Σ if p is odd. Either way, its intersection with each vertical annulus in the boundary of M' will be as p spanning arcs. We can thus take $\frac{p}{p_i}$ meridian discs in each of these singular fibre neighbourhoods and attach them to the degree p horizontal surface in M' to form a connected degree p horizontal surface in M. This surface is incompressible (as it is a finite cover of the base orbifold and hence is π_1-injective) and ∂-incompressible (by the same argument on the double of M), and does not contain any trivial spheres or discs, so let F be a minimal normal representative. By Lemma 5.2.1, $\chi(F)$ is negative.

Suppose that F is not fundamental, so $F = G_1 + G_2$. We will show that at least one of G_1 and G_2 is horizontal. By Theorem 2.3.3, G_1 and G_2 are incompressible and

∂-incompressible and not a trivial disc or sphere, so are horizontal or pseudo-vertical by Lemma 2.5.7.

Suppose both G_1 and G_2 are pseudo-vertical, so are either annuli or non-orientable surfaces with one boundary component. As ∂F is not a vertical curve, there must be at least one component of each of ∂G_1 and ∂G_2 on each boundary component of M.

Consider the case when p is odd. As there are no even multiplicity singular fibres, there are no non-orientable pseudo-vertical surfaces so G_1 and G_2 are both vertical annuli and in particular both have Euler characteristic zero. Otherwise p is even so F is orientable. Then $2F$, the double of F as a normal surface vector, is minimal in its isotopy class and is incompressible and ∂-incompressible as it is just two copies of F. As $2F = 2G_1 + 2G_2$, by Theorem 2.3.3 the doubles of G_1 and G_2 are incompressible and ∂-incompressible. As $2G_i$ is orientable, and $\partial(2G_i)$ is two copies of ∂G_i and so is a vertical fibre, each of the surfaces $2G_i$ is orientable and vertical and so has Euler characteristic 0. But as $\chi(2G_i) = 2\chi(G_i)$, each G_i has zero Euler characteristic. Then in either case $\chi(F) = \chi(G_1) + \chi(G_2) = 0$ but we know that $\chi(F)$ is negative.

We can thus assume that G_1 is horizontal. It remains to show that the degree of its induced covering of the base orbifold is p. Now, G_1 intersects the singular fibre neighbourhoods (which were cut out by vertical annuli) in a collection of meridian discs. A meridian disc around a multiplicity p_i fibre intersects the relevant vertical annulus in p_i spanning arcs. We know that up to isotopy G_1 intersects the circle bundle M' as a horizontal surface, and so in particular intersects each vertical annulus in the boundary component of M' along which we glued the singular fibres the same number of times. Thus this number must be a multiple of all of the multiplicities: that is, it is kp for some integer k, recalling that p is the lowest common multiple of the p_i, and $G_1|_{M'}$ is a degree kp cover of Σ. (If $n = 0$, taking the lowest common multiple of the empty set to be 1 by definition, this reasoning holds vacuously.)

Since k is at least one, and $\frac{1}{k}\chi(G_1) = \chi(F)$ which is negative, $\chi(G_1)$ is uniquely maximised when $k = 1$. As $\chi(G_1) \geq \chi(F)$ and this maximum achieves equality, k must be 1 and hence G_1 is a degree p horizontal surface. Thus there is a fundamental such surface. □

5.3 Split handle structures

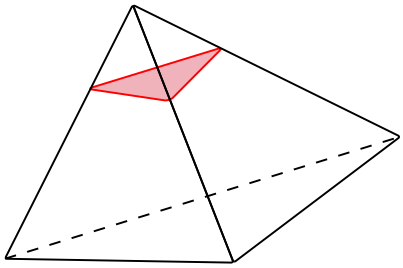
We wish to find a complete collection of normal vertical annuli such that the size of the corresponding normal vector is bounded by a polynomial in $|\mathcal{T}|$. Finding complete collections of surfaces has been considered before: in his doctoral thesis, Saul Schleimer

produced a collection of normal and almost normal fibres in a surface bundle over a circle [Sch01, Theorem 6.2.1] using an inductive argument and the notion of thin position. Our approach is more related to work of Simon King in his doctoral thesis in giving a maximal collection of normal spheres in a 3-manifold [Kin01, §3.2 and §3.3], and also of Marc Lackenby to give a certificate for product sutured manifolds [Lac21, Thm. 12.2]. We will generalise this approach and set it on rigorous foundations by giving a theory of normal surfaces in *split handle structures*. The main result we will obtain is the following. A normal surface in a split handle structure is *duplicate* (see Definition 5.3.18) if two of its components are normally isotopic or if one is isotopic into the forbidden region. The *size* of a normal surface F in a split handle structure (see Definition 5.3.6) is the number of components of the intersection of F with the 0-handles.

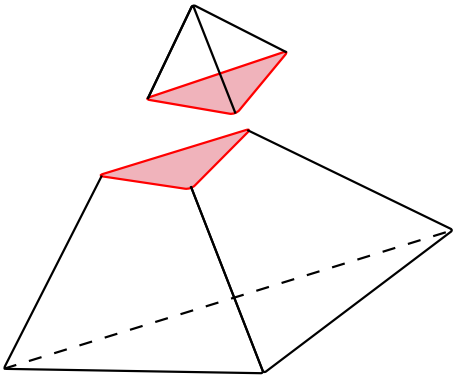
Corollary 5.3.35. *There exists a constant c_S, which we can take to be 10^{10^7}, such that the following holds. Let M be a manifold with a subtetrahedral split handle structure $\mathcal{H}$. Let $\{\Sigma_i\}$ be a collection of n disjoint normal surfaces in M such that, if we set $\mathcal{H}_0 := \mathcal{H}$ and $\mathcal{H}_i := \mathcal{H}_{i-1} \backslash\backslash \Sigma_i$, then Σ_{i+1} is a non-duplicate normal surface in $\mathcal{H}_i$ which is either fundamental or the double of a fundamental surface. There is a normal surface representative of the collection in $\mathcal{H}$ whose size is at most $c_S^{|\mathcal{H}|^2}$.*

Split handle structures naturally arise when we cut along normal surfaces in handle structures that are dual to triangulations. We keep track of the *forbidden region*, which is the part of the boundary that comes from the normal surface, and of *parallelity pieces*, which are I-bundles over surfaces from where the normal surface runs close to itself (compare to parallelity handles in Section 2.7.1). Figure 5.1 is a motivating example; it shows the result of cutting a tetrahedron (in some larger triangulation) along an elementary triangle, from the point of view of first, the triangulation, and second, the dual split handle structure.

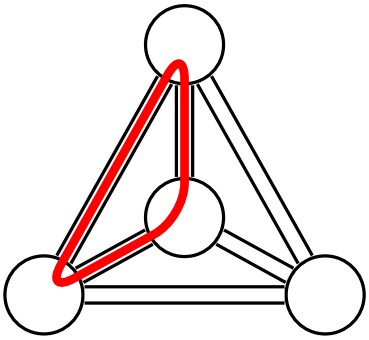
We will use the term "sutures" for the boundary of this forbidden region as split handle structures are reminiscent of sutured handle structures, which were devised by Lackenby (see §5 of [Lac21] for an exposition) to allow for normal-surface-type arguments in the context of Scharlemann's combinatorial approach to Gabai's sutured manifold decompositions [Sch89]. Normal surfaces in split handle structures are also evocative of the normal surface theory of handle structures with boundary pattern if we require that elementary discs do not intersect the pattern. However, in the boundary pattern case, we usually require that the pattern is contained in the 1-skeleton of the



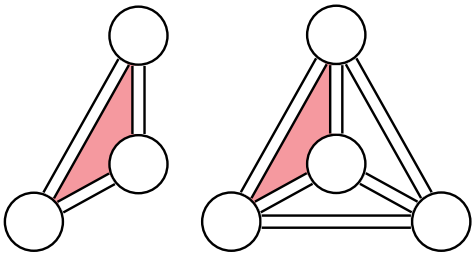
(a) An elementary disc in a tetrahedron.



(b) After cutting along it.



(c) Its boundary curve in the boundary graph of the dual handle.



(d) The boundary graphs of the resulting split handles. (We have rotated the first one so that the forbidden region is contained in the graph as we draw it on the plane, rather than being its complement in S^2.) These particular boundary graphs have no sutures.

Figure 5.1: The pieces resulting from cutting along an elementary triangle in a tetrahedron (in the interior of some larger triangulation), and the picture in the dual split handle structure. To depict the handles, we draw their boundary graphs (in S^2). The forbidden region is shaded in red.

induced handle structure on the boundary (see [Mat07, §3]). Here, the sutures are normal curves.

Definition 5.3.1. A *split handle structure* $\mathcal{H}$ for a compact orientable 3-manifold M is a partition of M into:

1. k-handles for k between 0 and 3, where each k-handle has a homeomorphism to $D^k \times D^{3-k}$, and

2. parallelity pieces, each with a homeomorphism to $\Sigma \overset{(\sim)}{\times} I$ for Σ a compact surface

and with a distinguished *forbidden region* $\mathcal{I} \subseteq \partial M$ such that the following conditions hold. Write $\mathcal{H}^k$ for the collection of k-handles, $\mathcal{H}^{\mathcal{P}}$ for the collection of parallelity pieces, and $\partial_h \mathcal{H}^{\mathcal{P}}$ or $\partial_v \mathcal{H}^{\mathcal{P}}$ respectively for the collection of the horizontal or vertical

boundaries of the parallelity pieces. The *boundary graph* of a 0-handle H in a split handle structure is the decorated graph in $\partial H \cong S^2$ whose vertices, which we call *islands*, are the components of $H \cap \mathcal{H}^1$, whose edges (which we call *bridges*) are the components of $H \cap \mathcal{H}^2$ and $H \cap \mathcal{H}^{\mathcal{P}}$, and which may also contain *sutures*, which are the arcs $H \cap \partial\mathcal{I} - (\mathcal{H}^1 \cup \mathcal{H}^2 \cup \mathcal{H}^{\mathcal{P}})$. We say that the boundary graph divides ∂H into islands, bridges, *lakes* (components of intersection between ∂H and $(\mathcal{H}^3 \cup \partial M) - \mathcal{I}$), and *forbidden regions* which are components of $H \cap \mathcal{I}$; the sutures are the intersections between the forbidden region and lakes.

We require that:

1. each k-handle $D^k \times D^{3-k}$ intersects handles of lower index in exactly $\partial D^k \times D^{3-k}$, and is disjoint from the other k-handles;
2. the boundary graph of each 0-handle is connected;
3. each parallelity piece is disjoint from the 2- and 3-handles and the other parallelity pieces;
4. the forbidden region $\mathcal{I} \subseteq \partial M$ contains $\partial_h \mathcal{H}^{\mathcal{P}}$;
5. each 1-handle $D^1 \times D^2$ intersects 2-handles $D^2 \times D^1$ in components that are of the form $D^1 \times \gamma$ in the 1-handle and $\lambda \times D^1$ in the 2-handle, where γ and λ collections of arcs in ∂D^2 in the respective product structures;
6. the intersection of any component $P \cong \Sigma \overset{(\sim)}{\times} I$ of $\mathcal{H}^{\mathcal{P}}$ with a 1-handle $D^1 \times D^2$ is as $D^1 \times \gamma$ in the 1-handle and $\lambda \times I$ in the parallelity piece, where γ is a collection of arcs in ∂D^2 and λ is a collection of arcs in $\partial\Sigma$.

If $\mathcal{I}$ is not empty, we require the following. Write $(\partial\mathcal{H})^k$ for each k for the components of $\partial M \cap \mathcal{H}^k$, and $(\partial\mathcal{H})^{\mathcal{P}}$ for the components of intersection of ∂M with $\mathcal{H}^{\mathcal{P}}$. Note that $(\partial\mathcal{H})^0$ and $(\partial\mathcal{H})^2$ are collections of discs, and $(\partial\mathcal{H})^1$ and $(\partial\mathcal{H})^{\mathcal{P}}$ are collections of discs and possibly some annuli. We require that $\partial\mathcal{I}$ avoids $(\partial\mathcal{H})^2$, runs through discs of $(\partial\mathcal{H})^0$ and $(\partial\mathcal{H})^1$ in arcs that each do not start and end on the same component of $(\partial\mathcal{H})^0 \cap (\partial\mathcal{H})^1$, and intersects each component of $\partial_v \mathcal{H}^{\mathcal{P}} \cap \partial M$ in exactly two arcs or curves, each of which is transverse to the I-bundle structure from the parallelity pieces.

If the forbidden region is empty (which implies that there are no parallelity pieces), this is the usual notion of handle structure.

Definition 5.3.2. A surface in a split handle structure is *∂-compressible* if it admits a non-trivial ∂-compression disc that is disjoint from the forbidden region.

5.3.1 Normal surfaces

We will now define normal surfaces with respect to a split handle structure, give a normalisation procedure, and prove that if we apply this procedure to an incompressible ∂-incompressible surface in an irreducible ∂-irreducible manifold then we obtain an admissibly isotopic normal surface.

Definition 5.3.3. A properly-embedded surface in M is *standard* with respect to a split handle structure if it satisfies the following conditions:

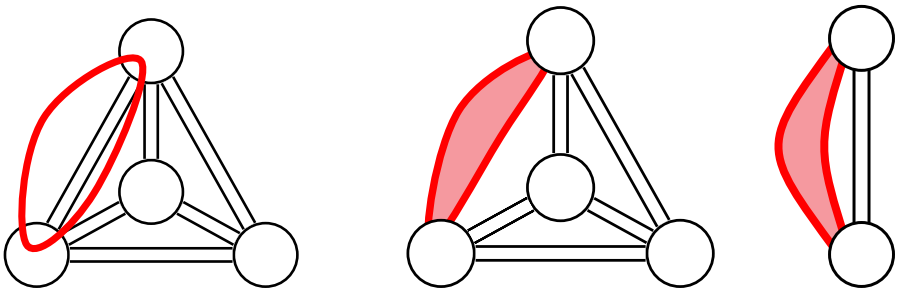
1. it is disjoint from the 3-handles and from the forbidden region;
2. it is transverse to the I-bundle structure of the 2-handles $D^2 \times I$ and the parallelity pieces $\Sigma \overset{(\sim)}{\times} I$, and is disjoint from their horizontal boundaries;
3. no component of it is contained in a parallelity piece;
4. it intersects the 1-handles $D^1 \times D^2$ in a union of D^1-fibres, where each component of this intersection is a disc;
5. it intersects the 0-handles in discs.

It follows from the definition that if F is a standard surface, then it intersects each 2-handle $D^2 \times D^1$ in sheets of the form $D^2 \times \{*\}$, and similarly intersects each parallelity piece in sheets that are a section of the I-bundle or, if Σ is non-orientable, the double cover of a section.

Definition 5.3.4. A standard surface F in M with respect to a split handle structure $\mathcal{H}$ is *normal* if it satisfies the following additional conditions:

1. F intersects each 1-handle $D^1 \times D^2$ in $D^1 \times \lambda$ where λ is a collection of disjoint proper arcs in the island $\{0\} \times D^2$, such that no component of λ starts and ends on the same connected component of the intersection of the island with a lake;
2. the intersection of F with each lake does not contain any closed curves or arcs that start and end on the same component of the intersection of the lake with an island;
3. F intersects each 0-handle in discs, called *elementary discs*, such that the boundary curve of each of these discs crosses each bridge and lake at most once, and if a bridge and a lake are adjacent, intersects only one of the pair.

This is a generalisation of the usual notion of a normal surface in a handle structure. Note that the boundary of an elementary disc of F in the boundary graph of H determines the disc. Figure 5.2 is an analogous picture to Figure 5.1. This time, we cut along an elementary disc in the dual split handle structure that does not correspond to an elementary disc in the triangulation, as its boundary passes through a lake (which is dual to a vertex in the dual tetrahedron). Each arc in a lake produces a suture when we cut along it. We would encourage the reader to satisfy themselves that the disc boundary shown is in fact that of an elementary disc, and that the resulting boundary graphs are possible boundary graphs of 0-handles in a split handle structure. We will generalise the examples in Figures 5.1 and 5.2 when we define *induced split handle structures* in Definition 5.3.14.



(a) The boundary of an elementary disc in the boundary graph in S^2.

(b) The two split 0-handles resulting from cutting along this elementary disc.

Figure 5.2: The induced split handle structure from cutting along an elementary disc in a split 0-handle that does not correspond to one in a tetrahedron. The forbidden region is shaded in red, and the sutures are depicted as thick red lines.

Definition 5.3.5. An *admissible isotopy* of a surface with respect to a split handle structure is an isotopy of the surface in the manifold that fixes the forbidden region $\mathcal{I}$ as a set.

A *normal isotopy* of a normal surface in a split handle structure is an isotopy of the surface in the manifold that fixes each of $\mathcal{I}$, $\mathcal{H}^k$ for each k, and $\mathcal{H}^{\mathcal{P}}$ as a set.

Definition 5.3.6. The size of a normal surface F in a split handle structure, $s(F)$, is the number of elementary discs in F; that is, it is the number of components of $F \cap \mathcal{H}^0$.

Definition 5.3.7. The weight of a normal surface F is $(p(F), b(F), |F \cap \partial M|)$, where $p(F)$, the *plate degree*, is $|F \cap (\partial\mathcal{H}^2 \cup \partial\mathcal{H}^{\mathcal{P}})|$ and $b(F)$, the *beam degree*, is $|F \cap \mathcal{H}^1|$.

We will now give a procedure to normalise a surface in a 3-manifold, and show that this procedure reduces weight.

Lemma 5.3.8. *Let S be an incompressible (connected) surface in the orientable I-bundle $\Sigma \overset{(\sim)}{\times} I$, other than a trivial disc or sphere, that is disjoint from the horizontal boundary and does not admit any ∂-compression discs with respect to the vertical boundary of $\Sigma \overset{(\sim)}{\times} I$. Then S is isotopic to $\Sigma \times \{*\}$ or its double cover.*

Proof. First, suppose that Σ is a disc. Then as $\Sigma \times I$ is a ball (so S is two-sided) and S is incompressible and hence π_1-injective, S is a sphere or a disc. If it is a sphere, it is a trivial one. If ∂S contains a trivial curve in the vertical boundary, then S must be a boundary-parallel disc. Thus S is a disc and its boundary is essential in the vertical boundary, so it is isotopic to $\Sigma \times \{*\}$.

For the general case, suppose that Σ has boundary. Decompose Σ into one 0-handle, which we view as a polygon, and some number of 1-handles, which we view as rectangles. Above each edge of the 0-handle and 1-handles in $\Sigma \overset{(\sim)}{\times} I$ is a quadrilateral with two boundary edges in the vertical boundary of $\Sigma \overset{(\sim)}{\times} I$ and two in the horizontal boundary. Isotope S so that it is transverse to these quadrilaterals and thus intersects each quadrilateral in a collection of arcs and closed curves. As S is incompressible, each of these closed curves bounds a disc in S, so by minimising the number of components of intersection between S and the quadrilaterals we can assume that the intersection is only of arcs. Since S is disjoint from the horizontal boundary, each of these arcs starts and ends on either the same or different vertical boundary arcs of a quadrilateral. If one starts and ends on the same vertical boundary arc, by taking an outermost such arc we can obtain a ∂-compression disc for S with respect to the vertical boundary, so by an isotopy we can remove this arc. Thus up to isotopy S intersects each quadrilateral in arcs that are transverse to the product structure.

Consider one of the handles of Σ, H, which is a disc. Consider the intersection of S with $H \overset{(\sim)}{\times} I \cong D^2 \times I$ in $\Sigma \overset{(\sim)}{\times} I$. We can assume (by minimising the number of components of intersection between S and the quadrilaterals) that S is incompressible in this copy of $D^2 \times I$, so as we have already discussed, each component of it is a (trivial) sphere (which we have ruled out) or disc. The boundary of $H \overset{(\sim)}{\times} I$ has three parts: its horizontal boundary, from which S is disjoint, quadrilaterals, and pieces of $\partial_v(\Sigma \overset{(\sim)}{\times} I)$. For similar reasoning as with the quadrilaterals, up to isotopy S also intersects the vertical boundary pieces in arcs that are transverse to the product structure. Thus the boundary of each disc of $S \cap H \overset{(\sim)}{\times} I$ is an essential curve in $\partial_v(H \overset{(\sim)}{\times} I)$. Thus S intersects $H \overset{(\sim)}{\times} I$ in a collection of horizontal discs for each H, and so intersects all of $\Sigma \overset{(\sim)}{\times} I$ in either $\Sigma \times \{*\}$ or (if Σ is non-orientable) possibly $\Sigma \tilde{\times} S^0$.

If Σ is closed, let D be a sub-disc of Σ. Isotope S to minimise the number of components of $S \cap \partial(D \times I)$. The intersection of S with $D \times I$ is incompressible since

otherwise, as S is incompressible, $|S \cap \partial(D \times I)|$ would not be minimal. For the same reason $S \cap (D \times I)$ contains no spheres or trivial discs. By the first part of this proof, $S \cap (D \times I)$ is a collection of discs of the form $D \times \{*\}$. If $S \cap (\Sigma - D) \times I$ admits a ∂-compression disc with respect to the vertical boundary, we can use it to isotope S to reduce $|S \cap \partial(D \times I)|$. Thus by the previous part of the proof, $S \cap (\Sigma - D) \times I$ is isotopic to $(\Sigma - D) \times \{*\}$ or its double cover, so we have the result. □

We wish to modify an arbitrary surface F in M, whose boundary is disjoint from the forbidden region, by a series of admissible isotopies and normalisation moves such that the result is normal (but may not be isotopic to F). Compare the following procedure to the proof of Theorem 3.4.7 in [Mat07]. The only substantial difference is in Move 2, as we need to consider parallelity pieces that are more complicated than 2-handles.

Let F be a properly-embedded surface in M that is disjoint from the forbidden region. We will see in Proposition 5.3.10 that almost all the *normalisation moves* which follow will reduce the weight of F. All isotopies in these moves are required to be admissible.

Procedure 5.3.9 (Normalisation)**.** The *normalisation procedure* is to perform Move 1 once, and then repeat Moves 2-7 in sequence as long as possible.

Move 1. The boundary of the forbidden region $\mathcal{I}$ is disjoint from the boundaries of the 2-handles. (Admissibly) isotope F so that F is transverse to the handle structure and ∂F is disjoint from the horizontal boundaries of the 2-handles. Each 3-handle contains an open ball that is disjoint from F, so by expanding these balls, we can isotope F to be disjoint from the 3-handles. Discard any components of F that are entirely contained in a parallelity piece.

Move 2. Consider each component $H \cong \Sigma \overset{(\sim)}{\times} I$ of $\mathcal{H}^2 \cup \mathcal{H}^{\mathcal{P}}$. If any of the components of $F \cap \partial H$ are trivial curves in the vertical boundary, compress F along the discs in ∂H that these curves bound. Similarly, if $F \cap H$ admits a compression disc, compress F along it. If $\partial_v H$ intersects ∂M (in which case H is a parallelity piece), and $F \cap H$ admits a ∂-compression disc D that is disjoint from the horizontal boundary – that is, where the arc of ∂D in ∂H, α, is in $\partial_v H$ – then compress along it. If F admits a ∂-compression disc with respect to $\partial_v H$ whose boundary is disjoint from ∂M isotope F through the compressing disc. Now, as each component of $F \cap \partial H$ is an essential curve in $\partial_v H$, up to an isotopy supported in a collar of $\partial_v H$ within H, we can arrange that each of these curves is transverse to the induced I-bundle structure on $\partial_v H$.

If we performed a ∂-compression, repeat the first two steps of this move. Finally, discard any components of $F \cap (\mathcal{H}^2 \cup \mathcal{H}^{\mathcal{P}})$ that are spheres or properly embedded ∂M-parallel discs.

Move 3. For each 1-handle $D^1 \times D^2$, take a disc D whose boundary is essential in $D^1 \times \partial D^2$, transverse to F, that minimises the number of components in $D \cap F$. Blow a regular neighbourhood of D out to be the whole 1-handle.

Move 4. If any component of intersection of F with a 1-handle is a tube $S^1 \times I$, compress it and isotope the two resulting discs out of the 1-handle. Similarly, if any component of intersection with a 1-handle $D^1 \times D^2$ is a disc whose intersection with $D^1 \times \partial D^2$ is contained in a single region of $(\partial M - \mathcal{I}) \cap (D^1 \times D^2)$, ∂-compress this tunnel and isotope the pieces out of the 1-handle.

Move 5. Compress any compressible pieces of $F \cap \mathcal{H}^0$ and discard any trivial spheres in the 0-handles.

Move 6. If F intersects a lake in a trivial curve, compress it and throw away the resulting ∂-parallel disc. If F intersects a lake in an arc that starts and ends on the same component of the intersection of the lake with an island, by an isotopy (again as $\partial\mathcal{I}$ intersects the cell structure on the boundary in a normal curve) push this piece F off the lake and into the 1-handle.

Move 7. If a disc of $F \cap \mathcal{H}^0$ crosses a bridge twice or a bridge and an adjacent lake, isotope the portion of disk between them into the bridge. If it crosses a lake twice, ∂-compress the resulting tunnel along both of its intersections with the boundary of the lake.

Proposition 5.3.10. *Let F be a properly-embedded surface in a split handle structure $\mathcal{H}$. Applying Procedure 5.3.9 terminates and the result (if non-empty) is a normal surface.*

Proof. After Move 1 F satisfies conditions 1 and 3 of Definition 5.3.3. If Move 2 has no effect, $F \cap H$ is incompressible and ∂-incompressible with respect to the vertical boundary, so by Lemma 5.3.8 F satisfies condition 2 of Definition 5.3.3. If Moves 3 and 4 have no effect then it satisfies condition 4 of Definition 5.3.3 and condition 1 of Definition 5.3.4. After Move 5 it satisfies condition 5 of Definition 5.3.3. Then Move 6 ensures condition 2 of Definition 5.3.4 and Move 7 ensures condition 3. As a result, if no more moves can be performed, then F is normal.

Let the normalisation complexity of F be

$$(p(F), b(F), |F \cap \partial M|, \gamma(F), \eta(F), n(F)),$$

ordered lexicographically, where the new terms are:

- $\gamma(F) = \sum_{i=1}^{m}(1 - \chi(F_i))$ where $\{F_i\}$ is the collection of connected components of $F \cap \mathcal{H}^0$ that are not spheres,
- $\eta(F) = \sum_{i=1}^{m}(1 - \chi(F_j))$ where $\{F_j\}$ is the collection of connected components of $F \cap (\mathcal{H}^2 \cup \mathcal{H}^{\mathcal{P}})$ that are not spheres,
- and $n(F)$ is the number of connected components of F.

We will show that each of Moves 2-7 reduces the normalisation complexity of F.

For Move 2, suppose we perform a ∂-compression on $F \cap H$. As $\mathcal{H}$ is a split handle structure, if $\partial\mathcal{I}$ intersects a component of $\partial_v H \cap \partial M$ then it does so in two curves or arcs. Thus this component is either contained in $\mathcal{I}$, or (as $\partial_h H$ is contained in $\mathcal{I}$) it intersects $\mathcal{I}$ in a neighbourhood of its boundary with the horizontal boundary of H. If $F \cap H$ admits a ∂-compression disc D that is disjoint from the horizontal boundary, then let $\bar{\alpha}$ be a properly-embedded arc in $\partial_v H$, such that α is a subarc of it, which is isotopic in $\partial_v H$ to one of the I fibres of $\partial_v H$ and has minimal intersection with each of the components of $F \cap \partial_v H$ – that is, it intersects each of the essential curves in this collection once. This is possible since α runs between two different components of $F \cap \partial_v H$; otherwise we would be able to upgrade our ∂-compression disc to a compression disc. (Note that $|\bar{\alpha} \cap F|$ is the degree of the projection map from $F \cap H$ to Σ. When Σ is orientable, this is the number of sheets of F in H.)

The effect of the ∂-compression or isotopy to the boundary on $F \cap \partial H$ is to remove $\partial\alpha \times I$ from it and add $\alpha \times \partial I$ to it, for some small thickening $\alpha \times I$ of α. This reduces $|\bar{\alpha} \cap F|$ by two, and thus we see that the number of essential curves in $F \cap \partial_v H$ has reduced by two, so either $p(F)$ has reduced or we reduce it in the next step of Move 2. After this move, $F \cap H$ does not admit any ∂-compression discs with respect to the vertical boundary of H. Compressing F along trivial curves in the vertical boundary of H reduces $p(F)$. Compressing $F \cap H$ reduces $\eta(F)$ and fixes the earlier terms in the normalisation complexity.

If Move 3 is non-trivial then it does not increase $p(F)$ and reduces $b(F)$, and the same applies for Move 4. Move 5 does not change $p(F)$ or $b(F)$. If there are compressible pieces then it reduces $\gamma(F)$, and otherwise if all the pieces are 2-spheres or discs then it reduces $n(F)$ and fixes everything else. Move 6 either reduces $|F \cap \partial M|$ and does not change $p(F)$ and $b(F)$, or fixes $p(F)$ and reduces $b(F)$. Move 7 either decreases $p(F)$ or fixes it and reduces $b(F)$. As a result, the procedure terminates with a normal surface. □

Definition 5.3.11. A properly-embedded disc in a 3-manifold M is *essential* if its boundary does not bound a disc in ∂M.

This definition does not change if there is forbidden region.

Lemma 5.3.12. *Suppose that M is irreducible with a split handle structure $\mathcal{H}$ such that there is some essential disc in M that avoids the forbidden region. Then there is a normal essential disc, disjoint from the forbidden region.*

Proof. Apply the normalisation procedure to this essential disc D. As M is irreducible and D is (trivially) incompressible, whenever we compress D in the normalisation procedure we will produce a surface admissibly isotopic to D and a trivial sphere, which we can discard. Whenever we ∂-compress D we will produce two discs, both of which are disjoint from the forbidden region. At least one of them must be essential as ∂D does not bound a disc in ∂M. Discard the other. □

Proposition 5.3.13. *Let F be an incompressible ∂-incompressible properly-embedded surface in an irreducible ∂-irreducible manifold M with split handle structure $\mathcal{H}$, disjoint from the forbidden region, such that no component is a trivial sphere or disc or is entirely contained in a parallelity piece. Then F is admissibly isotopic to a normal surface.*

Proof. Apply the normalisation procedure to F. As F is incompressible and ∂-compressible and M is irreducible and ∂-irreducible, each time we (∂-)compress F in the normalisation procedure we will produce two components: a surface that is admissibly isotopic to F and a trivial sphere or disc. Thus we produce a normal surface admissibly isotopic to F, as well as some collection of trivial spheres and discs, which we can discard. □

Induced split handle structures

We will now see that we obtain a split handle structure whenever we cut a (split) handle structure (dual to, say, a simplicial triangulation) along a normal surface. Figures 5.1 and 5.2 are motivating examples here.

Definition 5.3.14. Let F be a normal surface in M, with respect to a split handle structure $\mathcal{H}$. The *induced split handle structure* on $M\backslash\backslash F$ is constructed as follows.

Consider $\mathcal{H}\backslash\backslash F$. Set its forbidden region $\mathcal{I}$ to be the union of the forbidden region from M and the image of F in $M\backslash\backslash F$. (Since F was normal, these are disjoint.) As F is disjoint from the 3-handles, we can continue to view them as 3-handles in $\mathcal{H}\backslash\backslash F$.

Each component of a k-handle in $\mathcal{H}\backslash\backslash F$ will become either a k-handle or a parallelity piece. This is determined as follows. First, if $P \cong \Sigma \overset{(\sim)}{\times} I$ is a parallelity piece of $\mathcal{H}$, as F intersects P in sheets transverse to the I-bundle structure, each component of $P\backslash\backslash F$ inherits an I-bundle structure as either $\Sigma \overset{(\sim)}{\times} I$ or possibly, if Σ is not orientable, as $\tilde{\Sigma} \times I$ where $\tilde{\Sigma}$ is the double cover of Σ. Thus we turn each component of $P\backslash\backslash F$ into a new parallelity piece.

If $H \cong D^2 \times D^1$ is a 2-handle of $\mathcal{H}$, a component C of $H\backslash\backslash F$ is itself an I-bundle over D^2, with two horizontal boundary components, each of which arises from intersection with a 3-handle, with $\mathcal{I}$, or with $\partial M - \mathcal{I}$. If both components arise from intersection with $\mathcal{I}$, set C to be a parallelity piece; otherwise, view C as a 2-handle.

If $H \cong D^1 \times D^2$ is a 1-handle of $\mathcal{H}$, a component C of $H\backslash\backslash F$ is a parallelity piece if its boundary consists of the following components: first, two components in the forbidden region $\mathcal{I}$ (whether from the forbidden region in M or arising from intersection with F); second, two components arising from intersection with pieces of 0-handles from M; and finally, the two remaining components where each is a component of intersection with one of ∂M, a single 2-handle, or a single parallelity piece. In this case, C has an I-bundle structure by setting $\partial_h C$ to be the two components in the forbidden region, and choosing a product structure on the remaining boundary, $\partial_v C$, such that each of the four components described above is a union of fibres, and then interpolating. (This product structure is chosen to be compatible with the product structure on any parallelity pieces defined thus far that C intersects).

Finally, if $H \cong D^0 \times D^3$ is a 0-handle and C is a component of it, consider the boundary of C. Set C to be a parallelity piece if its boundary contains two components of intersection with $\mathcal{I}$ and if there is a product structure $D^2 \times I$ on C such that its intersection with $\mathcal{I}$ is $D^2 \times \partial I$, and each component of its intersection with any of the other handles created so far is of the form $\alpha \times I$, where α is an arc or curve in ∂D^2. Again, we can choose this product structure to be compatible with any parallelity pieces that C intersects. Otherwise, set C to be a 0-handle.

Now, take the parallelity pieces of $\mathcal{H}$ to be the union of the parallelity pieces described so far, which we equipped with compatible product structures where they intersected.

We write $\mathcal{H}\backslash\backslash F$ for the induced split handle structure on $M\backslash\backslash F$.

Lemma 5.3.15. *If F is a normal surface in a split handle structure $\mathcal{H}$, then the induced split handle structure $\mathcal{H}\backslash\backslash F$ is (as the name suggests) itself a split handle structure.*

Proof. If a piece from a k-handle in $\mathcal{H}\backslash\backslash F$ becomes a parallelity piece in the induced split handle structure, then any pieces from adjacent handles of higher index will also be parallelity pieces. As a result, each k-handle $D^k \times D^{3-k}$ will intersect handles of lower index in $\partial D^k \times D^{3-k}$ and will be disjoint from the other k-handles. The boundary graph of each 0-handle will be connected as the boundary graph of each 0-handle arises from taking a connected boundary graph, removing some discs from it, and adding in the boundary of these disc (as we add a new suture for each time the boundary of one of the discs runs through a lake). As F does not intersect any 3-handles, so each one becomes a 3-handle in $\mathcal{H}\backslash\backslash F$, each parallelity piece will be disjoint from the 2- and 3-handles. As intersections of parallelity pieces with 1-handles must necessarily arise from intersections between a 2-handle and a 1-handle, and the product structure of the 2-handle is compatible with that of the parallelity piece, the intersections between parallelity pieces and 1-handles are of the required form. The intersection of $\partial\mathcal{I}$ with the boundary has the required form as F is normal. □

Subtetrahedral split handle structures

As our motivations are from normal surface theory in triangulations, we are particularly interested in split handle structures that arise from taking handle structures dual to (suitable) triangulations and cutting along normal surfaces in them. We call these *subtetrahedral*, and can gain good control of normal surfaces in them using their combinatorics. Note that, if the forbidden region is empty, this definition of *subtetrahedral* agrees with Definition 2.2.6.

Definition 5.3.16. A 0-handle H is *semitetrahedral* if it is disjoint from the forbidden region and its boundary graph is a connected subgraph of the complete graph on four vertices. A split handle structure is *semitetrahedral* if all of its 0-handles are.

A 0-handle H is a *subtetrahedral* split 0-handle if there is some 0-handle H' in some semitetrahedral handle structure $\mathcal{H}$ with normal surface F such that H is homeomorphic to one of the non-parallelity pieces obtained from H' in the induced split handle structure $\mathcal{H}\backslash\backslash F$.

We say that a split handle structure is *subtetrahedral* if all of its 0-handles are subtetrahedral.

Lemma 5.3.17. *Let $\mathcal{T}$ be a (material) triangulation of a (compact) 3-manifold M, such that the intersection of each tetrahedron with ∂M is connected and contractible. The dual handle structure to $\mathcal{T}$, formed by taking a $(3-k)$-handle for each k-simplex*

of $\mathcal{T}$ that is not in $\partial\mathcal{T}$, satisfies the definition of a subtetrahedral split handle structure with empty forbidden region and no parallelity pieces.

If $\mathcal{H}$ is a subtetrahedral split handle structure, and F is a normal surface in $\mathcal{H}$, then the induced split handle structure on $\mathcal{H}\backslash\backslash F$ is also subtetrahedral.

Proof. The complement of the boundary graph of a 0-handle will deformation retract to the intersection of the corresponding tetrahedron with the boundary, so the handles of this dual handle structure have connected boundary graphs. The second part follows by definition. □

We will now give some bounds on the combinatorics of subtetrahedral split handle structures and normal surfaces in them. To be able to bound normal surfaces, we will only consider those that are not *duplicate*: those that do not have normally isotopic components, and do not bound a small neighbourhood of the forbidden region. We think of these as being "redundant": if we have a subtetrahedral split handle structure that was formed by cutting a subtetrahedral handle structure along normal surfaces, then we see duplicate normal surface when two components of these normal surfaces are parallel.

Definition 5.3.18. A normal surface in a split handle structure with forbidden region $\mathcal{I}$ is *duplicate* if it has two components that are normally isotopic, or it has a component F such that there is a connected component I of the forbidden region, with a collar $C(I) \cong I \times [0,1]$, such that F is normally isotopic to $I \times \{1\}$ in $C(I)$.

Lemma 5.3.19. *Let H be a subtetrahedral split 0-handle. Let G be its boundary graph in $\partial H \cong S^2$. Then G contains between one and four islands; each island has at most three components of intersection with bridges; and if b is the number of bridges of G, then the number of sutures of G at most $12 - 2b$ and an island intersecting v bridges has at most $6 - 2v$ intersections with sutures.*

Proof. Consider a semitetrahedral handle structure $\mathcal{H}$ and normal surface F such that H is homeomorphic to one of the pieces of $\mathcal{H}\backslash\backslash F$ from some 0-handle H' in $\mathcal{H}$. We can think of forming the pieces of H' in the induced split handle structure as occurring in two steps. First, we cut H' along a collection of elementary discs. This has the effect on the boundary of H', which we think of as a graph embedded in S^2, of cutting it along a collection of (separating) curves and filling the holes in with the forbidden region. Second, we possibly replace some of the 1-handles with parallelity pieces, in which case (as, if a 1-handle can be given a parallelity structure, so can any 2-handles

it borders) the effect on the boundary graph is to merge a valence one or two island with the bridge(s) adjacent to it.

Consider cutting along one elementary disc of F in H' at a time. Suppose that we have cut along some number of elementary discs already to obtain a piece H'', and the next disc in H'' to cut along is D. As $\mathcal{H}$ is semitetrahedral, the islands of H' were at most trivalent, so ∂D runs through each island of H'' at most once. As ∂D runs through each island or bridge at most once and is separating, the subgraph of islands and bridges after cutting along it is a subgraph of the complete graph on four vertices, so after cutting into pieces, there are between zero and four islands, each of which has valence at most three. We can rule out the case when the boundary graph contains no islands as then D would not have been elementary. Turning 1-handles into parallelity handles does not increase the number of islands or their valence. If all of the 1-handles become parallelity pieces (so there were zero islands) then in fact the entire 0-handle would have become a parallelity piece, so this is impossible.

Let G be the boundary graph of H'' (containing b bridges) and $\mathcal{I}$ its forbidden region, and assume that G is connected and contains at most $12 - 2b$ sutures. The elementary disc boundary ∂D separates $\partial H''$ into two components. Pick one of them, C, to consider. When we cut along D, the other component from $\partial H''$ has one less bridge than G for every bridge fully contained in C, one less suture for each suture in C, and one more suture for each arc of D that runs through a lake. Thus for the count of sutures it suffices to prove that twice the number of bridges fully contained in C plus the number of sutures of H'' in C is at least the number of arcs in ∂D that run through a lake. Consider one of these arcs α. It separates the lake into two components, one of which is in C, and each of which is bordered by at least one bridge or suture. If the component in C is bordered by a bridge, assign the arc α to that bridge. As the bridge is adjacent to at most two lakes, and ∂D crosses each lake only once, by doing this we will associate at most two arcs to the bridge. Also note that as the elementary disc does not run through an adjacent bridge and lake, this bridge is contained in C. Otherwise, this component in C is bordered by at least one suture, as α runs between two different components of intersection of an island and a lake, so assign α to this suture, which again we see is contained in C. The forbidden region is on the other side of this suture so we will not assign any other arcs to it. We can thus see that the count is as we claimed. Turning 1-handles into parallelity regions only reduces the number of bridges further.

For the final statement, fix an island in H'' of valence v that intersects at most $6 - 2v$ sutures. Consider some elementary disc boundary running through this island.

By the same procedure as for the total suture count, assign any new sutures that are created to a bridge or a suture that intersects the island to get the result: we start with three adjacent bridges, and add at most two sutures each time we remove one. □

Lemma 5.3.20. *The total number of bridges and lakes in a subtetrahedral split 0-handle H is at most 13.*

Proof. From Lemma 5.3.19, each 0-handle contains at most 6 bridges and $12 - 2b$ sutures, where b is the number of bridges. The number of lakes from the subgraph on islands and bridges alone is at most $b + 1$ by an Euler characteristic argument. Each pair of sutures can increase the number of lakes by one, if they bound a forbidden region that borders a lake on both sides. (If a suture bounds a forbidden region that borders a bridge on one side then it does not increase the complexity.) Thus having sutures does not in fact increase the count, so that maximum count is $6 + 6 + 1 = 13$. □

Lemma 5.3.21. *There are a finite number of subtetrahedral split 0-handles up to homeomorphism, and each of these can contain a finite number of elementary disc types up to normal isotopy.*

Proof. That there are a finite number of types of subtetrahedral split 0-handles follows from the combinatorial constraints of Lemma 5.3.19. An elementary disc D in such a 0-handle H is determined by its boundary ∂D, which is itself completely described (as the complement of the boundary graph is a collection of discs) by the ordered list of islands it runs through and, for each, which intersection with a bridge or lake it enters and leaves by. As it may run through each island, bridge or lake at most once, there are a finite number of possibilities. □

Notation 5.3.22. Write d_H for the maximum possible number of elementary disc types up to normal isotopy in a subtetrahedral split 0-handle.

Remark 5.3.23. By Lemma 5.3.20, each 0-handle has a total of at most 13 lakes and bridges, so, by labelling each disc type by the number of edges in it and then the order of them, d_H is at most $13 \cdot 13!$. We can do much better than this bound; as we prove in Lemma A.2 in Appendix A, d_H is 59.

Lemma 5.3.24. *If F is a (possibly disconnected) non-duplicate normal surface in a subtetrahedral split handle structure $\mathcal{H}$ then $|\mathcal{H} \backslash\backslash F|$ is at most $60|\mathcal{H}|$.*

Proof. We know that $\mathcal{H}\backslash\backslash F$ is a split handle structure from Lemma 5.3.15. The normal surface F intersects each 0-handle H as a collection of elementary discs. Consider cutting along the discs in H one isotopy class at a time. By Lemma A.2 there are at most 59 elementary disc types in F in any given 0-handle. The first disc we cut along in an isotopy class may separate off a non-parallelity 0-handle from H, but every disc after that will only separate off a parallelity handle, so does not contribute to $|H\backslash\backslash F|$. Thus $|H\backslash\backslash F|$ is at most the number of elementary disc types in $F \cap H$, plus one, which is at most $59 + 1 = 60$. □

There is a natural inclusion map from surfaces in $M\backslash\backslash F$, disjoint from the forbidden region, to surfaces in M, which takes normal surfaces to normal surfaces.

Lemma 5.3.25. *A normal surface F in a subtetrahedral split handle structure has non-zero $b(F)$, the beam degree.*

Proof. As no component of a normal surface is contained in a parallelity piece, every normal surface contains at least one elementary disc. It suffices to show that the boundary of each elementary disc intersects at least one island. Since the boundary graph of each 0-handle is connected and contains at least one island, every curve contained in a lake is trivial and each bridge starts and ends on an island. The boundary of the elementary disc runs through a bridge or lake and thus through some island. □

5.3.2 Normal surface sum

In this section, all split handle structures are subtetrahedral. As in the triangulation setting, we can view normal surfaces algebraically. Let $\mathcal{H}$ be a subtetrahedral split handle structure. Recall that $d_H = 59$ is the maximum number of types of elementary disc in a subtetrahedral split 0-handle. To a normal surface F in $\mathcal{H}$ we can associate a vector in $\mathbb{Z}^{d_H|\mathcal{H}|}$ by counting the number of elementary discs in F of each type. Since F must avoid any parallelity piece of $\mathcal{H}$ that does not intersect at least one 0-handle, this vector uniquely determines F. We can interpret summation in $\mathbb{Z}^{d_h|\mathcal{H}|}$ geometrically as follows.

Procedure 5.3.26 (Normal surface sum in split handle structures)**.** Suppose G_1 and G_2 are normal surfaces in a subtetrahedral split handle structure $\mathcal{H}$, such that if we view them as vectors in $\mathbb{Z}^{d_H|\mathcal{H}|}$, $F = G_1 + G_2$ is also a vector representing a normal surface. We can realise G_1 and G_2 so that they are disjoint in the 0-handles and transverse elsewhere. In the 1-handles, we can use the product structure to ensure

that any two components of $G_1 \cap \mathcal{H}^1$ and $G_2 \cap \mathcal{H}^1$ intersect in a single arc that is contained in some $\{*\} \times D^2$ in the 1-handle $D^1 \times D^2$. As we have two embedded surfaces, there are no triple points of intersection. We know that G_1 and G_2 intersect the 2-handles and parallelity pieces in sheets transverse to the I-bundle, which are determined by their boundaries. Fix a 2-handle or parallelity piece $H \cong \Sigma \overset{(\sim)}{\times} I$. So long as at most one of $G_1 \cap H$ and $G_2 \cap H$ contains a non-orientable piece, we can realise all the components of $G_1 \cap H$ and $G_2 \cap H$ as follows. Let H' be the complement in H of a collar of the boundary of H. Take the correct number of sheets of G_1 and G_2 in H'. Then interpolate in the collar to achieve the required intersection pattern of G_1 and G_2 in $\partial_v H$, such that $G_1 \cap G_2 \cap H$ contains no closed curves. If both of them contain a non-orientable sheet in H, then we can realise the two non-orientable sheets to intersect in H' in a single arc, and then similarly interpolate in the collar to achieve the required intersection pattern. At each arc of intersection of G_1 and G_2 in H, if we cut along the arc then we have two choices of how to reglue to resolve the intersection: a choice of *switch*. This is determined by the picture at one of the arc's endpoints p in $\partial_v H$. The *regular switch*, which is the choice of gluing that produces sheets transverse to the product structure, produces two sheets of the same types as those we started with (unless both the sheets were non-orientable, in which case it produces one sheet that is a double of one of them).

We need to show that choosing all regular switches is compatible. As the product structures of the 1-handles and the 2-handles or parallelity pieces are transverse, if we consider the arc of intersection in the 1-handle containing p, the choice of regular switch at p induces the regular switch on the other end of the 1-handle arc. Similarly, as the regular switch is the choice giving sheets transverse to the product structure on the 2-handle or parallelity piece, if we consider the arc in the 2-handle H containing p, the choice of regular switch at p induces the regular switch along the whole arc. Thus there is a global *regular switch* for $G_1 \cup G_2$, which produces F.

Lemma 5.3.27. *Let G_1 and G_2 be normal surfaces such that the vector $F = G_1 + G_2$ also corresponds to a normal surface. Let F' be a surface constructed by resolving the curves of intersection of $G_1 \cup G_2$ with at least one irregular switch. Normalising F' produces a surface of lower weight than F.*

Proof. Pick G_1 and G_2 to minimise $|G_1 \cap G_2|$ in their normal isotopy class. The plate degree $p(F)$ is $p(G_1) + p(G_2)$ as the intersection of F with the boundaries of the 0-handles is the union of the intersections of G_1 and G_2 with them. Suppose we take the irregular switch at some point p in the boundary of a 2-handle or parallelity piece

H. Let c be the number of essential curves of G_1 and G_2 in the component of $\partial_v H$ containing p. Write F' for the result of this irregular switch at p. The point p is in some component of intersection Q of H with a 1-handle or the boundary, as G_1 and G_2 are disjoint in the 0-handles.

Suppose that p is in the boundary of a 1-handle or a component of intersection with the boundary that is a disc. This disc Q naturally has a boundary comprised of four edges: two arcs on the boundary of $\partial_h H$ and two arcs in $\partial_v H$. The two edges of ∂Q in $\partial_v H$ each intersect $G_1 \cup G_2$ in c points. However in the irregular switch case, $Q \cap F'$ now contains at most $c-2$ spanning arcs, so it contains an arc that starts and ends on the same edge of ∂Q: a *return* in the language of Matveev [Mat07, §4]. But then this gives us a disc of $F' \cap \mathcal{H}^0$ whose boundary runs through the same bridge twice, so F' is not normal. When we normalise F' in Procedure 5.3.9, in Move 2 we will reduce its plate degree by at least two, so F' is isotopic to a normal surface of plate degree at most $p(F) - 2$.

If p is in a component of intersection with the boundary that is an annulus (that is, if the component of $\partial_v H$ containing p is entirely in the boundary) then in most cases – if at most one of the sheets containing p is non-orientable, or if p is one end of an arc of intersection that projects to an inessential arc in Σ – we could have isotoped G_1 and G_2 to remove this intersection. The only case when this is not possible is if both the sheets are non-orientable and p is an endpoint of an essential arc of intersection. If we take the irregular switch at p, then there is a spanning arc of this annulus of $\partial_v H$ that intersects the resulting configuration $c-2$ times, so again in Move 2 we will reduce the plate degree of F' by at least two. □

Lemma 5.3.28. *Let F be a connected non-duplicate normal surface in a subtetrahedral split handle structure $\mathcal{H}$. If $F = G_1 + G_2$ then neither G_1 nor G_2 is both connected and duplicate.*

Proof. If G_1 were connected and duplicate, which is to say that there was a normal isotopy taking it to the boundary of a collar of a forbidden region component, then this isotopy would also make it disjoint from G_2, so then $F = G_1 + G_2$ would have two components. □

Summands of incompressible normal surfaces

We will prove in this section that, as in standard normal surface theory, if $F = G_1 + G_2$ is a normal surface sum in an irreducible, ∂-irreducible manifold where F is incompressible, ∂-incompressible and of minimal weight in its admissible isotopy

class, then G_1 and G_2 are also incompressible and ∂-incompressible. We will be able to combine this with Lemma 5.3.25 to find fundamental representatives of interesting essential surfaces.

In Matveev's excellent book [Mat07], he carefully lays out the foundations of normal surface theory in the triangulation setting. Although we are in a more general setting, it is similar enough for the ideas in the proofs he gives to apply. See the proof of Theorem 4.1.36 in [Mat07] for his treatment of these results in the triangulation setting, drawing on work of Haken, Jaco, Oertel and Tollefson, among others. The analogues of the structure he relies on are Lemmas 5.3.25 and 5.3.27, which show respectively that no normal surface has zero weight, and that performing an irregular switch reduces the weight of a normal surface sum.

Definition 5.3.29. A surface F is *minimal* if it is of minimal weight in its admissible isotopy class.

We do *not* require that a minimal surface is normal.

Definition 5.3.30. Suppose that $F = G_1 + G_2$. Let γ be the collection of double curves of $G_1 \cup G_2$. A *patch* of $G_1 \cup G_2$ is a component of $G_1 \cup G_2 - \gamma$.

We first need a technical lemma.

Lemma 5.3.31. *Let M be an irreducible, ∂-irreducible 3-manifold with a subtetrahedral split handle structure $\mathcal{H}$. Let F be a minimal normal surface that is incompressible and ∂-incompressible, with $F = G_1 + G_2$. The patches of $G_1 \cup G_2$ are incompressible and ∂-incompressible, and none are disc patches.*

Proof. To set up, pick G_1 and G_2 in their admissible isotopy class to minimise $|G_1 \cap G_2|$, subject to the requirement that $F = G_1 + G_2$.

First, we show there are no disc patches. Suppose that a disc patch E did exist, say in G_1. Write s for its double curve in $G_1 \cup G_2$, and s_1 and s_2 for the corresponding trace curves in F, where s_1 bounds E and s_2 is the other curve. If the boundary of this patch is one-sided in G_1, then the connected component of G_1 containing E is a normal projective plane P. As M is irreducible, it is $\mathbb{R}P^3$ and P is incompressible. By Lemma 5.3.25 all normal surfaces in M have non-zero beam degree. In particular, neither $b(G_1)$ nor $b(G_2)$ is zero, so $b(P)$ is strictly less than $b(F)$ and $p(P)$ is at most $p(F)$. Now, $\mathbb{R}P^3$ contains (up to isotopy) only one closed incompressible surface that does not contain a sphere: this projective plane. Thus F was not minimal.

Suppose that s_2 does *not* bound a disc in F. But then taking a parallel copy of E, bounded by s_2, we get a compression or ∂-compression disc for F.

Suppose that s_2 bounds a disc E' in F that is opposite to E at the intersection curve. If E' does not contain E, then $E \cup E'$ is a sphere that (as M is irreducible) bounds a 3-ball, so E and E' are isotopic. We can take the irregular switch at s instead to obtain (after normalisation) an isotopic surface to F which by Lemma 5.3.27 is of lower weight than F, so F was not minimal.

If E' does contain E, again taking the irregular switch along s, we get, if E was an interior disc, two surfaces $F_1 \cup T$. The surface F_1 is homeomorphic as a surface to F and $\partial F_1 = \partial F$, while T is formed by taking $E' \backslash E$ and identifying the boundary circles, and so is a torus or Klein bottle. As the switch was irregular, at least one of F_1 and T contains a return and hence after normalising has its weight decreased. Let F' be the result of normalising F_1. If the return is in F_1 then the weight of F' is less than that of F; if the return is in T then the weight of T is non-zero and so we can draw the same conclusion. As, since M is irreducible, E' is isotopic to a parallel copy of E, F and F' are isotopic and so F was not minimal.

If E intersected the boundary, we can follow a similar line of thought. Now after an irregular switch we get $F_1 \cup A$ where A is an annulus or Möbius band. The same reasoning however still holds.

The remaining case in which there are disc patches is when every disc patch has an adjacent companion disc. This is the case considered in Lemma 4.1.4 of [Mat07], where Matveev shows the following result in the triangulation setting:

Suppose that every disc patch of $G_1 \cup G_2$ has an adjacent companion disc E'. Then either F can be presented as $F = F_1 + T$ where T has Euler characteristic zero and F_1 is admissibly isotopic to F and of lower weight, or there is a disc patch E whose adjacent companion disc is also a patch of F.

We give the idea of Matveev's proof, which also goes through in our setting. It is to associate to $G_1 \cup G_2$ an oriented graph, whose vertices are disc patches and has an edge from E_1 to E_2 if E_2 is contained in the adjacent companion disc E_1' of E_1. As every vertex has at least one outgoing edge, this graph contains a cycle. Resolve $G_1 \cup G_2$ with the irregular switch at the edges of the patches in this cycle, and the regular switch otherwise. We get a decomposition $F = F_1 + T$ where T is the union of the pieces $E_i' - E_{i+1}$. If T is non-empty then as by Lemma 5.3.25 T has non-zero weight, F_1 is of lower weight than F. Note that F is homeomorphic to F_1 as if we perform regular switches everywhere, then at ∂E_i we replace E_i with $A_i \cup E_{i+1}$. Now,

$A_i \cup E_{i+1} \cup E_i$ is a disc or a sphere, so as M is irreducible and ∂-irreducible they separate off balls. We can use this to construct the desired isotopy.

As F is minimal, T must be empty, so there is a disc patch E whose adjacent companion disc E' is also a patch. Let G_1' and G_2' be the surfaces resulting from swapping E and E' within G_1 and G_2. Now, $E \cup E'$ is either a sphere or a properly-embedded disc, so as M is irreducible and ∂-irreducible, we can isotope G_1 to G_1' and G_2 to G_2', reducing the number of intersections in $G_1 \cup G_2$. Thus $G_1 + G_2$ was not in reduced form, as we assumed.

We can therefore conclude that there are no disc patches. Suppose D is a compression disc for some patch P, which we perturb to be transverse to F. Let γ be an innermost curve of $F \cap D$, which, as F is incompressible, bounds a disc D' in F. If $P \subseteq D'$ then P is planar and all but one of its boundary components bound discs in F. But then either P is a disc patch or the disc bounded by one of these boundary components contains a disc patch. Thus D' is disjoint from P so we can therefore isotope D to remove all intersections with F other than its boundary. Now, ∂D bounds a disc in F. This disc must be contained in P as otherwise some component of ∂P bounds a disc, and so there is a disc patch. Thus P is incompressible.

Suppose that D is a ∂-compression disc for P. By the same reasoning, we can isotope D to remove any curves of intersection with F. If α is an outermost arc of $F \cap D$ other than the segment of ∂D in P, as F is ∂-incompressible, it bounds a disc in F together with an arc in the boundary. By the same reasoning, this disc is disjoint from P so we can isotope D to further remove these arcs. Thus $D \cap P$ bounds a disc in F, which, as there are no disc patches, is contained in P. □

Proposition 5.3.32. *Let M be an irreducible ∂-irreducible 3-manifold and let $\mathcal{H}$ be a subtetrahedral split handle structure for M. Let F be a normal minimal surface that is incompressible and ∂-incompressible, with $F = G_1 + G_2$. Then G_1 and G_2 are also incompressible and ∂-incompressible, and neither G_1 nor G_2 is a copy of S^2, $\mathbb{R}P^2$ or a disc.*

Proof. Pick G_1 and G_2 as in the proof of Lemma 5.3.31. As M is irreducible and ∂-irreducible, F is not a disc or a sphere, as if it were, we could isotope it to a (non-normal) surface of zero weight in M, so F would not be minimal.

As any decomposition of a disc, sphere or projective plane has a disc patch, each component of G_1 and G_2 has nonpositive Euler characteristic. If G_1 is compressible or ∂-compressible, we will construct a nontrivial compression or ∂-compression disc D for G_1 that is disjoint from G_2. Let the *weight* of a compressing disc D for G_1 be

$w(D) = |D \cap G_2| + |\partial D \cap (G_2 \backslash\backslash \partial M)|$. The idea is to give moves to reduce this weight and thus produce a compressing disc for G_1 that is disjoint from G_2.

Claim 1: Suppose that G_1 has a nontrivial compression or ∂-compression disc D, which we can assume is transverse to G_2. If $D \cap G_2$ contains a closed curve or an arc with both endpoints in ∂M, then there is a compression or ∂-compression disc D' for G_1, of lower weight, that does not contain such intersections with G_2.

The triangulation setting analogue of this claim is shown in the first part of Lemma 4.1.35 of [Mat07]. We sketch the idea of the proof, which does not in fact depend on its setting. Take an innermost such curve or outermost such arc, bounding a disc Δ in D such that $\Delta \cap G_2 = \partial\Delta$ or $\Delta \cap G_2$ is a single arc. As $\partial\Delta \cap G_2$ is connected, it is contained in a single patch of G_2. Now, as the patches of $G_1 \cup G_2$ are incompressible and ∂-incompressible and M is irreducible and ∂-irreducible, there is a disc in this patch of G_2 isotopic to Δ. We can replace Δ in D with this disc (pushed off a little) to remove this intersection with G_2 and reduce the weight of D.

Let D be a nontrivial compression or ∂-compression disc for G_1 that intersects G_2 minimally. By the claim, D intersects G_2 only in arcs that have at least one endpoint away from ∂M. Let Δ be a region of D cut out by G_2. Then Δ is a (∂)-compressing disc for the polyhedron $G_1 \cup G_2$, as it is transverse to the singular subcomplex of $G_1 \cup G_2$ and $D \cap \partial M$ is connected. Following [Mat07, §4.1.6], label the vertices of Δ as "good" or "bad" angles according to whether, in F, the patches forming the two edges are pasted together or not respectively. If all angles of Δ are good, then Δ is a compression disc for F.

When formulated in the triangulation setting, the following claim is Lemma 4.1.33 of [Mat07].

Claim 2: Suppose that Δ has exactly one bad angle. Then F is not minimal.

The idea of the proof is to consider the component l of $G_1 \cap G_2$ that passes through this bad angle. Let A be the annulus or quadrilateral joining the two copies of l in F. One can build a (∂-)compressing disc for F from, if l is a closed curve, two copies of Δ and one of the components of $A - \Delta$, or if l is an arc, Δ and a component of $A - \Delta$. Now as F is incompressible and ∂-incompressible and M is irreducible and ∂-irreducible, this disc bounds a ball with a disc of F and possibly a disc in ∂M. We can use this to isotope F through this region to a surface of lower weight.

By this claim, each region has either no bad angles or at least two bad angles. Now, each arc in $D \cap G_2$ contributes one bad angle for each endpoint it has that is not on ∂M. Let m be the number of arcs that have no endpoints on ∂M, and n be the number with one endpoint on ∂M. There are $m + n + 1$ regions in D, and $2m + n$

bad angles. Thus, by the pigeonhole principle, there is a region Δ with at most one bad angle, which hence has no bad angles.

This region is a compression or ∂-compression disc for F itself, which is incompressible and ∂-incompressible, so $\Delta \cap F$ bounds a disc Δ' in F. We sketch the ideas of Lemma 4.1.34 of [Mat07], which deals with this situation. First, Δ' is decomposed by the trace curves into smaller regions. As there are no disc patches, none of these trace curves are closed curves or are arcs with both endpoints on ∂M. Thus, taking an outermost trace curve, there is a region Δ'_0 of Δ' whose boundary is one trace curve segment, one segment of $\Delta \cap F$, and possibly one segment from ∂M.

Now, as shown in Lemma 4.1.35 of [Mat07], if this region is in G_1, so $\partial\Delta'_0 = \Delta'_0 \cap (G_2 \cup \Delta \cup \partial M)$, then as M is irreducible we can isotope D through Δ'_0 to remove an arc of intersection with G_2. If this region is in G_2, then compress Δ along Δ'_0, removing an intersection. At least one of the resulting discs is essential and gives a non-trivial compression or ∂-compression disc. Either of these modifications reduces the weight of Δ. Continuing this procedure, we eventually get a non-trivial compression or ∂-compression disc for G_1 that is disjoint from G_2. But then the boundary of this disc is contained in a patch of G_1, which is incompressible and ∂-incompressible, so this is a contradiction. □

5.3.3 Fundamental normal surfaces

Our goal is to obtain a bound on collections of normal surfaces. In the triangulation setting, the exponential bound on the weight of fundamental normal surfaces is well-established (Lemma 2.3.6). If we were to take a fundamental surface and then wish to bound the size of a fundamental normal surface in its complement, the naïve approach would be to cut along it, take the induced handle structure, and find a fundamental surface there. However, as the first surface was bounded by an expression of the form $c^{|\mathcal{T}|}$, the induced handle structure may be of exponential size in the size of the original triangulation, so applying the usual theory, the bound on the second surface is of the form $c^{c^{|\mathcal{T}|}}$. We obtain a tower of exponentials, which is useless for any sort of polynomial time certificate. In this section we will use the subtetrahedral split handle structure theory that we have developed to instead give a bound of the form $c^{|\mathcal{T}|^2}$. As we noted at the start of Section 5.3, similar approaches have been used by Simon King [Kin01] and Marc Lackenby [Lac21].

Let $\mathcal{H}$ be a subtetrahedral split handle structure for a manifold M. If F is a normal surface in $\mathcal{H}$, we can uniquely describe F via a vector in $\mathbb{Z}^{d_H|\mathcal{H}|}$ as with the usual normal surface theory: each coordinate of the vector corresponds to a type of

elementary disc in a 0-handle of $\mathcal{H}$, and the entry at that coordinate gives the number of elementary discs of that type in F. By Lemma A.2, d_H is 59. A vector v in $\mathbb{Z}^{d_H|\mathcal{H}|}$ corresponds to some normal surface precisely when it satisfies the following conditions:

1. the entries of v are non-negative;
2. if we fix a parallelity piece P and a component of intersection C_1 of P with a 0-handle, for any other component C_2 of intersection of P with some 0-handle, the number of elementary discs that intersect C_1 from that 0-handle is equal to the number that intersect C_2 from the other 0-handle;
3. for each 1-handle $D^1 \times D^2$, which has two components of intersection with 0-handles, for each arc type in $\{0\} \times D^2$ (up to normal isotopy), the number of elementary discs in each 0-handle that intersect the 1-handle in that arc type are equal; and
4. for each pair of elementary discs in each 0-handle, if they have an essential intersection then v has a positive entry for at most one of them.

If t, u and v all are in this cone, v satisfies the fourth condition, and $v = u + t$, then u and t also satisfy the fourth condition. As the first three conditions can be expressed as linear equations, they restrict us to a polyhedral cone C in $\mathbb{Z}^{d_h|\mathcal{H}|}$, to which we can apply linear integer programming techniques, as in the triangulation case (Section 2.3). Using these ideas, we can apply Proposition 2.3.5 to obtain a bound.

Lemma 5.3.33. *There exists a constant c_F, which we can take to be 2^{480}, such that if G is a fundamental normal surface in a subtetrahedral split handle structure $\mathcal{H}$ then the size of G is bounded by $c_F^{|\mathcal{H}|}$.*

Proof. Note that n is at most $d_H|\mathcal{H}|$. The constraints from the parallelity piece equations give at most $6|\mathcal{H}|$ equations (one for each bridge), where the form of each equation is to set the sum of two lists of elementary discs to be equal, where each list is contained in one 0-handle, and so its ℓ^2 norm is at most $2d_H$.

As each 0-handle has at most four islands, there are at most $2|\mathcal{H}|$ 1-handles. Fix one of these islands, I. By Lemma 5.3.19 it intersects at most three bridges – let v be this number – and at most $6 - 2v$ sutures. As each part of the forbidden region that intersects I is bounded on each side by sutures or sides of bridges, of which we have a total of at most six, I has at most three components of intersection with the forbidden region. We thus have divided the boundary of I into at most 9 possible

segments for elementary discs to enter and leave by (corresponding to intersections with bridges and lakes, but not the forbidden region), where an arc of intersection of an elementary disc with I is determined by its entry and exit segments. There are thus at most $8 \cdot 9 = 72$ arc types in I from elementary discs. Thus each 1-handle gives at most 72 equations, of the same form as in the parallelity piece case: setting some sums of elementary disc types to be equal, which will be a row of A of ℓ^2 norm at most d_H.

The number of elementary discs of any given type of G is thus bounded by $(d_H|\mathcal{H}|)^{3/2} d_H^{d_H|\mathcal{H}|-1}$. We can then find a c_F as follows:

$$\begin{aligned} |G| &\leq |\mathcal{H}|(d_H|\mathcal{H}|)^{3/2} d_H^{d_H|\mathcal{H}|} \\ &= 2^{3/2\log(d_H)+5/2\log(|\mathcal{H}|)+d_H\log(d_H)|\mathcal{H}|} \\ &\leq 2^{2(d_H+1)\log(d_H)|\mathcal{H}|} \end{aligned}$$

so we can take c_F to be $2^{2(d_H+1)\log(d_H)}$ which, using $d_H = 59$ from Lemma A.2, is bounded above by 2^{480}. □

Lemma 5.3.34. *There exists a constant c_B, which we can take to be 2^{826}, such that the following holds. Let M be a manifold with a subtetrahedral split handle structure $\mathcal{H}$. Let $\{\Sigma_i\}$ be a collection of n disjoint normal surfaces in M such that, if we set $\mathcal{H}_0 := \mathcal{H}$ and $\mathcal{H}_i := \mathcal{H}_{i-1} \backslash\backslash \Sigma_i$, then Σ_{i+1} is non-duplicate and of size at most $2^{f(|\mathcal{H}_i|)}$ in the induced split handle structure on $\mathcal{H}_i$ for some (increasing) function f that does not depend on i and has $f(1) \geq 1$. Then there is a normal surface representative of the collection in $\mathcal{H}$ whose size is at most $c_B^{|\mathcal{H}|f(60|\mathcal{H}|)}$.*

Proof. By Lemma 5.3.15 $\mathcal{H}_i$ is a split handle structure for each i. Let $s(i)$ be the size of surfaces $\Sigma_i, \ldots, \Sigma_n$ included into $\mathcal{H}_{i-1}$. As Σ_{i+1} avoids the forbidden region of $\mathcal{H}_i$, it is disjoint from Σ_i, so their normal sum is their disjoint union. When we include the surfaces $\Sigma_{i+1}, \ldots, \Sigma_n$ into $\mathcal{H}_{i-1}$ and consider their size, two things happen: some of the parallelity pieces of $\mathcal{H}_i$ may be broken up into many 0-handle pieces and some of the elementary disc types of $\mathcal{H}_i$ may be identified (which does not change the total count of them in the surfaces). As the parallelity pieces had to have been been created from cutting along Σ_i, each sheet of the surfaces in a parallelity piece in $\mathcal{H}_i$ decomposes into at most $|\Sigma_i|$ 0-handle pieces in $\mathcal{H}_{i-1}$. As each sheet is adjacent to some elementary disc in a 0-handle of $\mathcal{H}_i$, and (as there are most six bridges in each of these 0-handles) each elementary disc runs through at most six bridges, the number of sheets is at most $6s(i+1)$. Thus when we include $\Sigma_{i+1}, \ldots, \Sigma_n$ into $\mathcal{H}_{i-1}$, the new size $s(i)$ is at most $|\Sigma_i| + s(i+1) + 6s(i+1)|\Sigma_i| \leq 8s(i+1)|\Sigma_i|$.

By Lemma 5.3.24, $|\mathcal{H}_i|$ is at most $60|\mathcal{H}|$ in $\mathcal{H}_{i+1}$ for all i. Thus $|\Sigma_i| \leq 2^{f(60|\mathcal{H}|)}$ for all i, so $s(i) \leq s(n)(8 \cdot 2^{f(60|\mathcal{H}|)})^{n-i}$. Thus we have that $s(0)$ is at most $8^n 2^{(n+1)f(60|\mathcal{H}|)}$. By a Kneser-Haken finiteness argument, as Σ_{i+1} is non-duplicate either it and another Σ_i are a pair of a non-orientable surface and its orientable double cover, or it contains some elementary disc that is not in a parallelity region, so n is at most $4d_H|\mathcal{H}|$. Now, we bound c_B as follows:

$$\begin{aligned} s(0) &\leq 2^{12d_H|\mathcal{H}|} 2^{(d_H|\mathcal{H}|+1)f(2^{13}|\mathcal{H}|)} \\ &\leq 2^{14d_H|\mathcal{H}|f(60|\mathcal{H}|)} \end{aligned}$$

which, as by Lemma A.2 d_H is 59, gives us the bound. □

Corollary 5.3.35. *There exists a constant c_S, which we can take to be 10^{10^7}, such that the following holds. Let M be a manifold with a subtetrahedral split handle structure $\mathcal{H}$. Let $\{\Sigma_i\}$ be a collection of n disjoint normal surfaces in M such that, if we set $\mathcal{H}_0 := \mathcal{H}$ and $\mathcal{H}_i := \mathcal{H}_{i-1} \backslash\backslash \Sigma_i$, then Σ_{i+1} is a non-duplicate normal surface in $\mathcal{H}_i$ which is either fundamental or the double of a fundamental surface. There is a normal surface representative of the collection in $\mathcal{H}$ whose size is at most $c_S^{|\mathcal{H}|^2}$.*

Proof. This immediately follows from Lemmas 5.3.33 and 5.3.34, as well as the fact that $c_B^{|\mathcal{H}|(480 \cdot 60|\mathcal{H}|+1)} \leq 2^{(1+826 \cdot 28800)|\mathcal{H}|^2} < 10^{10^7}$. □

5.4 Bounding normal vertical annuli

We will use the machinery we have just developed to bound the weight of a complete collection of normal vertical annuli.

5.4.1 Fundamental Möbius bands

First, we consider fundamental normal pseudo-vertical Möbius bands in subtetrahedral split handle structures.

Lemma 5.4.1. *If S is a Möbius band in an orientable irreducible 3-manifold M where ∂S is non-trivial in $\pi_1(M)$, then the double of S is essential unless M is $S^1 \times D^2$.*

Proof. As the double of S, $\tilde{S}$, is an annulus, there is only one loop (up to isotopy) in $\tilde{S}$ that does not bound a disc in $\tilde{S}$: one of its boundary components. As this is isotopic to ∂S, it is non-trivial in $\pi_1(M)$ so cannot bound a disc in M either, so $\tilde{S}$ is incompressible. As $\tilde{S}$ is the boundary of a neighbourhood of S, it is separating. Since

M is irreducible, an annulus is ∂-parallel if and only if it is ∂-compressible. In this case, we can observe that the boundary of a neighbourhood of a Möbius band is not, in fact, boundary-parallel in this neighbourhood. So if $\tilde{S}$ is not essential, it cuts M into $S^1 \times I \times I$ and $S \tilde{\times} I$, so M is $S \tilde{\times} I \cong S^1 \times D^2$. □

Lemma 5.4.2 (Lemma 2.4 [Mij04])**.** *Suppose that M is irreducible, ∂-irreducible, and is not $T^2 \times I$ or $K \tilde{\times} I$. Let S be a toroidal boundary component of M. If A and B are properly-embedded incompressible and ∂-incompressible annuli in M (that are not necessarily disjoint), such that at least one boundary component of each is on S, then these boundary components are isotopic.*

Lemma 5.4.3. *Suppose that M is a Seifert fibered space with non-empty boundary, not homeomorphic to a solid torus or $K \tilde{\times} I$, and has a subtetrahedral split handle structure $\mathcal{H}$, where the forbidden region (if non-empty) is a collection of vertical annuli. If M has any multiplicity two singular fibres, then for each boundary component T of M, there is a fundamental normal pseudo-vertical Möbius band containing one of these fibres whose boundary is on T.*

Proof. No Seifert fibration of $T^2 \times I$ has multiplicity two singular fibres. Let S be a pseudo-vertical Möbius band that is incompressible, ∂-incompressible, and disjoint from $\mathcal{I}$. Up to isotopy, we can consider it to be a union of regular fibres and a multiplicity two singular fibre. By Lemma 5.4.1 the double of S, $\tilde{S}$, is essential in M with empty forbidden region, and so is essential in $\mathcal{H}$ as well.

Let F be a minimal normal surface representative of S. Suppose that $F = G_1 + G_2$. By Proposition 5.3.32, G_1 and G_2 are incompressible and ∂-incompressible, and each has Euler characteristic at most zero, and so (as $\chi(F) = 0$) equal to zero. By Lemma 2.5.7 they are both horizontal or (pseudo-)vertical. If one is a Möbius band, say G_1, taking the double and observing by the same reasoning as above that it is essential, we can see by Lemma 5.4.2 that ∂G_1 is isotopic to ∂F. It is thus a pseudo-vertical Möbius band, so is (up to isotopy) a union of regular fibres and a single multiplicity two singular fibre. As the boundary of G_1 is a summand of that of F, ∂G_1 is contained in T. By Lemma 5.3.25 the weight of G_2 is non-zero so the weight of G_1 is strictly less than that of F.

As at least one of G_1 and G_2 has boundary and we have dealt with the Möbius band case, either both are annuli or one is an annulus and the other is closed. In either of these cases, we can apply Lemma 5.4.2 to the boundary components of the essential annuli to see that their boundary curves are isotopic to ∂F. There are an even number of boundary curves among G_1 and G_2, so $\partial G_1 + \partial G_2$ is trivial in $H_1(\partial M; \mathbb{Z}/2)$. But

∂F is not as it is a single curve, and normal sum is additive on homology with $\mathbb{Z}_2$ coefficients, so we have a contradiction. Thus (by induction) there is a fundamental pseudo-vertical Möbius band. □

5.4.2 A complete family of vertical annuli

Lemma 5.4.4. *Let M be a Seifert fibered space with non-empty boundary, other than the solid torus. Let $\mathcal{H}$ be a subtetrahedral split handle structure for M where the forbidden region $\mathcal{I}$ (if it is non-empty) is a collection of vertical annuli. Then $\mathcal{H}$ contains a fundamental normal surface, disjoint from the forbidden region, that is either an essential Möbius band or an essential non-duplicate annulus. Furthermore, if M is not homeomorphic to $T^2 \times I$ or $K \tilde{\times} I$ or $\mathcal{I}$ is not empty, this surface is vertical with respect to the Seifert fibered structure.*

Proof. When M has a multiplicity two singular fibre and is not $K \tilde{\times} I$, we proved this in Lemma 5.4.3. If M does have a multiplicity two singular fibre, then, we can assume that it is $K \tilde{\times} I$ with the $[D^2, 1/2, 1/2]$ Seifert fibration.

Let A be an incompressible ∂-incompressible annulus or Möbius band in M that is disjoint from the forbidden region and is not isotopic into it. This exists as M is not the solid torus. By Proposition 5.3.13 A has a normal representative F. As A is incompressible, it cannot be isotoped to be inside a single handle, so we can let F be of minimal weight in the admissible isotopy class of A (that is, up to isotopy fixing $\mathcal{I}$). Suppose that $F = G_1 + G_2$ as a non-trivial sum of normal surfaces, where we choose G_1 and G_2 to minimise $|G_1 \cap G_2|$. By Proposition 5.3.32, G_1 and G_2 are also incompressible and ∂-incompressible, and have Euler characteristic at most 0. If (say) G_1 has multiple connected components, so $F = G_1' + G_1'' + G_2$, then by resolving the sum $G_1'' + G_2$ (which we note does not add any curves of intersection with G_1') we get $F = G_1' + (G_1'' + G_2)$, where $|G_1 \cap G_2| > |G_1' \cap (G_1'' + G_2)|$, so both G_1 and G_2 must be connected. As $\chi(G_1) + \chi(G_2) = \chi(A) = 0$, both have Euler characteristic exactly zero – that is, each component is an annulus, Möbius band, Klein bottle or torus. By Lemma 5.3.28 neither is duplicate.

If one is closed, the other, say G_1, has the same boundary as A, so is an incompressible and ∂-incompressible non-duplicate annulus or Möbius band also. If both are essential annuli, then we have two essential non-duplicate annuli of smaller weight than A.

If G_1 is a Möbius band then by Lemma 2.5.7 it is either horizontal, in which case $\mathcal{I}$ must be empty and M must be $K \tilde{\times} I$, or pseudo-vertical so M is again $K \tilde{\times} I$ with the Seifert structure $[D^2, 1/2, 1/2]$.

In any of these cases, we have produced an incompressible and ∂-incompressible non-duplicate annulus or Möbius band that is of smaller weight than F, so there is a fundamental such surface G. If $\mathcal{I}$ is not empty then G cannot be horizontal, so must be (pseudo-)vertical and, if a Möbius band, have an essential vertical double. Otherwise, if G is horizontal, then by Lemma 2.5.8 M is homeomorphic to $K \tilde{\times} I$ or $T^2 \times I$. □

We will use Lemmas 5.3.33 and 5.3.34 to prove that we can obtain a complete collection of these annuli of bounded size. Recall that by Lemma 5.3.17, cutting along a normal surface produces another subtetrahedral split handle structure.

Proposition 5.4.5. *There exists a constant c_A, which we can take to be 10^{10^7}, such that the following holds. Let M be a Seifert fibered space with non-empty boundary, other than the solid torus. Let $\mathcal{H}$ be a subtetrahedral handle structure for M that is dual to a triangulation $\mathcal{T}$ such that the intersection of each tetrahedron of $\mathcal{T}$ with ∂M is connected and contractible. There is a maximal collection of disjoint normal essential annuli in $\mathcal{H}$, so that no two are isotopic, of total size at most $c_A^{|\mathcal{H}|^2}$. If M is not $T^2 \times I$ or $K \tilde{\times} I$, then these annuli are vertical.*

Proof. By Lemma 5.3.17, $\mathcal{H}$ can be viewed as a subtetrahedral split handle structure with empty forbidden region. Then by Lemma 5.4.4 there is a fundamental normal surface F in $\mathcal{H}$ that is either an essential Möbius band or an essential non-duplicate annulus. Using Lemma 5.3.17, $\mathcal{H}\backslash\backslash F$ is also a subtetrahedral split handle structure. If M is not $T^2 \times I$ or $K \tilde{\times} I$, then we can take the first annulus to be vertical, so the forbidden region of $\mathcal{H}\backslash\backslash F$ will be a non-empty collection of vertical annuli. We can then continue to apply Lemma 5.4.4 to generate a collection of disjoint normal essential annuli A_i and handle structures $\mathcal{H}_{i+1} = \mathcal{H}_i\backslash\backslash A_i$, where the annulus A_i is fundamental in $\mathcal{H}_i$ or is the double of a fundamental surface. This is the situation of Corollary 5.3.35, so there is a normal surface representative of the whole collection of size at most $\left(10^{10^7}\right)^{|\mathcal{H}|^2}$.

If M is $T^2 \times I$, then by Proposition 2.5.4 any essential annulus in $T^2 \times I$ is $\gamma \times I$ for some essential curve γ in T^2, and then the complement of this annulus is a solid torus with two parallel annuli as its forbidden region, so there are no further isotopy classes of essential annuli. If M is $K \tilde{\times} I$, it has a Seifert fibered structure as $S \tilde{\times} S^1$,

where S is the Möbius band. An essential annulus is either horizontal or vertical. If horizontal, it separates M into two copies of $S \tilde{\times} I$, which are each homeomorphic to a solid torus with a single annulus as their forbidden region so contain no essential annuli that are not isotopic to the one we already cut along. If vertical, it cuts M into a single solid torus, with two parallel annuli as its forbidden region, so again there are no further annuli. In these cases we took a single fundamental annulus, or double of a fundamental Möbius band, which is of size at most $2c_F^{|\mathcal{H}|}$ for $c_F = 2^{480}$ by Lemma 5.3.33, which is certainly smaller than bound on c_A we have given. □

Lemma 5.4.6. *Let $\mathcal{H}$ be a handle structure for an irreducible ∂-irreducible 3-manifold M that is dual to a triangulation $\mathcal{T}$, such that the intersection of each tetrahedron of $\mathcal{T}$ with ∂M is connected and contractible. Let F be a non-duplicate incompressible ∂-incompressible normal surface in $\mathcal{H}$. There is a normal surface F' in $\mathcal{T}$, isotopic to F, such that the weight of F' (that is, its number of intersections with the 1-skeleton) is at most $16|\mathcal{T}|s(F)$.*

Proof. By Lemma 5.3.17, $\mathcal{H}$ is a subtetrahedral split handle structure. Each elementary disc of F runs over at most four bridges, so its number of intersections with the 1-skeleton of $\mathcal{T}$ will be at most four. The barrier to directly including F into $\mathcal{T}$ is that an elementary disc in a subtetrahedral handle does not necessarily correspond to one in the dual tetrahedron: the ones that do not are those that run over vertices in the boundary of $\mathcal{T}$. It suffices then to perturb F enough that it is transverse to $\mathcal{T}$, without increasing its weight too much, as then, as F is incompressible and ∂-incompressible, we can normalise it which does not increase its weight.

Map F into $\mathcal{T}$ by sending each elementary disc of F to a disc in a tetrahedron of $\mathcal{T}$ that is transverse to the triangulation except that it may run over vertices in $\partial\mathcal{T}$. Each elementary disc of F ran over at most four bridges, so this surface intersects the 1-skeleton of $\mathcal{T}$ in at most $4s(F)$ points. At each vertex in $\partial\mathcal{T}$, consider its link L, which is a disc as M is a manifold. The elementary discs of F intersect L in a set of disjoint arcs. Pick a coherent choice of *direction* transverse to each arc, so that no two arcs point towards each other. At each arc, replace the portion of F that runs through the vertex with the subset of L in the chosen direction, pushed off a little so we obtain an embedding. This produces a surface F' which is isotopic to F and transverse to the triangulation. As L intersects each edge of $\mathcal{T}$ at most twice, this operation adds at most two points of intersection of each elementary disc with each edge of $\mathcal{T}$. There are at most $6|\mathcal{T}|$ edges so the weight of F' is at most $(4 + 12|\mathcal{T}|)s(F)$. Normalising F' produces a normal surface isotopic to F and of no greater weight than F'. □

Corollary 5.4.7. *There exists a constant c_T, which we can take to be $10^{10^{13}}$, such that the following holds. Let $\mathcal{T}$ be a triangulation of a Seifert fibered space M with non-empty boundary, other than the solid torus. There is a collection of disjoint normal essential annuli in $\mathcal{T}$, such that their complement is a collection of solid tori, of total weight at most $c_T^{|\mathcal{T}|^2}$. We can take these annuli to be vertical so long as M is not $T^2 \times I$ or $K \tilde{\times} I$.*

Proof. Consider the handle structure $\mathcal{H}$ that is dual to the 1^{st} barycentric subdivision of $\mathcal{T}$, $\mathcal{T}^{(1)}$, which contains $24|\mathcal{T}|$ tetrahedra. In $\mathcal{T}^{(1)}$ the intersection of each tetrahedron with ∂M is connected and contractible. As $\mathcal{H}$ satisfies the hypotheses of Proposition 5.4.5, there is a maximal collection C of disjoint normal essential annuli in $\mathcal{H}$, such that no two are isotopic, of total size at most $c_A^{|\mathcal{H}|^2}$, and so that if M is not $T^2 \times I$ or $K \tilde{\times} I$ then these annuli are vertical. We can now apply Lemma 5.4.6 to this collection C to obtain an isotopic set of normal surfaces in $\mathcal{T}^{(2)}$ of weight at most $16 \cdot 24^2 |\mathcal{T}| c_A^{24^4 |\mathcal{T}|^2} \leq c_A^{2 \cdot 24^4 |\mathcal{T}|^2}$ where $c_A < 10^{10^7}$. We can thus set c_T to be $c_A^{2 \cdot 24^4}$, which gives the desired bound.

If M is $T^2 \times I$ or $K \tilde{\times} I$, we took a single essential annulus that cuts M into one or two solid tori, so let this be the minimal collection. Otherwise, all the annuli are vertical. They lie over essential arcs in the base orbifold of M. If there are n singular fibres, there must be n annuli separating neighbourhoods of singulars fibre from the remainder of M. Consider the dual graph to the rest of the collection: it has a vertex for each region of $M \backslash\backslash C$, and an edge for each annulus, connecting the vertices of the regions it bounds. Take the complement of a spanning tree of this dual graph. This collection of annuli will cut the remainder of the orbifold into a single disc, and so will cut the remainder of M into a single solid torus. Take these annuli as well as the ones that cut off singular fibre neighbourhoods as the minimal collection. □

5.5 The certificate

We first prove that recognising a circle bundle over a surface with boundary is in **NP**. We prove a slightly stronger result, which says that we can certify the manifold and additionally certify a normal section and a normal fibre on each boundary component.

Lemma 5.5.1. *Let $\mathcal{T}$ be a triangulation of the solid torus. Then there is a fundamental normal meridian disc in $\mathcal{T}$ and a normal curve in $\partial\mathcal{T}$, of weight at most exponential in $|\mathcal{T}|$, that intersects the disc once.*

Proof. Corollary 6.4 of [JT95] tells us that if M is ∂-irreducible then there is a vertex normal essential disc, and the only essential disc in $S^1 \times D^2$ is the meridian disc. Take a normal curve γ in $\partial\mathcal{T}$ that intersects this meridian disc once. If $\gamma = \alpha_1 + \alpha_2$ as a normal sum, then one of the α_i must also intersect the disc once, so there is a fundamental normal curve in the triangulation on the boundary that satisfies this property. □

Proposition 5.5.2. *Deciding whether a 3-manifold M is an orientable circle bundle over a surface with non-empty boundary is in **NP**, as is, for such a manifold, deciding if it is a circle bundle over a certain surface (given as genus, number of boundary components and orientability). Furthermore, unless $M \cong K \tilde{\times} I$ or $T^2 \times I$, there exists a normal section and one normal fibre on each boundary component, such that these properties can be certified in time polynomial in $|\mathcal{T}|$.*

As we stated in Theorem 4.1.5, Haraway and Hoffman previously showed that certifying the cirlce bundle over the Möbius band ($K \tilde{\times} I$) and $A \times S^1$ (the I-bundle over the Möbius band) is in **NP**, and Matveev gave an algorithm for the problem that (although he did not discuss its complexity) also shows the problem is in **NP** [Mat07, Lemma 6.4.20]. We are reproving these results as we need to also certify the normal section and fibres.

Proof. Let $\mathcal{T}$ be a triangulation of M. The data given in the certificate is the following:

1. A compatible orientation of each tetrahedron;
2. If M is the solid torus, a fundamental normal disc, a curve in ∂M intersecting the disc once, and the data from Theorem 4.1.4 to certify that M is the solid torus;
3. If M is $T^2 \times I$, a normal annulus F of weight at most exponential in $|T|$, a triangulation of $M \backslash\backslash F$ constructed using Corollary 2.7.10, and a fundamental normal essential disc in this triangulation;
4. If M is $K \tilde{\times} I$ (which we consider to have the Seifert structure of an S^1-bundle over a Möbius band), a fundamental annulus F, and:
   (a) If the fundamental annulus is horizontal, the triangulation of $M \backslash\backslash F \cong S^1 \times D^2 \sqcup S^1 \times D^2$ from Corollary 2.7.10 and a fundamental normal essential disc in each component of the result;

(b) If the fundamental annulus is vertical, the triangulation of $M \backslash\backslash F \cong S^1 \times D^2$ from Corollary 2.7.10 and a fundamental normal essential disc in it;

5. Otherwise:

   (a) a collection of non-isotopic normal vertical annuli F, of total weight at most exponential in $|\mathcal{T}|^2$, whose complement is a solid torus;

   (b) the triangulation of $M \backslash\backslash F$ from Corollary 2.7.10;

   (c) a fundamental normal essential disc in this triangulation of $M \backslash\backslash F$;

   (d) a fundamental normal section in M.

Claim 1: The data of this certificate exists and has size bounded by a polynomial in $|\mathcal{T}|$.

Proof: When M is a solid torus, the data exists by Lemma 5.5.1.

When M is $T^2 \times I$, by Lemma 5.4.4 there is a normal annulus of the required size. Corollary 2.7.10 allows us to construct a triangulation of $M \backslash\backslash F$ of polynomial size in $|\mathcal{T}|$ as the weight of F is at most exponential in $|\mathcal{T}|$. By Lemma 5.5.1 then there is a fundamental normal essential disc in this new triangulation.

When M is $K \tilde{\times} I$, by Lemma 5.4.4 there is a fundamental normal essential annulus, which (as discussed in the proof of Proposition 5.4.5) is either horizontal and cuts M into two solid tori, or is vertical and cuts it into one solid torus. Either way we can use Corollary 2.7.10 to construct a triangulation of its complement that is of polynomial size in $|\mathcal{T}|$, and by Lemma 5.5.1 there is a fundamental normal essential disc in each component of this new triangulation.

In the general case, by Corollary 5.4.7 there is a collection of vertical essential annuli C in $\mathcal{T}$, of total weight at most $c_T^{|\mathcal{T}|^2}$, whose complement is a single solid torus. By Corollary 2.7.10 we can construct a triangulation of their complement, recording the boundary of C, that is of size at most polynomial in $|\mathcal{T}|$. By Lemma 5.5.1 there is a fundamental normal essential disc in this new triangulation. By Proposition 5.2.2 there is a fundamental section in M. ■

Claim 2: The homeomorphism type of M can be verified in polynomial time.

Proof: Check that M has boundary; that is, that some face of the triangulation is not identified with any other. Check that the given orientations of the tetrahedra are compatible, and thus certify that M is orientable.

We know that recognising the solid torus is in **NP** by Theorem 4.1.4. To certify the data from Lemma 5.5.1, check that the given surface is in fact a disc using Proposition 2.7.9 then check (using Lemma 2.4.2) that the algebraic intersection

number of the curve and the boundary of the disc, given as normal curves, is ± 1. We thus know that the curve and disc are essential, so must be a fibre and meridian disc.

If M is $T^2 \times I$, certify that F is an annulus using Proposition 2.7.9. Apply Corollary 2.7.10 to produce a triangulation of $M \backslash\backslash F$ and record the normal curves from ∂F in the boundary triangulation. We have already seen that we can certify that $M \backslash\backslash F$ is a solid torus. Compute the algebraic intersection numbers of the boundary curves from F with the boundary of the meridian disc, which (as the meridian curve is from a fundamental normal surface so is of at most exponential weight in $|\mathcal{T}|$) we can do in polynomial time by Lemma 2.4.2. Verify that they are each ± 1, so the core curve of F in $M \backslash\backslash F$ is a longitude. The mapping class group of the annulus up to non-∂-preserving isotopy has two elements: the class of the identity, and the class that exchanges the two boundary components, so M is either (in the first case) $T^2 \times I$ or (in the latter) $K \tilde{\times} I$. The first homology of $T^2 \times I$ is $\mathbb{Z}^2$ and that of $K \tilde{\times} I$ is $\mathbb{Z} \times \mathbb{Z}_2$, so it suffices to compute the homology of M, which, as the dimension of M is fixed, can be done in polynomial time by work of Iliopoulos [Ili89] as explained in [KP03].

Suppose M is $K \tilde{\times} I \cong S \tilde{\times} S^1$ where S is the Möbius band. If the essential annulus F is horizontal it covers the Möbius band, so as it has two boundary components it is a degree two cover and separates M into two copies of $S \tilde{\times} I$ (where it is the horizontal boundary of this bundle), which is the solid torus. As in the $T^2 \times I$ case, we can efficiently certify that this fundamental surface separates M into two solid tori. We can use the normal curve vectors for the boundary of F and the boundaries of the meridian discs to verify (using Lemma 2.4.2) that these curves have algebraic intersection number ± 2. Up to choice of coordinates there is only one curve on the boundary of the solid torus that intersects a meridian disc twice, so this is enough to certify that $\mathcal{T}$ is a triangulation of $K \tilde{\times} I$. If F is vertical, it sits over a spanning arc in the Möbius band and cuts M into a single solid torus. As in the $T^2 \times I$ case, certify that $M \backslash\backslash F$ is a solid torus, and produce a normal vector for a meridian curve. Compute the algebraic intersection numbers of the boundary of the two copies of F with this meridian curve. We find that they are ± 1 and by the same homology computation as in the $T^2 \times I$ case we can certify that M is $K \tilde{\times} I$.

We are left with the general case. By Proposition 2.7.9 we can quickly check that each surface in F is an annulus and compute the number of components n in F. As we have already described, we can certify that $M \backslash\backslash F$ is a solid torus and that the given surface is indeed a meridian disc, and record the corresponding meridian curve. Compute the intersection number of each annulus boundary and this meridian curve (using Lemma 2.4.2) to certify that the annulus core curves are longitudes in the solid

torus. Now, when we glue up the annuli we will get a circle bundle over a surface Φ. This surface has Euler characteristic $1-n$. As M is not a circle bundle over a disc, annulus or Möbius band, $1-n$ will be negative.

Let Σ be the normal surface that we claim is a normal section. Verify (using Proposition 2.7.9) that $\chi(\Sigma) = 1-n$. Compute algebraic intersection numbers (using Lemma 2.4.2) to certify that the boundary curve of Σ intersects each boundary curve of each annulus in A exactly once, which implies that Σ intersects each annulus in one spanning arc and possibly some trivial curves. This shows that, after compressions and ∂-compressions (which increase Euler characteristic), Σ is horizontal; as $\chi(\Phi) < 0$ and $\chi(\Sigma) = \chi(\Phi)$, Σ must be a degree one horizontal surface: that is, a section. ■

For the collection of fibres, we have a normal section in the surface bundle and a complete collection of vertical essential annuli. Take one annulus boundary component on each boundary component of M. □

We can now handle a special case of Seifert fibered space recognition.

Proposition 5.5.3. *Deciding if a 3-manifold M with triangulation $\mathcal{T}$ is a Seifert fibered space with only multiplicity two singular fibres and non-empty boundary, other than $S^1 \times D^2$, $T^2 \times I$ and $K \tilde{\times} I$, is in **NP**. When M is such a Seifert fibered space, deciding if it admits a certain set of Seifert data is in **NP** $\cap$ co-**NP** and there is a degree two normal horizontal surface and one normal fibre on each boundary component such that these properties have polynomial time certificates.*

Proof. The data given in the certificate is the following:

1. the Seifert data of M;
2. two collections of disjoint normal essential vertical annuli C in $\mathcal{T}$, of total weight at most $c^{|\mathcal{T}|^2}$, for a fixed constant c, where the complement of C is the union of a solid torus neighbourhood of each singular fibre and a circle bundle over a surface;
3. a triangulation of $\mathcal{T}\backslash\backslash C$ from Corollary 2.7.10, with a record of the normal curves from the boundaries of the components of C;
4. the data from the certificate in Proposition 5.5.2 for the circle-bundle-over-a-surface component of $\mathcal{T}\backslash\backslash C$;
5. a fundamental normal essential disc in each solid torus component of $\mathcal{T}\backslash\backslash C$;
6. a degree two horizontal fundamental normal surface F in $\mathcal{T}$;

Claim 1: The data of this certificate exists and is of size at most polynomial in $|\mathcal{T}|$.
Proof: Corollary 5.4.7 gives us the collection of normal vertical annuli. Use Corollary 2.7.10 to build the required triangulation of the complement, and then by Lemma 5.5.1 we can find a fundamental normal essential disc in each solid torus component of this triangulation. Proposition 5.5.2 gives us the certificate for $\mathcal{T}\backslash\backslash C$. By Proposition 5.2.2, the desired degree two horizontal fundamental normal surface exists. ■

Claim 2: The homeomorphism type of M can be verified from the certificate in polynomial time.
Proof: As in Proposition 5.5.2, certify that M has boundary and is orientable.

Build the triangulation of $\mathcal{T}\backslash\backslash C$ using Corollary 2.7.10, verifying that it agrees with the one in the certificate. We can then certify that all but one of the resulting pieces are solid tori using Theorem 4.1.4, and certify the homeomorphism type of the remaining piece M' using Proposition 5.5.2. By computing algebraic intersection numbers using Lemma 2.4.2, we can check that for each of the solid torus meridian discs, each annulus (algebraically) intersects it twice or not at all. We thus know that M is a copy of M' with solid tori glued on by gluing a $(1, 2)$ slope curve in the boundary of the solid torus to a fibre of M'; this is enough to certify that these solid tori are neighbourhoods of singular fibres of multiplicity two. ■

We know that the boundary curves of the annuli are fibres by construction, so take one on each boundary component. It remains to certify that F is a degree two horizontal surface. Check that ∂F is connected using Theorem 2.7.8, and then check that the algebraic intersection number of ∂F with each annulus is ± 2 using Lemma 2.4.2. From this information, we can deduce that F compresses and ∂-compresses to a degree two horizontal surface by Lemma 2.5.7, so $\chi(F) \leq 2\chi(\Sigma) - n$. Now check that $\chi(F) = 2\chi(\Sigma) - n$, so we could not have done any compressions or boundary compressions since they increase Euler characteristic. □

For the general case, we need one more result concerning computing singular fibre data.

Lemma 5.5.4. *Let M be an oriented Seifert fibered space with n of its singular fibres drilled out. Let F be a degree k horizontal surface in M, and let η_i, $1 \leq i \leq n$ be the collection of its curves of intersection with the boundary component arising from drilling out the i^{th} singular fibre. Let γ_i, $1 \leq i \leq n$, be a regular fibre on each of these boundary components. Pick orientations of η_i and γ_i such that the algebraic intersection number $i(\eta_i, \gamma_i)$ is positive with respect to the orientation of ∂M induced*

by the orientation of M. Then if μ_i is a meridian from the drilled out solid torus, the Seifert data of this singular fibre q/p (with respect to this orientation of M and the basis of $H_1(\partial M, \mathbb{Z})$ induced by $[\frac{1}{k}\eta_i]$ and γ_i for each i) is $(i(\eta_i, \mu_i)/k)/i(\gamma_i, \mu_i)$.

Proof. If we flip the orientation of both η_i and γ_i, this ratio of intersection numbers does not change sign.

Note that η_i is k copies of a curve, and $(\frac{1}{k}\eta_i, \gamma_i)$ is a positive basis for the homology of this torus, with its induced orientation. By the definition of the construction of Seifert fibered spaces (for example, see [Mar16, Defn. 10.3.1]), a (p_i, q_i) singular fibre is when the meridian curve is $\frac{p_i}{k}\eta_i + q_i\gamma_i$. □

Theorem 5.1.1. *The problem* SEIFERT FIBERED SPACE WITH BOUNDARY RECOGNITION *is in* ***NP****.*

Theorem 5.1.2. *The problem* NAMING SEIFERT FIBERED WITH BOUNDARY *is in* ***NP*** $\cap$ *co-****NP****.*

Proof of Theorems 5.1.1 and 5.1.2. Let M be a Seifert fibered space with non-empty boundary. Our certificate will be of the following form. If M is a circle bundle over a surface or a Seifert fibered space with only multiplicity two singular fibres it will be the relevant data from Proposition 5.5.2 or Proposition 5.5.3, respectively. Otherwise, it will be:

1. the Seifert data of M;
2. a compatible orientation of each tetrahedron of $\mathcal{T}$;
3. the triangulation $\mathcal{T}^{(82)}$, constructed by subdividing each tetrahedron of $\mathcal{T}$ in the order given;
4. the non-multiplicity-two singular fibres in $\mathcal{T}^{(82)}$, a solid torus neighbourhood of each one, a meridian disc for it with boundary of length at most 48, a longitude curve in its boundary that intersects the meridian once, a triangulation $\mathcal{T}'$ of $\mathcal{T}^{(82)}$ with these neighbourhoods of singular fibres removed and length at most 48 meridian curves marked and a compatible choice of orientations of the tetrahedra of $\mathcal{T}'$;
5. if there are singular fibres of multiplicity two, the certificate from Proposition 5.5.3 for $\mathcal{T}'$, or otherwise, the certificate from Proposition 5.5.2 for $\mathcal{T}'$.

Assume that M has at least one non-multiplicity-two singular fibre. It is straightforward to see that the certificate exists, as giving a triangulation of $\mathcal{T}^{(82)}$ is constructive, and Theorem 3.3.2 gives us the required singular fibre neighbourhoods.

To verify the certificate, first, as discussed in the proof of Proposition 5.5.2, check that M is orientable and has boundary.

Construct $\mathcal{T}^{(82)}$ by barycentrically subdividing the tetrahedra in order and verify that it agrees with the given triangulation. Certify that the removed regions are solid tori using Proposition 5.5.2 and use Proposition 2.7.9 to verify that the given meridian discs are indeed discs. Check that the longitude intersects the meridian curve once for each singular fibre neighbourhood using Lemma 2.4.2, thus certifying that the given meridian discs are essential.

Note that $\mathcal{T}' \ncong K \tilde{\times} I$ as M' has more than one boundary component. Since the only Seifert fibered structure for $T^2 \times I$ is $A \times S^1$, if $\mathcal{T}' \cong T^2 \times I$ then M was a solid torus as we can only have removed one singular fibre, so we can rule out this case. As we have the data of Proposition 5.5.2 or Proposition 5.5.3 to certify the Seifert data of $\mathcal{T}'$, since $\mathcal{T}' \ncong K \tilde{\times} I$ or $T^2 \times I$, take the normal horizontal surface and complete collection of fibres in the boundary contained in this certificate. Push the simplicial meridian ν off the 1-skeleton to get a normal meridian μ. The length of μ is at most the length of ν plus the total valence of the vertices of this boundary torus: that is, at most three times the number of edges in this boundary torus, which is at most $3 \cdot 24^{82} \cdot 6|\mathcal{T}|$.

Check that the given orientation of the tetrahedra of $\mathcal{T}'$ is in fact consistent and compute the orientation of each boundary triangle induced by it. In each boundary component of M' that bounds a singular fibre, consider the boundary of the horizontal surface, η, and the boundary of a normal vertical annulus fibre, γ, each of whose length is at most $c_T^{|\mathcal{T}|^2}$. By Lemma 2.4.2, we can arbitrarily orient each of the three curves μ, η and γ and then compute $i(\eta, \gamma)$ in polynomial time in the original input, with respect to the induced orientation of the boundary triangles. Set the orientation of η such that $i(\eta, \gamma)$ is positive and then compute $q' = i(\mu, \eta)$ and $p = i(\mu, \gamma)$. To correct for the fact that the horizontal surface may not have been a section, set $q = \frac{q'}{i(\eta,\gamma)}$. By Lemma 5.5.4 this singular fibre has Seifert data q/p. With the certificate for M', this certifies the homeomorphism type of M, and we can check that this matches the given Seifert data. For Theorem 5.1.2, note that as M has boundary it is not a lens space, so by Lemma 2.5.3 we can quickly check if the Seifert data of M describes a manifold that is homeomorphic to a given Seifert fibration. □

Chapter 6

Further questions

We will briefly present some questions that arise from the work in this thesis. First, we consider the recognition problem as discussed in Chapters 4 and 5. Recall the motivating question given in the introduction, which is well-known in the field:

Question A. *Is* 3-MANIFOLD HOMEOMORPHISM *in* ***NP****?*

We discussed a potential proof of a positive answer in Section 4.1.2 by first showing that the recognition problem for geometric 3-manifolds is in **NP**, and we stated there, the remaining geometric cases are hyperbolic, Sol, and closed manifolds with $\mathbb{H}^2 \times \mathbb{R}$ or $\widetilde{SL_2}$ geometries. Some special cases, then, of this question, are:

Question B. *Are the recognition and homeomorphism problems for closed Seifert fibered spaces with essential tori in* ***NP****?*

Question C. *Are the recognition and homeomorphism problems for Seifert fibered spaces without essential tori in* ***NP****?*

In Question B, we can take a fundamental normal essential torus (using much the same argument as in Lemma 5.4.4). We can cut along this torus using Proposition 2.7.7, and then by Theorem 5.1.1 we can certify the Seifert data of the resulting Seifert fibered space with boundary. This approach lets us reconstruct the Seifert data of the original closed manifold, up to the Euler number. Finding the Euler number is more subtle: it requires that we keep very careful track of the gluing map from the vertical torus. However, when we find a new triangulation using Proposition 2.7.7, we may dramatically change the boundary triangulation of the cut-open manifold, which makes this difficult.

In Question C we are considering Seifert fibered spaces whose base orbifolds are a sphere with at most three cone points or $\mathbb{R}P^2$ with at most one cone point. As

by work of Lackenby and Schleimer the recognition and naming problems for elliptic manifolds is in **NP**, we can restrict to the case of the sphere with three cone points. The difficulty here is that our approach so far has relied on taking certain essential surfaces, taking their minimal normal representative, and carefully analysing the complement. If there are no essential surfaces, we have no place to start.

For the recognition problem for the Sol geometry, we might hope to use the techniques of Section 4.2: considering the homology of torus bundles over S^1. Given such a manifold, there is a fundamental essential torus (by Lemma 4.2.4). Using Proposition 2.7.7 to build a triangulation of its complement, we can apply Theorem 4.1.5 to certify that it is $T^2 \times I$, and thus that the original manifold was a mapping torus over the torus. Thus the problem reduces to certifying Sol manifolds among Euclidean and Nil ones. As we discuss in the proof of Theorem 4.2.1, we can compute second homology to distinguish the Sol case from T^3 and from the case where the gluing map is conjugate to $\phi_e = \begin{pmatrix} 1 & e \\ 0 & 1 \end{pmatrix}$. We still need to distinguish it from the other Euclidean torus bundles, and from the gluing map being $\bar{\phi}_e = \begin{pmatrix} -1 & e \\ 0 & -1 \end{pmatrix}$.

The recognition and homeomorphism problems for hyperbolic manifolds are quite mysterious. If a hyperbolic manifold is fibered, we might hope to use this fibered structure to at least certify that the manifold is hyperbolic, if not to tackle the homeomorphism problem in full. Certifying that an irreducible, atoroidal manifold is fibered is known to be in **NP** [Sch11, Corr. 1.4]. When the hyperbolic manifold under consideration is not fibered, we know by Agol's proof of the Virtual Fibering Conjecture that it has a finite cover that is fibered. However, this proof is not effective: there is no known bound on the degree of this cover in general, so it would be very challenging to try to reduce the question to the fibered case.

Question D. *Is* 3-MANIFOLD HOMEOMORPHISM ***NP****-hard?*

This is the obvious companion problem to Question A, and is surely not original to this text. As we discussed in Section 2.1.1, several problems in 3-manifold topology are known to be **NP**-hard. The usual proof strategy for these results is to reduce some variant of SAT to them: the problem of, given a Boolean expression, determining if it is satisfiable. For example, Agol-Hass-Thurston (for bounding knot genus) and de Mesmay-Rieck-Sedgwick-Tancer (for recognising n-component unlink sublinks) prove their decision problems are **NP**-hard by using each term of a 3-SAT expression to construct, respectively, a piece of a branched surface or an arc of a link. The best

known result is Lackenby's proof that 3-manifold homeomorphism is at least as hard as graph isomorphism [Lac17, Theorem 1.2].

We now turn to the topic of triangulation complexity, as dealt with in Chapter 3.

Question E. *Is there a constant k such that, given a triangulation of a Seifert fibered space M where ∂M is not empty, all singular fibres of M are simplicial in the k^{th} barycentric subdivision of the triangulation? What about for a closed Seifert fibered space?*

Recall from Theorem 3.3.2 that given a triangulation $\mathcal{T}$ of a Seifert fibered space with non-empty boundary, the singular fibres that are not of multiplicity two are simplicial in $\mathcal{T}^{(79)}$. The technical issue with extending to multiplicity two in the case with boundary was with the analysis of the parallelity bundle, as an annulus separating such a singular fibre from the rest of the manifold is naturally the horizontal boundary of an I-bundle over a Möbius band. To answer Question E one would need to either extend the result to the general case, or give a set of examples where the length of a multiplicity two singular fibre transverse to the triangulation grows exponentially in the size of the triangulation.

In the closed case, the difficulty is that there is not an essential surface separating each singular fibre from the remainder of M: instead of an annulus, a surface bounding a singular fibre is a compressible torus. When M is a lens space (and thus is a Seifert fibered space over a sphere with two cone points) and is not $\mathbb{R}P^3$ or a prism manifold (that is, the union of $K \times I$ and a solid torus, glued along their boundaries), Lackenby and Schleimer proved that there are simplicial singular fibres in $\mathcal{T}^{(139)}$ for any triangulation $\mathcal{T}$.

Question F. *What is a good approximation of the triangulation complexity of a closed Seifert fibered space? What about the triangulation complexity of a graph manifold?*

This question asks for generalisations of Theorem 3.3.1. The natural upper bound for $\Delta(M)$ when M is closed and Seifert fibered, by an analogous construction to Proposition 3.3.15, is $\Delta(M) \leq c(|\chi(\Sigma) + 1 + \sum_{i=1}^{n} \|q_i/p_i\| + \|e\|)$, where e is the Euler number. While (so long as M has an essential vertical torus) we can use Theorem 3.3.1 and the ideas in this thesis to lower bound $\Delta(M)$ in terms of $\chi(\Sigma)$ and $\|q_i/p_i\|$, understanding its relation to the Euler number does not immediately follow. Graph manifolds present similar challenges: lower-bounding their complexity in terms of the sum of the lower bounds on the complexities of their JSJ decompositions is straightforward, but this does not take the gluing maps into account. For the closed

case, Lackenby and Purcell gave formulas for the triangulation complexity of elliptic manifolds (implicitly in terms of their Seifert data) and showed that this naïve upper bound is approximately correct (up to a multiplicative constant) [LP22a, Theorems 1.1 and 1.2].

Question G. *What is a bound for k in Theorem 3.3.1?*

The reason that the constant in Theorem 3.3.1 is not constructive is that the lower bound uses the constant in Lemma 2.8.9, which uses Theorem 9.1 of [LP22a] which itself relies on Theorem 1.6 of the same paper, which uses Theorem 8.2 of that paper, which finally relies on the number $Q(S)$ in Theorem 6.1 of [MMS12]. Answering Question G would involve bounding $Q(S)$ in terms of the surface S.

Appendix A

Bounds on elementary disc types in subtetrahedral split 0-handles

We wish to bound the number of elementary disc types (up to normal isotopy) in a subtetrahedral split 0-handle. Recall from Definition 5.3.16 that a 0-handle H in a split handle structure is *subtetrahedral* if we can obtain it by taking a semitetrahedral 0-handle H' (that is, one whose boundary graph is a connected subgraph of the complete graph on four vertices, so is shown in Figure A.1) and cutting it open along elementary discs.

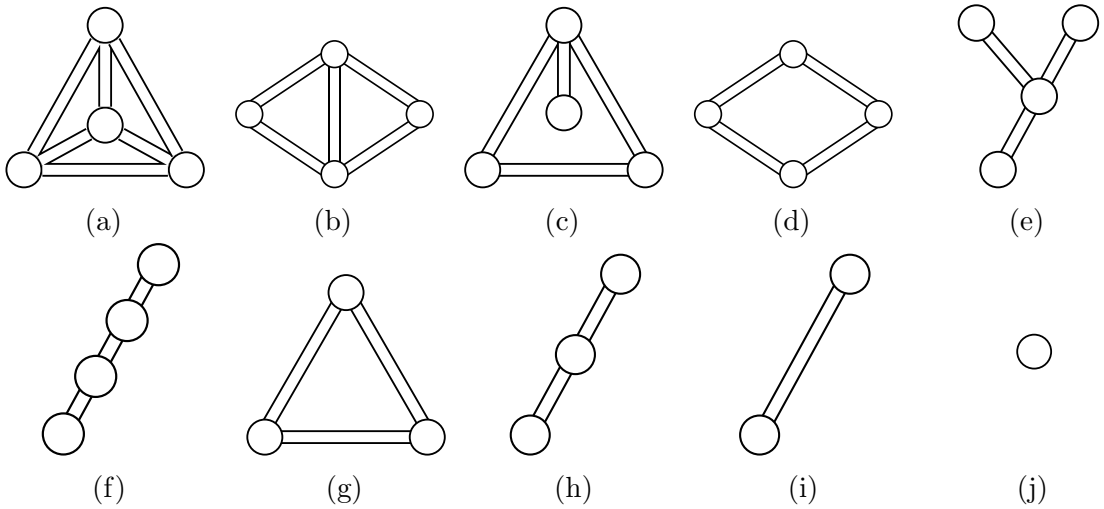


Figure A.1: The semitetrahedral split 0-handles.

An elementary disc is determined by its boundary in ∂H. From Definitions 5.3.3 and 5.3.4, a properly-embedded disc in a 0-handle is elementary if the following statements are true of its boundary:

1. it is disjoint from the 3-handles and the forbidden region,

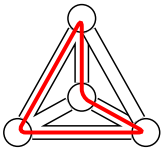
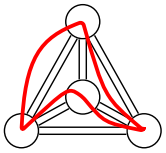
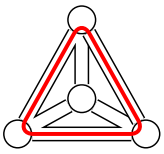
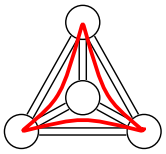
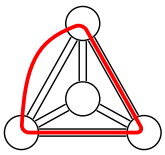
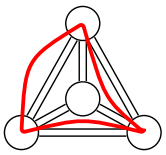
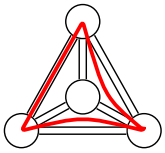
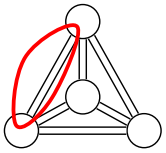
Elementary disc homeomorphism type	Number of types up to normal isotopy	Elementary disc homeomorphism type	Number of types up to normal isotopy
	3		6
	4		4
	12		12
	12		6

Table A.1: Homeomorphism types of elementary discs with boundary graph the complete graph on four vertices, and the number of normal isotopy types of each.

2. it is not a closed curve contained in a single bridge, lake or island,
3. it runs along the bridges (not doubling back or leaving through a lake),
4. it crosses each bridge or lake at most once,
5. if a bridge and lake are adjacent, it intersects at most one of the pair,
6. no arc of its intersection with an island starts and ends on the same connected component of the island with a lake, and
7. it does not contain any arc in a lake that starts and ends on the same connected component of the intersection of the lake with an island.

Lemma A.1. *The number of elementary disc types in H is bounded by the number in H'.*

Proof. Let C be the collection of elementary discs in H' that we cut along to obtain H. Any elementary disc D in H naturally includes into H' as some properly-embedded

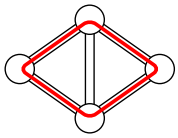
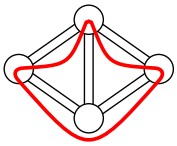
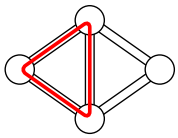
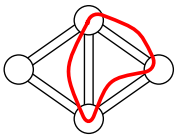
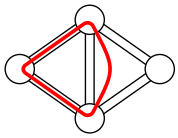
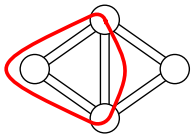
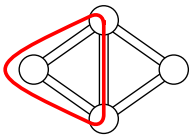
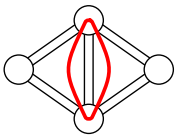
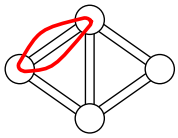
Elementary disc homeomorphism type	Number of types up to normal isotopy	Elementary disc homeomorphism type	Number of types up to normal isotopy
	1		2
	2		4
	2		2
	1		1
	4		

Table A.2: Homeomorphism types of elementary discs in a 0-handle with boundary graph the complete graph on four vertices with an edge removed, and the number of normal isotopy types of each.

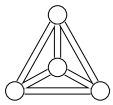
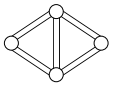
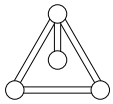
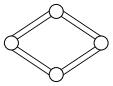
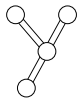
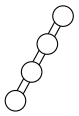
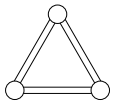
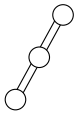
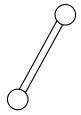
Boundary cell structure	Disc types up to homeomorphism	Disc types up to normal isotopy	Boundary cell structure	Disc types up to homeomorphism	Disc types up to normal isotopy
	8	59		9	19
	4	6		3	7
	0	0		0	0
	2	4		0	0
	0	0		0	0

Table A.3: Counts of elementary disc types in semitetrahedral 0-handles.

disc D', as ∂D is disjoint from the forbidden region. We claim that it does so as an elementary disc.

Certainly, $\partial D'$ satisfies conditions (1) and (3). As ∂D is not contained in a single bridge, lake or island, it intersects them in arcs, so in H' it also runs through the bridges, lakes and islands in arcs and thus satisfies (2). As the discs in C are elementary, the image of any single bridge or lake in H' is connected in H so $\partial D'$ satisfies condition (4). If $\partial D'$ runs through both a bridge and an adjacent lake, the corresponding two arcs of ∂D must run through a bridge and lake which are not adjacent. The only way this is possible is if (since the image of the bridge and lake are connected) there is a disc in C whose boundary runs through that bridge and lake (so cutting along the disc in C replaces the boundary between the bridge and lake with a patch of the forbidden region), but as the discs in C are elementary, this is impossible. Thus $\partial D'$ satisfies condition (5). For the last two conditions, consider a component of intersection of an island with a lake. As no disc in C runs through this component twice, and every disc in C is separating, we can see by induction on C that the image of this component is connected in H, and so $\partial D'$ satisfies conditions (6) and (7).

Second, we claim that if two elementary discs D_1 and D_2 in H are not normally isotopic, then their images D_1' and D_2' in H are also not normally isotopic. Suppose

that D'_1 and D'_2 were normally isotopic. As D'_1 and D'_2 are disjoint, the region of $\partial H' - (\partial D'_1 \cup \partial D'_2)$ that lies between the two of them is an annulus, and as they are normally isotopic, it is a product handle. If there were a disc of C in this product region, it would also be normally isotopic to D'_1 and D'_2 and so it would separate them, and one of them would not be in H. Thus there is no such disc, and so this annulus also exists in H, so the two discs are normally isotopic there as well. □

To bound the number of elementary disc types in a subtetrahedral split 0-handle, it thus suffices to bound the number in a semitetrahedral split 0-handle. We do this by enumerating them. In Figure A.1 we give the possible boundary graphs of a semitetrahedral 0-handle in a split handle structure. Most of these handles allow for a very small number of elementary discs; for example, those whose boundary graphs are trees support none at all. The ones with more than four elementary disc types up to homeomorphism are the complete graph on four vertices, and the same graph with one edge removed. In Tables A.1 and A.2 we enumerate the elementary disc types up to homeomorphism in these cases, and give, for each of these, the number of types up to normal isotopy. Through these techniques, we obtain Table A.3, which gives the counts of elementary discs up to homeomorphism and up to normal isotopy for each semitetrahedral split 0-handle. From this enumeration we obtain the following bound.

Lemma A.2. *A subtetrahedral split 0-handle supports at most 59 elementary disc types up to normal isotopy.*

Bibliography

[AHT06] Ian Agol, Joel Hass, and William Thurston. "The computational complexity of knot genus and spanning area". In: *Trans. Amer. Math. Soc.* 358.9 (2006), pp. 3821–3850.

[AB09] Sanjeev Arora and Boaz Barak. *Computational complexity*. New York, NY, USA: Cambridge University Press, 2009.

[BDS17] David Bachman, Ryan Derby-Talbot, and Eric Sedgwick. "Computing Heegaard genus is NP-hard". In: *A journey through discrete mathematics: a tribute to Jiří Matoušek*. Ed. by Martin Loebl, Jaroslav Nešetřil, and Robin Thomas. Cham: Springer International Publishing, 2017, pp. 59–87. DOI: 10.1007/978-3-319-44479-6_3.

[Bak] Kenneth L. Baker. *The one tetrahedron, one vertex triangulation of the solid torus*. https://sketchesoftopology.wordpress.com/2008/05/15/the-one-tetrahedron-one-vertex-triangulation-of-the-solid-torus/.

[BS19] John A. Baldwin and Steven Sivek. "On the complexity of torus knot recognition". In: *Trans. Amer. Math. Soc.* 371.6 (2019), pp. 3831–3855.

[BF23] Federica Bertolotti and Roberto Frigerio. *Integral filling volume, complexity and integral simplicial volume of 3-dimensional mapping tori*. 2023. arXiv: 2303.07730 [math.GT].

[Bes+10] Laurent Bessières, Gérard Besson, Michel Boileau, Sylvain Maillot, and Joan Porti. *Geometrisation of 3-manifolds*. Vol. 13. EMS Tracts in Mathematics. Zürich, Switzerland: European Mathematical Society, 2010.

[Bra21] H. Roy Brahana. "Systems of circuits on two-dimensional manifolds". In: *Ann. of Math. (2)* 23.2 (1921), pp. 144–168.

[Bri93] Matthew G. Brin. *Seifert fibered spaces: notes for a course given in the spring of 1993*. 1993. arXiv: 0711.1346 [math.GT].

[BFP15] Michelle Bucher, Roberto Frigerio, and Cristina Pagliantini. "The simplicial volume of 3-manifolds with boundary". In: *J. Topol.* 8.2 (2015), pp. 457–475.

[Bur11] Benjamin A. Burton. "Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations". In: *Proceedings of the 36th international symposium on symbolic and algebraic computation*. ISSAC '11. San Jose, California, USA: Association for Computing Machinery, 2011, pp. 59–66. DOI: 10.1145/1993886.1993901.

[Bur20] Benjamin A. Burton. "The next 350 million knots". In: *36th international symposium on computational geometry (SoCG 2020)*. Ed. by Sergio Cabello and Danny Z. Chen. Vol. 164. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020, 25:1–25:17. DOI: 10.4230/LIPIcs.SoCG.2020.25.

[BBP+22] Benjamin A. Burton, Ryan Budney, William Pettersson, et al. *Regina: software for low-dimensional topology*. http://regina-normal.github.io/. 1999–2022.

[Cas65] B.G. Casler. "An imbedding theorem for connected 3-manifolds". In: *Proc. Amer. Math. Soc.* 16 (1965), pp. 559–566.

[EN13] Jeff Erickson and Amir Nayyeri. "Tracing compressed curves in triangulated surfaces". In: *Discrete Comput. Geom.* 49.4 (2013), pp. 823–863.

[FV11] E. A. Fominykh and A. Yu. Vesnin. "Exact values of complexity for Paoluzzi-Zimmermann manifolds". In: *Dokl. Math.* 84.1 (2011), pp. 542–544.

[FW13] E.A. Fominykh and Bert Wiest. "Upper bounds for the complexity of torus knot complements". In: *J. Knot Theory Ramifications* 22.10 (2013), p. 1350053.

[Fro86] Charles Frohman. "One-sided incompressible surfaces in Seifert fibered spaces". In: *Topology Appl.* 23 (1986), pp. 103–116.

[GH75] Cameron McAllen Gordon and Wolfgang Heil. "Cyclic normal subgroups of fundamental groups of 3-manifolds". In: *Topology* 14.4 (1975), pp. 305–309. DOI: 10.1016/0040-9383(75)90014-2.

[Gro82] Michael Gromov. "Volume and bounded cohomology". In: *Publ. Math. Inst. Hautes Études Sci.* 56 (1982), pp. 5–99.

[Hak61a] Wolfgang Haken. "Ein Verfahren zur Aufspaltung einer 3-Mannigfaltigkeit in irreduzible 3-Mannigfaltigkeiten". In: *Math. Z.* 76 (1961), pp. 427–467.

[Hak61b] Wolfgang Haken. "Theorie der Normalflächen". In: *Acta Math.* 105.3 (1961), pp. 245–375.

[HH22] Robert Haraway III and Neil R Hoffman. *On the complexity of cusped non-hyperbolicity*. 2022. arXiv: 1907.01675 [math.GT].

[HLP99] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. "The computational complexity of knot and link problems". In: *J. ACM* 46.2 (1999), pp. 185–211.

[HP22] Neil R Hoffman and Kathleen L Petersen. *Small $PSL(2,\mathbb{F})$ representations of Seifert fiber space groups*. 2022. arXiv: 2209.05478 [math.GT].

[Ili89] Costas S. Iliopoulos. "Worst-case complexity bounds on algorithms for computing the canonical structure of finite abelian groups and the Hermite and Smith normal forms of an integer matrix". In: *SIAM J. Comput.* 18.4 (1989), pp. 658–669.

[Iva08] S.V. Ivanov. "The computational complexity of basic decision problems in 3-dimensional topology". In: *Geom. Dedicata* 131 (2008), pp. 1–26.

[Jac23a] Adele Jackson. *Minimal triangulation size of Seifert fibered spaces with boundary*. 2023. arXiv: 2301.02085 [math.GT].

[Jac23b] Adele Jackson. *Recognition of Seifert fibered spaces with boundary is in NP*. 2023. arXiv: 2306.04612 [math.GT].

[JJST16] William Jaco, Jesse Johnson, Jonathan Spreer, and Stephan Tillmann. "Bounds for the genus of a normal surface". In: *Geom. Topol.* 20.3 (2016), pp. 1625–1671.

[JRST19] William Jaco, Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann. "On minimal ideal triangulations of cusped hyperbolic 3-manifolds". In: *J. Topol.* 13.1 (2019), pp. 308–342.

[JRT09] William Jaco, Hyam Rubinstein, and Stephan Tillmann. "Minimal triangulations for an infinite family of lens spaces". In: *J. Topol.* 2.1 (2009), pp. 157–180.

[JRT11] William Jaco, J Hyam Rubinstein, and Stephan Tillmann. "Coverings and minimal triangulations of 3–manifolds". In: *Algebr. Geom. Topol.* 11.3 (2011), pp. 1257–1265.

[JR06] William Jaco and J. Hyam Rubinstein. *Layered-triangulations of 3-manifolds*. 2006. arXiv: math/0603601 [math.GT].

[JRST20] William Jaco, J. Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann. "$\mathbb{Z}_2$-Thurston norm and complexity of 3-manifolds, II". In: *Algebr. Geom. Topol.* 20 (2020), pp. 503–529.

[JRST22] William Jaco, J. Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann. *Complexity of 3-manifolds obtained by Dehn filling*. 2022. arXiv: 2207.12066 [math.GT].

[JT95] William Jaco and Jeffrey L. Tollefson. "Algorithms for the complete decomposition of a closed 3-manifold". In: *Illinois J. Math.* 39.3 (1995), pp. 358–406.

[KP03] Volker Kaibel and Marc E. Pfetsch. "Some algorithmic problems in polytope theory". In: *Algebra, geometry and software systems*. Ed. by Michael Joswig and Nobuki Takayama. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, pp. 23–47.

[Kin01] Simon A. King. “Polytopality of triangulations”. PhD thesis. Université Louis Pasteur, Strasbourg, June 2001.

[KT21] Dale Koenig and Anastasiia Tsvietkova. “NP-hard problems naturally arising in knot theory”. In: *Trans. amer. math. soc. ser. b* 8 (2021), pp. 420–441.

[Kup19] Greg Kuperberg. “Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization”. In: *Pacific J. Math.* 301.1 (2019), pp. 189–241.

[Lac09] Marc Lackenby. “The crossing number of composite knots”. In: *J. Topol.* 2.4 (2009), pp. 747–768.

[Lac14] Marc Lackenby. “Core curves of triangulated solid tori”. In: *Trans. Amer. Math. Soc.* 366.11 (2014), pp. 6027–6050.

[Lac17] Marc Lackenby. “Some conditionally hard problems on links and 3-manifolds”. In: *Discrete Comput. Geom.* 58.3 (Aug. 2017), pp. 580–595.

[Lac21] Marc Lackenby. “The efficient certification of knottedness and Thurston norm”. In: *Adv. Math.* 387 (2021), p. 107796.

[LP22a] Marc Lackenby and Jessica S. Purcell. *The triangulation complexity of elliptic and sol 3-manifolds*. 2022. arXiv: `2212.04683 [math.GT]`.

[LP22b] Marc Lackenby and Jessica S. Purcell. *The triangulation complexity of fibred 3-manifolds*. 2022. arXiv: `1910.10914 [math.GT]`.

[LS22] Marc Lackenby and Saul Schleimer. *Recognising elliptic manifolds*. 2022. arXiv: `2205.08802 [math.GT]`.

[LY23] Marc Lackenby and Mehdi Yazdi. “The computational complexity of knot genus in a fixed 3-manifold”. In: *Proc. London Math. Soc.* 126.3 (2023), pp. 837–879. DOI: `https://doi.org/10.1112/plms.12500`. eprint: `https://londmathsoc.onlinelibrary.wiley.com/doi/pdf/10.1112/plms.12500`.

[Li11] Tao Li. “An algorithm to determine the Heegaard genus of a 3–manifold”. In: *Geom. Topol.* 15.2 (2011), pp. 1029–1106. DOI: `10.2140/gt.2011.15.1029`.

[Löh18] Clara Löh. “Odd manifolds of small integral simplicial volume”. In: *Ark. Mat.* 56.2 (2018), pp. 351–375.

[Mar58] A. A. Markov. “The insolubility of the problem of homeomorphy”. In: *Dokl. Akad. Nauk SSSR* 121.2 (1958), pp. 218–220.

[Mar16] Bruno Martelli. *An introduction to geometric topology*. 2016. arXiv: `1610.02592 [math.GT]`.

[MMS12] Howard A. Masur, Lee Mosher, and Saul Schleimer. “On train-track splitting sequences”. In: *Duke Math. J.* 161.9 (2012), pp. 1613–1656.

[Mat90] Sergei Vladimirovich Matveev. “Complexity theory of three-dimensional manifolds”. In: *Acta Appl. Math.* 19.2 (1990), pp. 101–130.

[Mat07] Sergei Vladimirovich Matveev. *Algorithmic topology and classification of 3-manifolds*. second. Vol. 9. Algorithms and Computation in Mathematics. NY: Springer-Verlag, 2007.

[MT20] Sergei Vladimirovich Matveev and Vladimir Viktorovich Tarkaev. "Recognition and tabulation of 3-manifolds up to complexity 13". In: *Chebyshevskiĭ Sb.* 21.2 (2020), pp. 290–300.

[MRST21] Arnaud de Mesmay, Yo'av Rieck, Eric Sedgwick, and Martin Tancer. "The unbearable hardness of unknotting". In: *Adv. Math.* 381 (2021), p. 107648. DOI: https://doi.org/10.1016/j.aim.2021.107648.

[MSS20] Arnaud de Mesmay, Marcus Schaefer, and Eric Sedgwick. "Link crossing number is np-hard". In: *Journal of knot theory and its ramifications* 29.06 (2020), p. 2050043. DOI: 10.1142/S0218216520500431. eprint: https://doi.org/10.1142/S0218216520500431.

[Mij04] Aleksandar Mijatović. "Triangulations of Seifert fibred manifolds". In: *Math. Ann.* 330 (2004), pp. 235–273.

[Per02] Grisha Perelman. *The entropy formula for the Ricci flow and its geometric applications*. 2002. arXiv: math/0211159 [math.DG].

[Per03a] Grisha Perelman. *Finite extinction time for the solutions to the Ricc flow on certain three-manifolds*. 2003. arXiv: math/0307245 [math.DG].

[Per03b] Grisha Perelman. *Ricci flow with surgery on three-manifolds*. 2003. arXiv: math/0303109 [math.DG].

[Prz19] Józef H. Przytycki. "Nonorientable, incompressible surfaces in punctured-torus bundles over S^1". In: *Rev. R. Acad. Cienc. Exactas Fís. Nat. Ser. A Mat. RACSAM* 113 (2019), pp. 1–26.

[Ran96] Richard Rannard. "Incompressible surfaces in Seifert fibered spaces". In: *Topology Appl.* 72 (1996), pp. 19–30.

[Rub78] J. Hyam Rubinstein. "One-sided Heegaard splittings of 3-manifolds". In: *Pacific J. Math.* 76.1 (1978), pp. 185–200.

[RST21] J. Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann. *A new family of minimal ideal triangulations of cusped hyperbolic 3-manifolds*. 2021. arXiv: 2112.01654 [math.GT].

[SSŠ02] Marcus Schaefer, Eric Sedgwick, and Daniel Štefankovič. "Algorithms for normal curves and surfaces". In: *Computing and combinatorics*. Ed. by Oscar H. Ibarra and Louxin Zhang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 370–380.

[Sch89] Martin Scharlemann. "Sutured manifolds and generalized Thurston norms". In: *J. Differential Geom.* 29.3 (1989), pp. 557–614.

[Sch01] Saul Schleimer. "Almost normal Heegaard splittings". PhD thesis. University of California, Berkeley, 2001.

[Sch11] Saul Schleimer. "Sphere recognition lies in NP". In: *Low-dimensional and symplectic topology*. Ed. by Michael Usher. Vol. 82. Providence, Rhode Island: American Mathematical Society, 2011, pp. 183–213.

[Sco83a] Peter Scott. "The geometries of 3-manifolds". In: *Bull. Lond. Math. Soc.* 15.5 (1983), pp. 401–487.

[Sco83b] Peter Scott. "There are no fake Seifert fibre spaces with infinite π_1". In: *Ann. of Math. (2)* 117.1 (1983), pp. 35–70.

[SS14] Peter Scott and Hamish Short. "The homeomorphism problem for closed 3–manifolds". In: *Algebr. Geom. Topol.* 14.4 (2014), pp. 2431–2444.

[Scu21] Joe Scull. *The homeomorphism problem for hyperbolic manifolds I*. 2021. arXiv: 2108.00779 [math.GT].

[Sel95] Zlil Sela. "The isomorphism problem for hyperbolic groups I". In: *Ann. of Math. (2)* 141.2 (1995), pp. 217–283.

[Sno90] Jack Snoeyink. "A trivial knot whose spanning disks have exponential size". In: *Proceedings of the sixth annual symposium on computational geometry*. SCG '90. Berkley, California, USA: Association for Computing Machinery, 1990, pp. 139–147. DOI: 10.1145/98524.98555.

[Tan23] Martin Tancer. *Simpler algorithmically unrecognizable 4-manifolds*. 2023. arXiv: 2310.07421 [math.GT].

[VKF74] I. A. Volodin, V. E. Kuznetsov, and A. T. Fomenko. "The problem of discriminating algorithmically the standard three-dimensional sphere". In: *Russian Math. Surveys* 29.5 (1974), pp. 71–172. DOI: 10.1070/RM1974v029n05ABEH001296.

[Wal67a] Friedhelm Waldhausen. "Ein Klasse von 3-dimensionalen Mannigfaltigkeiten I". In: *Invent. Math.* 3 (1967), pp. 308–333.

[Wal67b] Friedhelm Waldhausen. "Ein Klasse von 3-dimensionalen Mannigfaltigkeiten II". In: *Invent. Math.* 4 (1967), pp. 87–117.

[Wel93] Dominic J.A. Welsh. "The complexity of knots". In: *Quo vadis, graph theory?* Ed. by John Gimbel, John W. Kennedy, and Louis V. Quintas. Vol. 55. Annals of Discrete Mathematics. Elsevier, 1993, pp. 159–171. DOI: https://doi.org/10.1016/S0167-5060(08)70385-6.

[Zen18] Raphael Zentner. "Integer homology 3-spheres admit irreducible representations in $\mathrm{SL}(2,\mathbb{C})$". In: *Duke Math. J.* 167.9 (2018), pp. 1643–1712.