Certifying hyperbolicity of fibred 3-manifolds

Filippo Baroni

St Hilda's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

June 2025

To Mum and Dad.
In the darkest of times,
your candle never flickered.

Acknowledgements

It is customary, when writing a thesis, to relegate the acknowledgements to the last few frantic hours before submission. Whatever time remains after countless rounds of panicked proofreading and systematic second-guessing is hastily allocated to condensing many years of unique experiences, struggles, and connections into a few uninspired and painfully generic paragraphs. Or so they tell me.

This approach, like all approaches, has its merits and drawbacks. On the one hand, the odds of misspelling one's supervisor's name increase from negligible to "I don't have time to check but I'm sure it's fine"; of course, this can be mitigated by omitting one's supervisor from the acknowledgements altogether. On the other hand, the feeling of time marching inexorably towards the deadline, like sand slipping through one's fingers, gives an unrivalled clarity of mind. Like the memories of one's life flash before one's eyes in the moments preceding death, so do the memories of one's DPhil in the moments preceding submission. It is in these very moments that one most appreciates all the people who have made this possible, and who make the final stretch worth enduring. With the intent of turning this authentic feeling of appreciation into equally genuine words, it is time to write my own uninspired and painfully generic paragraphs.

Firstly, I will take my chances at the spelling game and thank my supervisor, Marc Lackenby, for his invaluable guidance from the very first day of my DPhil. I immediately learnt to fully trust his broad knowledge, his quick thinking, and his ability to retrieve the most obscure three-decade-old article from his memory and explain the core concepts with astounding clarity. I relied on him to direct my research, measure my progress, find problems in my solutions and solutions to my problems. Academia is a very intimidating place, but Marc opened the door for me, held my hand, and supported me while I took my first tentative steps. I am scared to let go now, but I hope he will be proud of the path I will take on my own.

I owe a significant portion of my knowledge to the senior researchers who have been willing to share their insights and expertise with me; I would like to single out Mehdi

Yazdi, Ric Wade, and Saul Schleimer, with whom I have had enriching conversations. I would also like to thank the first two for agreeing to read this lengthy manuscript in detail, and to apologise to the third for stranding him in Madrid.

I am grateful to the EPSRC for funding my DPhil, and asking nothing in return but a 2000-character summary of my research (spaces included). To St Hilda's College, for taking care of me when I first moved to Oxford, and for forcing social interactions that would later turn into long-term friendships. And to St Catherine's College, for having some of the most amazing first-year students Oxford has to offer: Cal, Laura, Sam, Srijan, Ashil, Ethan, Maya, and Oliver, I was lucky to have the opportunity of teaching you.

There are quite a few people who have endured my presence more than any reasonable person should want to. Given that the only reliable thing about research in mathematics is how often it does not work out, it has been immensely helpful to have someone to reliably turn to. I am particularly grateful to (little) Marc for being an amazing flatmate, a 10/10 friend, and an overall incredible person; and to Shan, for completing Marc in such a perfect way. To Ran, for her unmatched wit and inspiring spontaneity. To Mick, for his genuine kindness and (most importantly) his unending catalogue of board games. And to Tommi and Isaac, for demonstrating that friendship does not have to be grounded in shared research interests. To Adele, for imparting her wisdom to me. And to Misha, for accepting my wisdom, and occasionally forcing it out of me despite my feeble attempts at hiding behind a closed door. To Ismael, for being an OK flatmate. And to Pawel, Ognjen, Joe, Sam (but which one?), Soheil, Colin, Cameron, Dario (a fellow Italian), Davide (another fellow Italian), Giovanni (a fellow Italiano), and the whole Oxford Topology research group, for providing the most welcoming and stimulating environment a topologist could wish for. I shall not forget my non-Oxford friends, especially Yongsheng and Panos (for stalking me at every single conference I attended), Sayantika and Meenakshy (for sharing my love for pictures, while being infinitely better than me at producing them), Anna (who temporarily became an Oxford friend), and Andrea and Davide (with their *fontane mortali*).

Finally, I wish to extend my gratitude to my family. I decided to mention them last, not because their support has been less significant, but simply because I know that only they would be willing to endure this entire monologue before seeing their names mentioned. Mum, Dad, and Riccardo, thank you for providing a warm, welcoming, and safe haven; no matter how far I travel, I will always feel at home with you. *Nonna* Luisa, *nonno* Emilio, thank you for being my number one fans; you are more proud of me than I could ever be myself, and knowing this is one of my greatest sources of motivation.

Miriam, you essentially wrote your entire thesis about me, and finding out was one of the most emotional moments of my life; I cannot say to have done quite the same, but at least you made it into the acknowledgements. *Nonna* Giuliana, thank you for lighting the spark of my mathematical education, of which this thesis is, in some sense, the culmination; I do not know if you were thinking of fibred 3-manifolds while teaching me multiplication tables, but hopefully you are not disappointed. *Nonno* Gianni, you did not get to see this, but I hope you would be proud; not quite mechanical engineering, but close enough, right?

Jessica, I do not know whether I am more grateful to you for guiding me through my DPhil, or to my DPhil for guiding me towards meeting you. The number of times I would have completely broken down without your kindness and support far exceeds every other number that appears in this thesis (if you were to read it, you would find some really large ones). You have helped me complete this chapter of my life, and now I get the privilege of starting a new one with you. Thanks to you, I can conclude this section with a pun: my proposal must have been in NP, because your answer was "yes", and there is a certificate of this.

Abstract

In this thesis, we study the algorithmic problem of deciding whether a 3-manifold fibres over the circle, and if so, whether it is hyperbolic. We are not only concerned with decidability of these questions, but also with their computational complexity: a 3-manifold is described by a triangulation, and the efficiency of an algorithm is measured against the number of tetrahedra in the triangulation.

We prove that the problem of deciding whether a triangulated orientable 3-manifold fibres over the circle lies in the complexity class NP, generalising a result of Schleimer that only applied to atoroidal 3-manifolds. By design, our certificate for fibredness can be used to recover the monodromy of a fibration. Building on our previous work on algorithmic Nielsen-Thurston classification of surface mapping classes, we can decide whether the monodromy is pseudo-Anosov, and thus whether the fibred 3-manifold is hyperbolic. More precisely, we show that hyperbolicity of a triangulated orientable fibred 3-manifold can be certified in polynomial time in the number of tetrahedra in the triangulation and the Euler characteristic of a fibre. In the special case of knots in the 3-sphere, where the input is given as a planar diagram, our result implies that the problem of deciding whether a fibred knot is hyperbolic lies in NP.

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.

Contents

Chapter 1

Introduction

Can we decide if a given 3-manifold is hyperbolic? Perelman's proof of Thurston's geometrisation conjecture (see [31, 32, 33]) established that every closed orientable 3-manifold admits a canonical decomposition into pieces that fall into eight geometrically defined categories. It is then only natural that the effort to understand and classify 3-manifolds should be preceded by an effort to understand and classify these geometric pieces. The most basic question we can ask is whether we can recognise the geometric categories, that is, tell them apart from other 3-manifolds and between each other. To be precise, we are asking for a general, systematic procedure; in other words, an *algorithm*.

Of the eight geometries, hyperbolic geometry is undoubtedly the richest, most varied, and least amenable to a systematic classification. For these reasons, it is also the geometry we will focus on in this thesis. The general goal we will be working towards is providing an algorithm to decide whether a given orientable 3-manifold is hyperbolic. In an interesting turn of events, the geometrisation theorem, in addition to justifying our interest in the algorithmic recognition of hyperbolic 3-manifolds, also provides the theoretical foundation required to solve this problem. In [25], Manning gives an algorithm to decide whether a closed orientable 3-manifold is hyperbolic, assuming that its fundamental group has solvable word problem; that this hypothesis is always satisfied is a consequence of the geometrisation theorem (see [25, Section 1.1]). Alternatively, a self-contained solution to the hyperbolicity detection problem, assuming the geometrisation theorem, is given by Kuperberg in [20, Theorem 5.4].

While the question of *whether* hyperbolicity of 3-manifolds is decidable is settled, *how efficiently* it can be decided remains unanswered. For our purposes, a 3-manifold is described by a triangulation, and the efficiency of an algorithm is measured against the *size* of the triangulation, that is, the number of tetrahedra in it. Heuristically, the correct computational complexity class to aim for is NP. Informally speaking, a decision problem is in NP if, whenever the answer to the problem is "yes", there is a certificate of this fact that can be verified in polynomial time in the size of the input; for a formal discussion (and an introduction to computational complexity), see [2, Chapter 2]. Many interesting decision problems on 3-manifolds are known to be in NP:

- recognising the 3-sphere, by [16, 35];
- recognising other "simple" 3-manifolds, such as the real projective space, the solid torus, and handlebodies, by [16];
- recognising elliptic 3-manifolds, by [24];
- recognising Seifert-fibred spaces with boundary, by [17].

In particular, we see that elliptic geometry can be certified in polynomial time. In the case of non-empty boundary, the same is true for the six geometries represented by Seifert-fibred spaces (see [26, Theorem 12.1.1]).

It is natural to ask whether hyperbolic geometry can be included in the list above. Unfortunately, hyperbolicity seems to be much harder to certify than the other geometries. In [11], Haraway and Hoffman show that, if an irreducible 3-manifold with boundary is not hyperbolic, then this can be certified in polynomial time. However, there is no result in the literature that addresses efficient certification of hyperbolicity. Our main contribution in this thesis is a partial result in this direction, in the case of fibred 3-manifolds. In Theorem 5.18, we prove the following.

Theorem A. *Hyperbolicity of a triangulated orientable fibred* 3*-manifold can be certified in polynomial time in the size of the triangulation and the Euler characteristic of a fibre.*

It is important to note that this result does not imply that recognising fibred hyperbolic 3-manifolds is in NP. In fact, the size of our certificate depends polynomially on the size of the input – that is, the number of tetrahedra in the triangulation – but also on an additional parameter, the Euler characteristic of a fibre, which is not polynomially bounded by the size of the triangulation. There is a specific setting, however, in which this parameter does have a bound in terms of the input size. Consider a knot K in the 3-sphere, described by a diagram with n crossings. Thurston's geometrisation theorem (see [38]) specialises in this case to the following trichotomy: the knot K is either a torus knot, a satellite knot, or a hyperbolic knot. From the point of view of algorithmic recognition, the first two cases already have a satisfactory solution: the recognition problems for torus knots and satellite knots are both in NP, by [3] and [11] respectively; here, the size of the input is measured by the number n of crossings in the diagram. If we assume that K is fibred, then the Euler characteristic of a fibre is bounded below by $-n$. Therefore, as an immediate consequence of Theorem A, we obtain the following.

Corollary B. *The problem of deciding whether a fibred knot in the* 3*-sphere is hyperbolic is in* NP*.*

Our decision to focus on fibred 3-manifolds is mostly pragmatic. While hyperbolicity of arbitrary 3-manifolds is quite mysterious, for fibred 3-manifolds a very clear strategy emerges. By Thurston's geometrisation theorem (see [29]), a fibred 3-manifold is hyperbolic if and only if its monodromy is pseudo-Anosov. Therefore, certifying hyperbolicity of a fibred 3-manifold is, at least conceptually, quite simple: we need to provide a fibre, its monodromy, and a proof that the monodromy is pseudo-Anosov. These three steps correspond to Chapters 3, 4, and 5 of this thesis respectively, culminating with the proof of Theorem A.

In Chapter 3, we show that a normal least-weight fibre F of a triangulated orientable 3-manifold M admits a certificate of it being a fibre, that can be verified in polynomial time in the size of the triangulation and the logarithm of the weight of F. In particular, we obtain the following theorem (see Theorem 3.17), which may be of independent interest; we remark that, at this stage, there is no dependence on the complexity of the fibre.

Theorem C. *The problem of deciding whether a given triangulated orientable* 3*-manifold is fibred is in* NP*.*

This fact was already known for atoroidal 3-manifolds, by [35, Corollary 1.4]. Even though the rest of this thesis is only concerned with hyperbolic – and hence atoroidal – 3-manifolds, our proof Theorem C is motivated by more than the desire to fill a gap in the literature. In fact, our certificate is substantially different from that of [35, Corollary 1.4], and it is specifically designed to recover the monodromy of the fibration. We carry out this task in Chapter 4, by showing that the image of a normal curve in F under the monodromy can be certified in polynomial time; the relevant parameters here are the size of the triangulation of M, the logarithm of the weight of F, and the "size" of the curve (which is measured by the logarithm of its weight plus the number of triangles of

F it crosses). Precise statements of these results can be found in Propositions 4.11, 4.12, and 4.13.

Finally, in Chapter 5, we show how to certify the fact that the monodromy of F is pseudo-Anosov. This part relies heavily on our previous work in [4], which we could not include in this thesis due to space constraints; we remark, however, that the article has been published in a peer-reviewed journal (see the bibliography for details). In [4], we give an algorithm to decide whether a given mapping class φ of a surface F is pseudo-Anosov. The algorithm works by taking an essential normal curve a in F, computing the image $\varphi^n(a)$ of a under a sufficiently high power of φ, and then estimating the distance between a and $\varphi^n(a)$ in the curve graph of F. By work of Masur and Minsky (see [27]), and the quantitative refinements in [10] and [40], the mapping class φ is pseudo-Anosov if and only if this distance is sufficiently large. Adapting the algorithm of [4] to our setting is quite straightforward. Somewhat surprisingly, the hardest task is proving (as we do in Section 5.4) that the fibre F, when M is hyperbolic, contains an essential normal curve a of small weight. Once this is established, we can use the certificate from Chapter 4 to certify the image of a under a sufficiently high power of the monodromy, and then apply the algorithm of [4] to verify that the monodromy is indeed pseudo-Anosov.

Notation. Throughout this thesis, we use the following notation.

- If X is a finite set, then we denote by $|X|$ the number of elements of X.
- If X, instead, is a topological space, then by $|X|$ we mean the number of connected components of X.
- If Y is a subspace of a topological space X, then we denote by $\text{clos}(Y)$ the closure of Y in X; the ambient space X will always be clear from the context. If Y is a manifold, or at least a CW-complex, by $\text{int}(Y)$ we mean $Y \setminus \partial Y$ (which is not necessarily equal to the interior of Y in X).

Chapter 2

Triangulations of surfaces and 3-manifolds

2.1 Surfaces and normal curves

2.1.1 Triangulations of surfaces

Triangulations are one of the simplest way to describe compact surfaces in a combinatorial way that is amenable to algorithmic manipulation. In practical terms, a triangulation of a surface can be described by a list of vertices, a list of edges, and a list of triangles. Each edge is oriented, and its starting and terminal vertices (not necessarily distinct) are specified. The boundary of each triangle is specified by an ordered list of three (not necessarily distinct) edges, together with information about whether the edges are traversed consistently or inconsistently with their orientation.

Naturally, there are some conditions that must be satisfied in order for the topological space underlying this data to be a surface. For instance, the specified boundaries of the triangles must correspond to actual continuous paths in the 1-skeleton of the triangulation. Moreover, each edge must lie on the boundary of at least one and at most two triangles; in particular, edges that lie on the boundary of exactly one triangle define the boundary of the surface.

Our combinatorial description of a triangulation induces a natural orientation on each triangle. An oriented triangulation is a triangulation in which we enforce that no edge lies with the same orientation on the boundary of two triangles; in such a triangulation, the orientations of the triangles patch together consistently to define a global orientation on the surface.

A natural measure of complexity of a triangulation is its *size*, that is, the number of triangles it contains: if $\mathcal{T}$ is a triangulation of a compact surface F, we denote by $|\mathcal{T}|$ the number of triangles of $\mathcal{T}$. The number of edges and the number of vertices of $\mathcal{T}$ are then bounded above by $3|\mathcal{T}|$. In particular, we remark that the number of binary digits required to completely describe $\mathcal{T}$ is at most polynomial (in fact, *quasi-linear*) in $|\mathcal{T}|$.

For technical reasons, we will sometimes want to avoid situations in which a triangle has two or more edges lying on the boundary of the surface. Hence, we give the following definition.

Definition 2.1 (Flapless triangulation)**.** A triangulation of a compact surface F is *flapless* if each triangle has at most one edge lying on ∂F. ×

An easy counting argument shows that the number of edges in a flapless triangulation is at most twice the number of triangles.

2.1.2 Simplicial subsets of surfaces

We will often need to deal with subsets of a triangulated surface F that are simplicial – that is, unions of simplices – but are not necessarily submanifolds. A *sub*-1-*complex* of F is a union X of edges of F; if X is a 1-manifold, then we say it is a *simplicial sub*-1-*manifold* of F. Either way, we denote by $\ell(X)$ the number of edges in X.

Similarly, a *sub*-2-*complex* of F is a union X of triangles of F; if X is a surface, then we say it is a *simplicial subsurface* of F. Either way, we denote by $\mathrm{area}(X)$ the number of triangles in X. See Figure 2.1a for an example of a sub-2-complex of a triangulated surface. Even though X may have singular points that prevent it from being a subsurface – namely, when it contains two triangles that intersect in a vertex and not in one of the adjacent edges – there is an abstract triangulated surface $\mathrm{ab}(X)$, together with an immersion $\mathrm{emb}_X\colon \mathrm{ab}(X) \to F$, such that $\mathrm{emb}_X(\mathrm{ab}(X)) = X$ and emb_X is an embedding except at the vertices of $\mathrm{ab}(X)$. The surface $\mathrm{ab}(X)$ is obtained by taking the triangles of X and gluing them along shared edges, but not along isolated shared vertices. We will often blur the distinction between X and $\mathrm{ab}(X)$; for instance, when we talk about the components of X, we will actually mean the images under emb_X of the components of $\mathrm{ab}(X)$. It will always be clear from the context whether we are talking about X as a topological subspace of F or as an abstract triangulated surface.

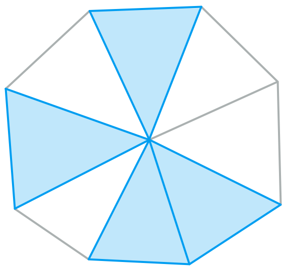

(a)

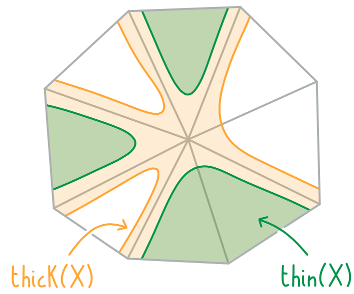

(b)

Figure 2.1. (a) A sub-2-complex X of a triangulated surface. (b) The subsurfaces thick(X) and thin(X).

Sometimes, for convenience, we will want to transform a sub-2-complex X of F into an actual subsurface; this can be done in two ways. The first way is to "thicken" X by adding a small regular neighbourhood of X in F; the result is a subsurface of F, that we denote by thick(X), as shown in Figure 2.1b. Note that thick(X) is defined up to isotopy and, as a topological space, it is homotopy equivalent to X. We will always pick thick(X) in its isotopy class so that it does not intersect any other relevant objects in F that are disjoint from X. The second way to obtain a subsurface from X is to "carve out" a regular neighbourhood of the boundary of X from X. More precisely, we let

$$\mathrm{thin}(X) = \mathrm{clos}(X \setminus \mathrm{thick}(\mathrm{clos}(F \setminus X))),$$

as shown in Figure 2.1b. This is a subsurface of X, and it is homeomorphic to ab(X). Like thick(X), the subsurface thin(X) is defined up to isotopy; whenever possible, we will pick thin(X) so that it contains every relevant object in F that is contained in X.

2.1.3 1-manifolds in surfaces

Let F be a compact surface. By *collared annulus* in F we mean a subsurface C of F that is homeomorphic to an annulus, with one component of ∂C contained in ∂F and the other disjoint from ∂F. A curve properly embedded in F is *essential* if it does not bound a disc or a collared annulus in F. A multicurve is *essential* if all its components are.

Fix now a triangulation $\mathcal{T}$ of F. We say that a 1-manifold a embedded in F is *in general position* if it is disjoint from the vertices of $\mathcal{T}$ and transverse to the edges of $\mathcal{T}$.

The *weight* of a is the number

$$w(a) = |a \cap \mathcal{T}^{(1)}|$$

of points in the intersection of a with the 1-skeleton of $\mathcal{T}$.

Finally, we define the *support* of a to be the union of all triangles of $\mathcal{T}$ whose interiors intersect a, and we denote it by $\operatorname{supp}(a)$. Note that the support of a is a sub-2-complex of F.

2.1.4 Normal 1-manifolds

A *normal arc* is an arc properly embedded in a triangle T of $\mathcal{T}$ whose endpoints lie on two distinct edges of T; there are three types of normal arcs in T, up to isotopy fixing the vertices of T. A general position 1-manifold a properly embedded in F is *normal* if it intersects each triangle of $\mathcal{T}$ in a union of normal arcs; the typical intersection of a normal 1-manifold with a triangle is shown in Figure 2.2a. Every "interesting" 1-manifold is isotopic to a normal 1-manifold; this claim will be clarified and justified in Proposition 2.3. On the other hand, normal 1-manifolds are intrinsically combinatorial objects, which makes them particularly well-suited for algorithmic manipulation.

There are two essentially equivalent ways of representing a normal 1-manifold a. The first way is to encode a as a vector of non-negative integers, so that each coordinate describes the number of intersections of a with a given edge of $\mathcal{T}$. Alternatively, we can represent a as a vector of non-negative integers, so that each coordinate describes the number of normal arcs of a given type in the intersection of a with a given triangle of $\mathcal{T}$. The first representation might seem more natural, but for consistency with the 3-dimensional setting we will adopt the second representation. More precisely, denote by t the number of triangles in $\mathcal{T}$. After fixing an order of the $3t$ types of normal arcs of $\mathcal{T}$ (we can assume that this order is part of the data that describes a triangulation), to every normal 1-manifold a we can associate a vector $\mathbf{v}_a \in \mathbb{Z}^{3t}_{\geqslant 0}$, where the i-th coordinate of $\mathbf{v}_a$ is the number of normal arcs of the i-th type in the intersection of a with the triangles of $\mathcal{T}$.

It is not the case that every vector $\mathbf{w} \in \mathbb{Z}^{3t}_{\geqslant 0}$ represents a normal 1-manifold. A necessary and sufficient condition for $\mathbf{w}$ to have a normal 1-manifold a such that $\mathbf{v}_a = \mathbf{w}$ is given by the *matching equations*. Fix and edge e of $\mathcal{T}$ that does not lie on ∂F. Let i and j be the types of normal arcs that intersect e on one side, and let k and l be the types of normal arcs that intersect e on the other side. Then the matching equation for

(a) (b)

Figure 2.2. (a) The intersection between a normal 1-manifold and a triangle of the triangulation. (b) The matching equation for the edge e reads $\mathbf{w}_i + \mathbf{w}_j = \mathbf{w}_k + \mathbf{w}_l$.

---

e reads

$$\mathbf{w}_i + \mathbf{w}_j = \mathbf{w}_k + \mathbf{w}_l,$$

and it is depicted in Figure 2.2b. The vector $\mathbf{w}$ represents a normal 1-manifold if and only if the matching equations are satisfied for every edge of $\mathcal{T}$ that does not lie on ∂F.

A vector $\mathbf{w}$ satisfying the matching equations represents exactly one normal 1-manifold up to *normal isotopy*, that is, up to isotopy that fixes the vertices of $\mathcal{T}$ and preserves its 1-skeleton. It will be implicitly understood that, when we algorithmically manipulate normal 1-manifolds – in particular, when we take them as input or return them as output – we are actually manipulating their associated vectors. The underlying ambiguity will never be an issue, since all our algorithmic statements will be invariant under normal isotopy. Finally, we remark that this representation allows us to encode a normal 1-manifold a with a list of $3t$ integers, each of which has at most $\log(2w(a) + 1)$ binary digits; here and in the following, the logarithm is taken to base 2.

2.1.5 Intersection number

The *intersection number* of two multicurves a and b in a compact surface is the integer

$$i(a, b) = \min \{|a' \cap b'| : a' \text{ is isotopic to } a,\ b' \text{ is isotopic to } b\}.$$

The following elementary result gives a bound on the intersection number of normal multicurves in terms of their weights.

Proposition 2.2 (Bounding intersection number with weight)**.** *Let $\mathcal{T}$ be a triangulation*

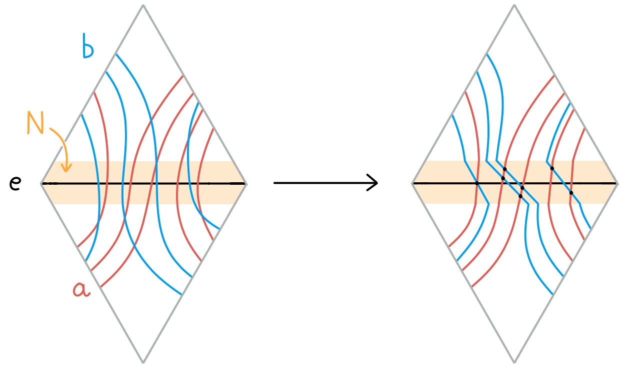

Figure 2.3. The multicurves a and b can be normally isotoped to be disjoint outside N, and intersect at most $|a \cap e| \cdot |b \cap e|$ near the edge e.

of a compact surface F, and let a and b be normal multicurves in F. Then

$$i(a, b) \leqslant w(a) \cdot w(b).$$

Proof. Let N be a small regular neighbourhood of the 1-skeleton of $\mathcal{T}$ in F. Up to normal isotopy, we can assume that b is disjoint from a outside N; this can be guaranteed by isotoping b so that its normal arcs run parallel to those of a in each triangle of $\mathcal{T}$. Importantly, note that a normal isotopy does not change the weight of b. Near each edge e of $\mathcal{T}$, the intersections $a \cap N$ and $b \cap N$ can be isotoped to unions of "straight arcs", so that $a \cap N$ intersects $b \cap N$ at most $|a \cap e| \cdot |b \cap e|$ times near e. The normal isotopies of a and b that realise this bound are depicted in Figure 2.3. Finally, we conclude that

$$i(a, b) \leqslant \sum_e |a \cap e| \cdot |b \cap e| \leqslant w(a) \cdot w(b),$$

where the sum is taken over all edges e of $\mathcal{T}$. ▢

2.1.6 Normalising 1-manifolds

We claimed previously that every "interesting" 1-manifold is isotopic to a normal 1-manifold. We now make this statement precise, and we provide a proof of this well-known fact.

Proposition 2.3 (Normalising 1-manifolds)**.** *Let F be a triangulated compact surface, and let a be a general position 1-manifold properly embedded in F. Suppose that no* *component of a is a curve bounding a disc in F or an arc cobounding a disc with a subarc of an edge of ∂F. Then a is isotopic fixing ∂F to a normal 1-manifold a'. Moreover, the normal 1-manifold a' can be chosen so that, for each edge e of F, the inequality*

$$|a' \cap e| \leqslant |a \cap e|$$

holds; in particular, this implies that $w(a') \leqslant w(a)$.

Proof. This is a 2-dimensional analogue of Proposition 2.7 below, which is carefully proved in [28, Theorem 3.3.21]. This setting is much simpler than the 3-dimensional one, but the statement proved in [28, Theorem 3.2.2] for curves in surfaces is not general enough for our purposes. Therefore, we provide a proof of this well-known result for the sake of completeness.

There exists a general position 1-manifold a' properly embedded in F that is isotopic to a fixing ∂F and satisfies the following properties:

- $|a' \cap e| \leqslant |a \cap e|$ for every edge e of F;
- there is no general position 1-manifold b properly embedded in F that is isotopic to a' fixing ∂F and satisfies $|b \cap e| \leqslant |a' \cap e|$ for every edge e of F, with at least one strict inequality.

There are two ways in which a' can fail to be normal. The first way is if some component of a' is a curve fully contained in a triangle T of F. If this were the case, then this component would bound a disc in T, which contradicts the hypothesis. The second way in which a' can fail to be normal is if there is a triangle T of F such that a component of $a' \cap T$ is an arc with both endpoints on the same edge e of T. Let c be an innermost such component of $a' \cap T$. In particular, this means that c cobounds a disc D with a subarc of e, such that $\mathrm{int}(D)$ is disjoint from a'. The edge e cannot lie on ∂F by hypothesis; therefore, we can isotope c across D to the other side of e. This produces a new general position 1-manifold b, isotopic to a' fixing ∂F, such that $|b \cap e| = |a' \cap e| - 2$, and $|b \cap e'| = |a' \cap e'|$ for every edge e' of F different from e; this contradicts the minimality of a'. Therefore, we conclude that a' is normal. ▢

2.1.7 Least-weight curves

A multicurve a in a compact surface F is *least-weight* if it has minimal weight amongst all general position multicurves isotopic to a. It is an immediate consequence of Proposition 2.3 that every general position multicurve in F is isotopic to a least-weight multicurve.

We conclude this section with an argument showing that being least-weight is a local property of curves. In fact, it is a consequence of Corollary 2.5 below that every subarc of a least-weight curve has minimal weight amongst all arcs isotopic to it fixing the endpoints.

Lemma 2.4 (Local isotopies of curves)**.** *Let F be a compact orientable surface, and let a be a curve in F. Let c be a non-empty subarc of a, and let c' be an arc that is isotopic to c in F fixing ∂c. Suppose that c' is transverse to $a \setminus c$, and let N be a neighbourhood of $c' \cap (a \setminus c)$ in F. Then there is a curve a' in F that is isotopic to a in F and coincides with $(a \setminus c) \cup c'$ outside N.*

Proof. Let $F' = F \setminus \partial c$, and let $a_0 = a \setminus c$, $c_0 = c' \setminus \partial c'$. Note that a_0 and c_0 are arcs in the punctured surface F', and by assumption they can be isotoped off each other in F'. For $i \geqslant 0$, we iterate the following procedure. If a_i and c_i are disjoint, then we set $n = i$ and terminate the procedure. Otherwise, they must form a bigon or a half-bigon. Let D be an innermost bigon or half-bigon; more precisely, it is a disc embedded in F such that ∂D is the union of two arcs r and s, with $r \subseteq a_i \cup \partial c$ and $s \subseteq c_i \cup \partial c$, and such that the interior of D is disjoint from a_i and c_i. Figure 2.4 shows how to use D to isotope a_i and c_i to arcs a_{i+1} and c_{i+1} satisfying the following properties:

- a_{i+1} and c_{i+1} intersect transversely, if at all;
- $a_{i+1} \cup c_{i+1}$ coincides with $a_i \cup c_i$ outside N;
- $|a_{i+1} \cap c_{i+1}| < |a_i \cap c_i|$.

This procedure must eventually terminate, yielding two disjoint arcs a_n and c_n in F'. It is then easy to see that the curve $a' = a_n \cup c_n \cup \partial c$ satisfies the required properties. ▢

Corollary 2.5 (Least-weight curves are locally least-weight)**.** *Let $\mathcal{T}$ be a triangulation of a compact orientable surface F, and let a be a general position curve in F. Let c be a non-empty subarc of a, and let c' be a general position arc that is isotopic to c in F fixing ∂c. Then there is a general position curve a' in F that is isotopic to a in F, such that*

$$|a' \cap e| = |(a \setminus c) \cap e| + |c' \cap e|$$

for each edge e of $\mathcal{T}$. In particular, we have that

$$w(a') = w(a) + w(c') - w(c).$$

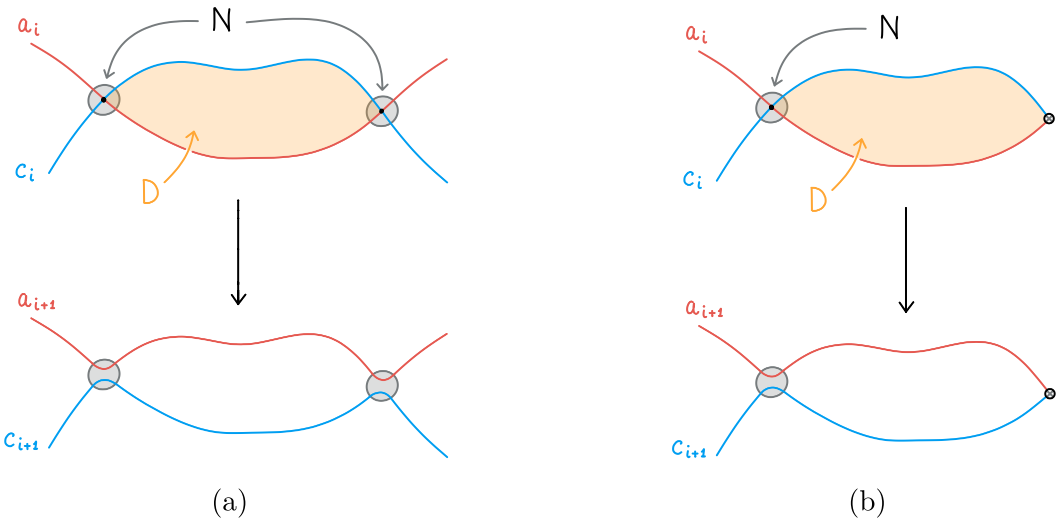

Figure 2.4. (a) How to construct the arcs a_{i+1} and c_{i+1} from a_i and c_i in the case where D is a bigon. (b) How to construct the arcs a_{i+1} and c_{i+1} from a_i and c_i in the case where D is a half-bigon.

Proof. Up to a small perturbation of c', that does not change the number of intersections between c' and edges of $\mathcal{T}$, we can assume that c' is transverse to $a \setminus c$, and moreover that $c' \cap (a \setminus c)$ is disjoint from $\mathcal{T}^{(1)}$. Therefore, there is a neighbourhood N of $c' \cap (a \setminus c)$ in F that is disjoint from $\mathcal{T}^{(1)}$. Applying Lemma 2.4 yields a curve a' that is isotopic to a and coincides with $(a \setminus c) \cup c'$ outside N; in particular, we see that

$$|a' \cap e| = |(a \setminus c) \cap e| + |c' \cap e|$$

for each edge e of $\mathcal{T}$. The second conclusion in the statement follows by summing this equality over all edges e of $\mathcal{T}$. ▢

2.2 3-manifolds and normal surfaces

2.2.1 Fibred 3-manifolds

The main results of this thesis concern fibred 3-manifolds. A compact 3-manifold M is *fibred* if it fibres over the circle, that is, if it admits a fibration $p\colon M \to S^1$ with fibre a compact surface. A *fibre* of M is a surface F of the form $p^{-1}(x)$ for some point $x \in S^1$. Let us consider the case where M is orientable, and F is *transversely oriented*

– that is, endowed with a transverse orientation in M. This transverse orientation induces an orientation on S^1, which we can use to define a map $f\colon [0,1] \to S^1$ such that $f(0) = f(1) = x$, and $f(t)$ runs around S^1 once in the positive direction as t ranges from 0 to 1. The map f lifts to an isotopy $(g_t)_{t\in[0,1]}$, where each g_t is an embedding of F in M such that $g_t(F) = p^{-1}(f(t))$; we also ask that $g_0\colon F \to M$ is simply the inclusion of F in M. Since $g_0(F) = g_1(F)$, we can define a homeomorphism $\varphi\colon F \to F$ by setting $\varphi = g_0^{-1} \circ g_1$. The homeomorphism φ is called the *monodromy* of F, and it is orientation-preserving. Since it is only defined up to isotopy, it is usually thought of as an element of the mapping class group $\mathrm{Mcg}(F)$ of F; note that flipping the transverse orientation of F amounts to replacing φ with φ^{-1}. The 3-manifold M can be reconstructed from the monodromy φ via the following procedure: take the interval bundle $F \times [0,1]$, and glue $F \times \{1\}$ to $F \times \{0\}$ via the homeomorphism φ. In other words, the 3 manifold M is homeomorphic to the mapping torus

$$F \times [0,1] \big/ \{(y,1) \sim (\varphi(y),0) : y \in F\}.$$

2.2.2 3-manifolds and embedded surfaces

The theory of 3-manifolds with boundary patterns that we use is consistent with that of Matveev's book [28]. We will restate the main definitions here, but we refer the reader to [28, §3.3] for more details.

A 3-*manifold with boundary pattern* is a pair (M,Γ), where M is a compact 3-manifold and Γ is a graph embedded in ∂M with no isolated vertices. If the boundary pattern is not specified, we will assume that it is the empty graph. An isotopy of M is *admissible* if it preserves Γ (not necessarily pointwise). A subset of M is *clean* if it does not intersect Γ.

The 3-manifold with boundary pattern (M,Γ) is *irreducible* if every sphere embedded in M bounds a 3-ball in M. It is *boundary-irreducible* if every clean disc properly embedded in M cobounds a clean 3-ball with ∂M.

We say that a surface F is *properly embedded* in (M,Γ) if F is properly embedded in M and ∂F intersects Γ transversely. A *compressing disc* for F is a disc D embedded in M such that $D \cap F = \partial D$; we say that D is *trivial* if ∂D bounds a disc in F. The surface F is *incompressible* in M if it has no non-trivial compressing discs; note that this notion does not depend on Γ. It is a fact (see [28, Lemma 3.3.5]) that every connected orientable surface properly embedded in a compact orientable 3-manifold is incompressible if and only if it is π_1-injective.

A *boundary-compressing disc* for F is a clean disc D embedded in M such that ∂D is the union of two arcs a and b with $\partial a = \partial b$, where $a = D \cap F$ and $b = D \cap \partial M$; we say that D is *trivial* if a cobounds a clean disc with ∂F in F. The surface F is *boundary-incompressible* in (M, Γ) if it has no non-trivial clean boundary-compressing discs.

2.2.3 Triangulations of 3-manifolds

Like in the case of surfaces, triangulations offer a simple way to describe compact 3-manifolds combinatorially. Without going into as much detail as we did for surfaces, we can describe a triangulation of a compact 3-manifold by a list of vertices, a list of edges, a list of triangles, and a list of tetrahedra; each i-simplex is equipped with information about its attaching map to the $(i-1)$-skeleton, for $1 \leqslant i \leqslant 3$. We remark that, with this combinatorial description, each tetrahedron is endowed with an orientation (up to arbitrarily and universally fixing an orientation on the "standard" tetrahedron). We say that the triangulation is oriented if the orientations of the tetrahedra patch together consistently to define a global orientation on the 3-manifold.

A natural measure of complexity of a triangulation $\mathcal{T}$ is the number $|\mathcal{T}|$ of tetrahedra it contains, which we also call the *size* of the triangulation. The complexity of an algorithm that takes as input a triangulated 3-manifold will always be measured against this quantity. In fact, the number of triangles, edges, and vertices in a triangulation $\mathcal{T}$ is bounded above by a linear function of $|\mathcal{T}|$. Therefore, the number of binary digits required to completely describe $\mathcal{T}$ is at most polynomial (in fact, *quasi-linear*) in $|\mathcal{T}|$.

2.2.4 Sub-3-complexes

Similarly to the 2-dimensional setting, we will sometimes need to deal with subsets of a 3-manifold that are submanifolds away from a lower-dimensional singular locus. Let M be a compact orientable 3-manifold with a polyhedral cell structure – meaning that M is obtained by gluing a collection of polyhedral pieces along their faces. This polyhedral cell structure might be a triangulation, or it might arise from cutting a triangulated 3-manifold along a normal surface, as we will see in Section 2.3.3. A *sub-3-complex* of M is a union X of cells of M such that the *singular locus* of X – that is, the set of points of $x \in X$ such that no open neighbourhood of x in M intersects X in an open 3-ball or half-ball – is a compact 1-manifold properly embedded in M. Like in the 2-dimensional case, the sub-3-complex X can be thought of as an abstract compact 3-manifold $\mathrm{ab}(X)$, together with an immersion $\mathrm{emb}_X\colon \mathrm{ab}(X) \to M$ such that $\mathrm{emb}_X(\mathrm{ab}(X)) = X$. The 3-manifold $\mathrm{ab}(X)$ can be obtained by taking the cells of M contained in X and gluing them along their common faces (but not along common isolated edges). The immersion emb_X is an embedding, except at the points of $\mathrm{ab}(X)$ that get mapped to the singular locus of X. Like in the 2-dimensional case, we will often blur the distinction between X and $\mathrm{ab}(X)$; in particular, by "components of X", we will mean the components of $\mathrm{ab}(X)$, or possibly their images under emb_X.

Similarly to what we did for sub-2-complexes of surfaces, we can define $\mathrm{thick}(X)$ and $\mathrm{thin}(X)$ for a sub-3-complex X of M. The definition is the same: we let $\mathrm{thick}(X)$ be the union of X with a small regular neighbourhood of ∂X in M, and we let

$$\mathrm{thin}(X) = \mathrm{clos}(M \setminus (\mathrm{thick}(\mathrm{clos}(M \setminus X))))$$

(note that $\mathrm{clos}(M \setminus X)$ is also a sub-3-complex of M).

2.2.5 Normal surfaces

Let $\mathcal{T}$ be a triangulation of a compact orientable 3-manifold M; we admit the possibility that M has a boundary pattern Γ, but this will not be relevant for the moment (note, however, that [28] assumes that boundary patterns are simplicial, and we will do the same). We say that a surface F embedded in M is *in general position* if it is disjoint from the vertices of $\mathcal{T}$ and transverse to the edges and triangles of $\mathcal{T}$. For such a surface F, we define the *weight* of F to be the number

$$w(F) = |F \cap \mathcal{T}^{(1)}|$$

of points in the intersection of F with the 1-skeleton of $\mathcal{T}$.

There is a theory of normal surfaces in triangulated 3-manifolds that is analogous to the theory of normal curves in triangulated surfaces, and is developed in great detail in [28, Section 3.3]. We again refer the reader to this source, and we only offer a summary of the main definitions and results. A *normal disc* is a disc properly embedded in a tetrahedron T of $\mathcal{T}$ such that its boundary crosses at least one edge of T, and no edge of T more than once. It is not hard to see that, up to *normal isotopy* – that is, up to isotopy that fixes the vertices and preserves the 1-skeleton and the 2-skeleton – there are exactly seven types of normal discs in a tetrahedron: four "triangles" and three "quadrilaterals". A general position surface F properly embedded in M is *normal* if it intersects each

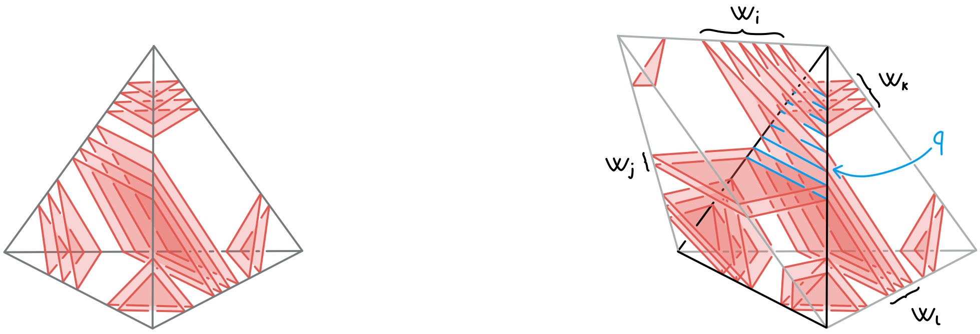

(a) (b)

Figure 2.5. (a) The intersection between a normal surface and a tetrahedron of the triangulation. (b) The matching equation for the normal arc of type q reads $\mathbf{w}_i + \mathbf{w}_j = \mathbf{w}_k + \mathbf{w}_l$.

tetrahedron of $\mathcal{T}$ in a union of normal discs; the typical intersection of a normal surface with a tetrahedron is shown in Figure 2.5a.

Similarly to normal curves in surfaces, a normal surface F in M is uniquely identified up to normal isotopy by a vector $\mathbf{v}_F \in \mathbb{Z}^{7t}_{\geqslant 0}$, where t is the number of tetrahedra in $\mathcal{T}$. In this setting, there is one matching equation for each type of normal arc in each triangle of $\mathcal{T}$ that does not lie on ∂M. More precisely, the matching equation for a given type q of normal arc in a triangle R of $\mathcal{T}$ not contained in ∂M reads

$$\mathbf{w}_i + \mathbf{w}_j = \mathbf{w}_k + \mathbf{w}_l,$$

where i, j, k, and l are the four types of normal discs (two triangles and two quadrilaterals) that intersect R in normal arcs of type q, with i and j lying in one of the two tetrahedra adjacent to R, and k and l lying in the other; a graphical representation of the matching equation for q is shown in Figure 2.5b.

There are other constraints, called *consistency equations*, that a vector $\mathbf{w}$ must satisfy in order to represent a normal surface. These equations enforce that each tetrahedron must contain at most one type of normal quadrilateral; in fact, it is easy to see that two normal quadrilateral of different types in the same tetrahedron must intersect. More precisely, the consistency equation for a tetrahedron T of $\mathcal{T}$ reads

$$\text{at most one of } \mathbf{w}_i,\ \mathbf{w}_j, \text{ and } \mathbf{w}_k \text{ is non-zero,}$$

where i, j, and k are the three types of normal quadrilaterals in T. One can show that there is a one-to one correspondence between normal isotopy classes of normal surfaces in M and vectors in $\mathbb{Z}^{7t}_{\geqslant 0}$ satisfying the matching and consistency equations.

Given two normal surfaces G and H in M, suppose that $\mathbf{v}_G + \mathbf{v}_H$ satisfies the consistency equations. We can then define the *normal sum* $G + H$ of G and H as the normal surface with normal vector $\mathbf{v}_G + \mathbf{v}_H$. A normal surface is called *fundamental* if it cannot be written as the normal sum of two non-empty normal surfaces. A key property of fundamental surfaces is that their weights are uniformly bounded by a constant that depends only on the size of the triangulation $\mathcal{T}$.

Proposition 2.6 (Weight of fundamental surfaces)**.** *Let $\mathcal{T}$ be a triangulation of a compact 3-manifold M with t tetrahedra, and let F be a fundamental normal surface in M. Then*

$$w(F) \leqslant t^2 \cdot 2^{7t+5}.$$

Proof. This result is an immediate consequence of [13, Lemma 6.1], which states that each entry of $\mathbf{v}_F$ is bounded above by $t \cdot 2^{7t}$. Since the vector $\mathbf{v}_F$ has $7t$ entries, and each normal disc contributes to at most 4 points in $F \cap \mathcal{T}^{(1)}$, the desired inequality follows. ▢

From now on, the boundary pattern Γ will be relevant. The following result shows how "interesting" surfaces in irreducible boundary-irreducible 3-manifolds with boundary pattern can be admissibly isotoped to be normal. Following [28, Definition 3.3.19], we say that a disc D properly embedded in (M, Γ) is an *inessential semi-clean disc* if it is isotopic fixing ∂D to a disc in ∂M whose intersection with Γ is an arc.

Proposition 2.7 (Normalising surfaces)**.** *Let F be an incompressible boundary-incompressible general position surface properly embedded in a triangulated compact irreducible boundary-irreducible 3-manifold with boundary pattern (M, Γ). Suppose that no component of F is a sphere, a clean disc, or an inessential semi-clean disc. Then F is isotopic fixing Γ to a normal surface F'. Moreover, the surface F' can be chosen so that, for each edge e of M, the inequality*

$$|F' \cap e| \leqslant |F \cap e|$$

holds; in particular, this implies that $w(F') \leqslant w(F)$.

Proof. This is essentially conclusion 4 of [28, Proposition 3.3.24], although that statement is less general than what we are claiming. One can, however, upgrade the proof therein, by noting that the eight normalisation moves described in [28, §3.3.3], when applicable, preserve the property of a surface to contain a subsurface that is isotopic to F fixing Γ; note that irreducibility and boundary-irreducibility of (M, Γ), together with incompressibility and boundary-incompressibility of F, are crucial to this claim. If we let F' be such a subsurface of the normal surface obtained through the normalisation procedure of [28, Theorem 3.3.21], then the fact that $|F' \cap e| \leqslant |F \cap e|$ for each edge e of M is observed in [28, Remark 3.3.22]. ▢

We say that a general position surface F properly embedded in (M, Γ) is *least-weight* if it has minimal weight amongst all general position surfaces admissibly isotopic to it. As an immediate consequence of Proposition 2.7, we see that every incompressible boundary-incompressible general position surface properly embedded in a triangulated compact irreducible boundary-irreducible 3-manifold with boundary pattern is admissibly isotopic to a least-weight normal surface, provided it does not have any component that is a sphere, a clean disc, or an inessential semi-clean disc.

2.2.6 Triangulations and submanifolds of normal surfaces

Let $\mathcal{T}$ be a triangulation of a compact orientable 3-manifold M, and let F be a normal surface in M. The normal surface F inherits a triangulation from $\mathcal{T}$, where each normal disc of F is a sub-2-complex. More precisely, normal triangles are triangulated with 3 triangles, and normal quadrilaterals are triangulated with 4 triangles, as shown in Figure 2.6. We say that two triangles in this triangulation of F are of the same *type* if they lie in the same tetrahedron T of M, they lie on normal discs of the same type, and they intersect the same face of ∂T (here, we think of T as an *abstract tetrahedron*, that is, ignoring the face identifications). Therefore, there are 3 types of triangles for each type of normal triangle, and 4 types of triangles for each type of normal quadrilateral. From now on, we will always implicitly assume that normal surfaces are endowed with this triangulation.

Remark 2.8 (The triangulation of a normal surface is flapless). Even though the triangulation of a normal surface we have defined might seem "suboptimal", in the sense that it uses more triangles than necessary, it has the convenient property of being flapless. ×

We now describe how to compactly encode normal 1-manifolds in F and sub-2-complexes of F, with respect to the triangulation of F inherited from $\mathcal{T}$. Let a be a normal 1-manifold in F. Then a can be fully described (up to normal isotopy) by listing the triangles in $\mathrm{supp}(a)$, and then encoding a as a normal 1-manifold in $\mathrm{ab}(\mathrm{supp}(a))$. We remark that a triangle of F can be pinpointed with a number of binary digits that is

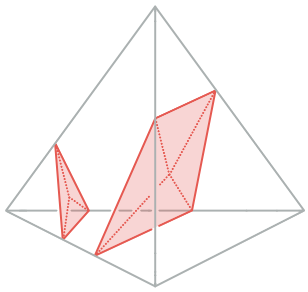

Figure 2.6. Normal triangles of a normal surface are triangulated with 3 triangles, and normal quadrilaterals with 4 triangles.

linear in $(\log t + 1) \cdot (\log w(F) + 1)$, where t is the number of tetrahedra in $\mathcal{T}$. Therefore, this encoding of a uses a number of binary digits that is linear in

$$(\log t + 1) \cdot (\log w(F) + 1) \cdot \operatorname{area}(\operatorname{supp}(a)) \cdot \log(w(a) + 1).$$

As far as sub-2-complexes are concerned, they are almost determined by their boundary, and in fact by the portion of their boundary that is not contained in ∂F. Let us introduce the following notation: if F is any compact surface, and G is any subset of F, we define

$$\partial_F G = \operatorname{clos}(\partial G \setminus \partial F).$$

If F is triangulated and G is a sub-2-complex of F, then we can encode G as follows: we list the edges in $\partial_F G$, and we include an additional list of triangles of F, one for each component of G; note that this information is enough to fully reconstruct G. It is easy to see that this encoding uses a number of binary digits that is linear in

$$(\log t + 1) \cdot (\log w(F) + 1) \cdot (\ell(\partial_F G) + 1).$$

We will see in Section 2.4.3 how to extract topological information about G from this compressed encoding.

2.3 Guts and parallelity bundle

2.3.1 Pre-sutured manifolds and subcomplexes

Sutured manifolds are a well-established tool in the study of 3-manifolds, introduced by Gabai in [9]. In this thesis, we will need sutured manifolds to describe candidate interval bundle structures on 3-manifolds. In particular, some of the items in our certificates will be "3-manifolds with sutures" that, in valid certificates, will always be suture-preservingly homeomorphic to product interval bundles; however, if the certificate is invalid, this may not be the case, and in fact the sutures may not even be annuli or tori. For this reason, we introduce a very lax variant of sutured manifolds, admitting arbitrary sutures, to fit our needs.

Definition 2.9 (Pre-sutured manifold)**.** A *pre-sutured manifold* is a compact oriented 3-manifold M together with a decomposition of ∂M into three compact surfaces $\partial_0 M$, $\partial_1 M$, and $\partial_v M$; we require that $\partial_0 M$ and $\partial_1 M$ are disjoint, and that $\partial_v M$ and $\partial_0 M \cup \partial_1 M$ intersect precisely along their boundaries. The subsurface $\partial_v M \subseteq \partial M$ is called the *vertical boundary* of M. We also write $\partial_h M$ for $\partial_0 M \cup \partial_1 M$, and we call it the *horizontal boundary* of M. ×

The notation we use is borrowed from interval bundles, rather than from standard sutured manifold theory. For a compact connected surface F, there is a unique orientable interval bundle M over F. Its horizontal boundary $\partial_h M$ is a two-sheeted cover of F, while its vertical boundary $\partial_v M$ is a disjoint union of annuli. Whenever we say that a 3-manifold is an interval bundle over a surface, we will always implicitly endow it with this natural pre-sutured manifold structure. We will mostly be dealing with *product interval bundles*, that is, 3-manifolds of the form $M = F \times [0, 1]$, where F is a compact orientable surface. In this case, note that $\partial_v M = \partial F \times [0, 1]$ and $\partial_i M = F \times \{i\}$ for $i \in \{0, 1\}$. When M is a *twisted interval bundle* over a connected non-orientable surface, there is no natural way to split its horizontal boundary into $\partial_0 M$ and $\partial_1 M$; in this case we will allow both possible assignments – that is, $\partial_0 M = \partial_h M$ and $\partial_1 M = \varnothing$ or vice versa.

Two pre-sutured manifolds M and N are *homeomorphic as pre-sutured manifolds* if there is an orientation-preserving homeomorphism $M \to N$ sending $\partial_i M$ to $\partial_i N$ for $i \in \{0, 1\}$ – and, consequently, $\partial_v M$ to $\partial_v N$.

A *pre-sutured triangulation* is a triangulation of a pre-sutured manifold M, such that $\partial_h M$ and $\partial_v M$ are simplicial subsurfaces of M. Note that the additional data – that is, the surfaces $\partial_0 M$, $\partial_1 M$, and $\partial_v M$ – can be encoded with a number of binary digits that is linear in the size of the triangulation. We say that a pre-sutured triangulation of a compact 3-manifold M is *suitable* if M is homeomorphic as a pre-sutured manifold to $\partial_0 M \times [0,1]$.

These notions can be extended to sub-3-complexes of 3-manifolds with minimal changes. Let M be a compact oriented 3-manifold with a polyhedral cell structure, and let X be a sub-3-complex of M. A *pre-sutured subcomplex* structure on X is defined analogously to Definition 2.9, with the additional requirements that the singular locus of X should be properly embedded in $\partial_v X$. Note that a pre-sutured subcomplex structure on X induces a pre-sutured manifold structure on $\mathrm{ab}(X)$, by pull-back via emb_X. Conversely, a pre-sutured manifold structure on $\mathrm{ab}(X)$ induces a pre-sutured subcomplex structure on X, provided that the preimage of the singular locus of X under emb_X is properly embedded in $\partial_v \mathrm{ab}(X)$. If X is a pre-sutured subcomplex of M, by *suitable pre-sutured triangulation* of X we mean the image of a suitable pre-sutured triangulation $\mathcal{T}$ of $\mathrm{ab}(X)$ under emb_X, provided that the preimage of the singular locus of X under emb_X is a union of edges of $\mathcal{T}$.

2.3.2 Vertical boundary of a pre-sutured manifold

Let M be a pre-sutured manifold; we say that a subset Ω of $\partial_v M$ is *vertical* if it is a disjoint union of (possibly infinitely many) arcs properly embedded in $\partial_v M$, each of which intersects $\partial_0 M$ and $\partial_1 M$ in a single point. Let Ω be a disjoint union of finitely many closed connected vertical subsets of $\partial_v M$; we require that no component of Ω is an arc – in other words, each component of Ω is either a disc or an annulus. For $i \in \{0,1\}$, denote by Ω_i the intersection $\Omega \cap \partial_i M$. A partition of Ω into vertical arcs induces a homeomorphism

$$\uparrow : \Omega_0 \longrightarrow \Omega_1,$$

sending a point in Ω_0 to the other endpoint of the vertical arc containing it. We call this homeomorphism the *upward shift map*, noting that it is only defined up to isotopy.

We say that a triangulation $\mathcal{T}$ of Ω is *square-tiled* if, when looking at M from the outside, with Ω_0 at the bottom and Ω_1 at the top, each component of Ω is triangulated as shown in Figure 2.7a; in other words if, from the same perspective, each component of Ω is a union of vertical squares, each of which is triangulated as a square with a diagonal going from the bottom left to the top right corner. Note that whether $\mathcal{T}$ is square-tiled or not depends in general on the orientation of M.

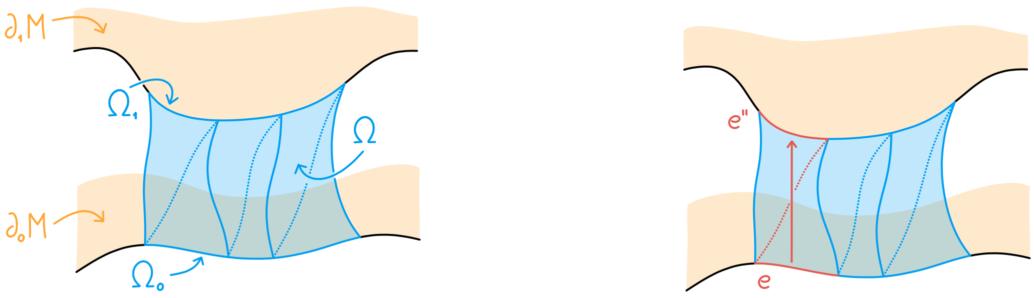

(a) (b)

Figure 2.7. (a) A square-tiled triangulation of a vertical subsurface Ω of $\partial_v M$. (b) The upward shift map of Ω with a square-tiled triangulation is a simplicial isomorphism.

If $\mathcal{T}$ is a square-tiled triangulation of Ω, then there is a well-defined simplicial isomorphism from Ω_0 to Ω_1 that is an upward shift map. This simplicial isomorphism is depicted in Figure 2.7b, and it can be described purely combinatorially as follows. Taking the same perspective as before, let e be an edge of Ω_0. Pivot counter-clockwise around its left endpoint to find the edge e' (this will be a diagonal edge). Then, pivot clockwise around the endpoint of e' lying on Ω_1 to find the edge e'', that will be entirely contained in Ω_1. The combinatorial upward shift map will then send e to e'' simplicially, with the appropriate orientation. The point of this (perhaps overly detailed) description is to show that the upward shift map is a combinatorial object that can be computed algorithmically. In particular, if M is described by a pre-sutured triangulation that restricts to a square-tiled triangulation of a simplicial vertical surface $\Omega \subseteq \partial_v M$, then the upward shift map $\uparrow\colon \Omega_0 \to \Omega_1$ is a simplicial isomorphism that can be computed explicitly in polynomial time in the size of the triangulation.

2.3.3 Cutting along a normal surface

If F is a surface properly embedded in a compact 3-manifold M, we can *cut M along F* to obtain a new compact 3-manifold. More precisely, we define $M \setminus\!\setminus F$ to be the closure in M of $M \setminus \mathcal{N}(F)$, where here and in the following $\mathcal{N}$ denotes a closed regular neighbourhood. If, additionally, the 3-manifold M is oriented and the surface F is transversely oriented, then $M' = M \setminus\!\setminus F$ has a natural pre-sutured manifold structure. In particular, we set $\partial_v M' = M' \cap \partial M$, so that $\partial_h M'$ is the union of two parallel copies of F; by convention, we set $\partial_0 M'$ to be one of these copies, and $\partial_1 M'$ to be the other, such that the transverse orientation of F points from $\partial_1 M'$ to $\partial_0 M'$ at each point of F. When M is a compact

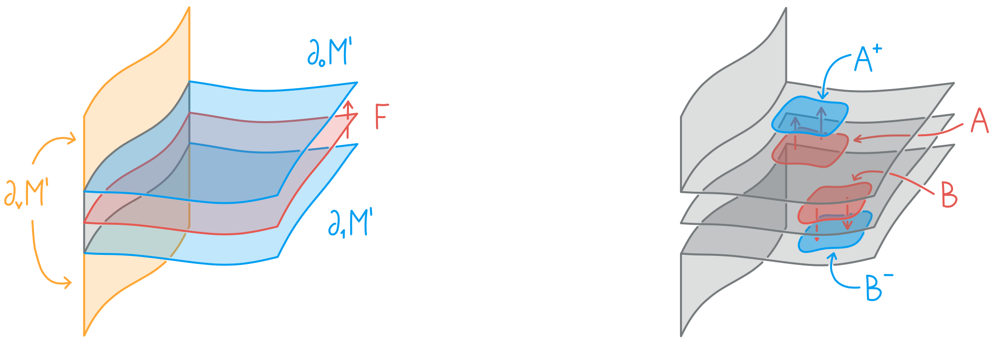

(a) (b)

Figure 2.8. (a) The subsurfaces $\partial_0 M'$, $\partial_1 M'$, and $\partial_v M'$ of $\partial M'$ inducing the pre-sutured manifold structure on $M \setminus\!\setminus F$; the transverse orientation of F is represented by an arrow. (b) The positive and negative push-offs of the normal surface F; two arbitrary subsets A and B of F are depicted, together with their push-offs $A^+ \subseteq \partial_0 M'$ and $B^- \subseteq \partial_1 M'$.

oriented 3-manifold and F is a transversely oriented surface properly embedded in M, we will always implicitly endow $M \setminus\!\setminus F$ with this pre-sutured manifold structure, as depicted in Figure 2.8a.

The fact that $\mathcal{N}(F)$ is an interval bundle over F provides two homeomorphisms

$$(-)^+ \colon F \longrightarrow \partial_0 M' \quad \text{and} \quad (-)^- \colon F \longrightarrow \partial_1 M',$$

that we call the *positive push-off* and the *negative push-off* respectively; they are only defined up to isotopy, and they are depicted in Figure 2.8b.

When M is described by a triangulation $\mathcal{T}$ and F is normal, we can put more structure on $M \setminus\!\setminus F$. Firstly, we remark that a transverse orientation on F can be encoded combinatorially by a list of points in $F \cap \mathcal{T}^{(1)}$, one for each component of F; each of these points – say x – is equipped with a sign, depending on whether the transverse orientation at x agrees with the orientation of the edge of $\mathcal{T}$ containing x. We will see in Section 2.4.3 how to algorithmically transport the transverse orientation to a point different from the one where it is presented combinatorially. If M is oriented and F is transversely oriented, then F inherits a natural orientation; we will not emphasise this, but we will always implicitly assume that F is oriented in this way.

Let T be a tetrahedron of $\mathcal{T}$. A *piece* of M' is a component of $T \cap M'$ (here, we think of T as an abstract tetrahedron). A *parallelity piece* of M' is a piece of M' that lies between two normal discs of the same type. If a piece is not a parallelity piece, then

Figure 2.9. (a) The parallelity pieces of M' in a tetrahedron T of $\mathcal{T}$ are the components of $T \cap M'$ that lie between two normal discs of the same type. (b) The gut pieces of M' in a tetrahedron T of $\mathcal{T}$ are the components of $T \cap M'$ that are not parallelity pieces.

we call it a *gut piece*. See Figure 2.9 for an example of parallelity and gut pieces in a tetrahedron. The union of all parallelity pieces in M' is called the *parallelity bundle*, and the union of all gut pieces is called the *guts*.

2.3.4 Guts and parallelity bundle as pre-sutured manifolds

Let M, $\mathcal{T}$, F, and M' be as above. Denote by X the guts of M', and by Y the parallelity bundle. Note that X and Y are sub-3-complexes of M' (with the polyhedral cell structure whose cells are the pieces of M'), and the same holds for every component of X and Y; we call a component of X a *gut component*, and a component of Y a *parallelity component*. We remark that $\mathrm{ab}(Y)$ has a natural interval bundle structure over some compact surface, since it is obtained by gluing together parallelity pieces, each of which is a product interval bundle over a normal disc. We call this interval bundle structure the *intrinsic interval bundle structure* of $\mathrm{ab}(Y)$. Each parallelity component, endowed with its intrinsic interval bundle structure, is either a product interval bundle over a compact orientable surface, or a twisted interval bundle over a compact non-orientable surface; we will mostly be dealing with the former case, but we cannot a priori exclude the latter. Note that, importantly, the intrinsic interval bundle structure does not necessarily agree with the pre-sutured manifold structure of M', in the sense that $\mathrm{emb}_Y(\partial_0 \,\mathrm{ab}(Y))$ might intersect $\partial_1 M'$ (or vice versa). It is, however, true that

$$\partial_h \,\mathrm{ab}(Y) = \mathrm{emb}_Y^{-1}(\partial_h M').$$

There is some ambiguity in the choice of intrinsic interval bundle structure for $\mathrm{ab}(Y)$ – namely, for each component C of $\mathrm{ab}(Y)$, the surfaces $\partial_0 C$ and $\partial_1 C$ can be interchanged. The only requirement we impose is that $\mathrm{emb}_Y(\partial_0 \mathrm{ab}(Y))$ and $\mathrm{emb}_Y(\partial_1 \mathrm{ab}(Y))$ are disjoint; it is easy to see that this can always be arranged.

Even though $\mathrm{ab}(X)$ is not necessarily an interval bundle, we can still endow it with a pre-sutured manifold structure, by setting

$$\begin{aligned}\partial_0 \mathrm{ab}(X) &= \mathrm{emb}_X^{-1}(\partial_0 M'),\\ \partial_1 \mathrm{ab}(X) &= \mathrm{emb}_X^{-1}(\partial_1 M'), \text{ and}\\ \partial_v \mathrm{ab}(X) &= \mathrm{emb}_X^{-1}(\partial_v M' \cup Y).\end{aligned}$$

We will always endow X and Y with the pre-sutured subcomplex structure induced by the pre-sutured manifold structures of $\mathrm{ab}(X)$ and $\mathrm{ab}(Y)$ described above. We will also say that a subset $\Omega \subseteq \partial_v X$ is a vertical subset of $\partial_v X$ to mean that $\mathrm{emb}_X^{-1}(\Omega)$ is a vertical subset of $\partial_v \mathrm{ab}(X)$, and similarly for Y.

Finally, we extend these pre-sutured manifold structures to the submanifolds $\mathrm{thin}(X)$ and $\mathrm{thick}(Y)$ of M'. Suppose – as we will always implicitly assume – that $\mathrm{thin}(X)$ and $\mathrm{thick}(Y)$ are chosen so that their intersection is contained in $\partial\,\mathrm{thin}(X)$ and in $\partial\,\mathrm{thick}(Y)$. We can then give $\mathrm{thin}(X)$ the following pre-sutured manifold structure:

$$\begin{aligned}\partial_0 \mathrm{thin}(X) &= \mathrm{thin}(X) \cap \partial_0 M',\\ \partial_1 \mathrm{thin}(X) &= \mathrm{thin}(X) \cap \partial_1 M', \text{ and}\\ \partial_v \mathrm{thin}(X) &= \mathrm{thin}(X) \cap (\partial_v M' \cup \mathrm{thick}(Y)).\end{aligned}$$

The submanifold $\mathrm{thick}(Y)$ of M' is an interval bundle over $\mathrm{thick}(\partial_0 Y)$, and it is therefore endowed with an intrinsic pre-sutured manifold structure – which, once again, does not necessarily agree with that of M'. We resolve the ambiguity of how to choose $\partial_0 C$ and $\partial_1 C$ for each component C of $\mathrm{thick}(Y)$ by requiring that $\partial_i Y \subseteq \partial_i \mathrm{thick}(Y)$ for $i \in \{0, 1\}$; this choice is made so that the natural retraction of $\mathrm{thick}(Y)$ onto Y sends $\partial_i \mathrm{thick}(Y)$ to $\partial_i Y$ for $i \in \{0, 1\}$ and $\partial_v \mathrm{thick}(Y)$ to $\partial_v Y$.

2.3.5 Triangulating the guts

Let us keep the notation of the previous paragraph. Recall that F inherits a natural triangulation from $\mathcal{T}$. We now describe how to triangulate the horizontal boundary and the guts of M'. Firstly, note that $\partial_h M'$ is itself a normal surface, since it intersects each tetrahedron of $\mathcal{T}$ in a collection of normal discs; more precisely, there is a natural one-to-two correspondence between the normal discs of F and the normal discs of $\partial_h M'$. Therefore, we can triangulate $\partial_h M'$ as a normal surface, using 3 triangles for each normal triangle and 4 triangles for each normal quadrilateral. With these canonical triangulations of F and $\partial_h M'$, the positive and negative push-offs of F can be naturally chosen to be simplicial isomorphisms. More precisely, for each tetrahedron T of $\mathcal{T}$, the positive push-off can be taken to map each normal disc $D \subseteq F \cap T$ to $\partial_0 M' \cap \mathcal{N}(D)$ simplicially; this is shown in Figure 2.10a. A similar statement holds for the negative push-off.

In this setting, we can define the *transfer map*. This is a simplicial map, defined on the largest sub-2-complex of F whose positive push-off is contained in Y. To define it, let D be a normal disc of F contained in a tetrahedron T of $\mathcal{T}$; suppose that D^+ is contained in a parallelity piece of T. Then there is a unique normal disc E in T, of the same type as D, such that either E^+ or E^- cobounds a parallelity piece of T with D^+. Suppose, for concreteness, that E^+ cobounds a parallelity piece P with D^+ (the case where E^+ is replaced by E^- is analogous). Let $f\colon D^+ \to E^+$ be the simplicial isomorphism induced by the interval bundle structure of P. Then the transfer map $\triangle$ sends a point $x \in D$ to the unique point $\triangle(x) \in E$ such that

$$\triangle(x)^+ = f(x^+).$$

We now construct a triangulation of the guts of M', building it piece by piece. Let P be a gut piece of M'; in other words, it is a component of $X \cap T$ for some (abstract) tetrahedron T of $\mathcal{T}$. We can think of P as a polyhedron, with some faces coming from $P \cap \partial_h X$, and the other faces coming from $P \cap \partial T$. We triangulate the faces of $P \cap \partial_h X$ according to the triangulation of $\partial_h M'$ we have already defined, as shown in Figure 2.10b. Then, we note that $P \cap Y$ is a union of four-sided faces. Some of these – namely, those that intersect both $\partial_0 M'$ and $\partial_1 M'$ – are vertical subsets of $\partial_v X$; denote the union of these faces by V_P. We add a diagonal edge to each face in V_P, so that the triangulation of V_P as a vertical subset of $\partial_v X$ is square-tiled, as depicted in Figure 2.10c; note that the choice of diagonal edge is unique, given that P is oriented. We then triangulate the remaining faces of P – that is, those in that are contained in ∂T but not in V_P – by coning over a vertex. More precisely, let t be the number of tetrahedra of $\mathcal{T}$. We fix a total order on the vertices of the gut pieces of M'; here, we mean that if two vertices come from possibly different gut pieces but are identified in M, then they should only be counted once. Since there are at most $6t$ gut pieces, independently of F, we can assume this order is part of the data describing the triangulation $\mathcal{T}$. Then, each relevant face of P can be triangulated by coning over its smallest vertex, as shown in Figure 2.10d. Finally, we triangulate the interior of P by coning the triangulation of ∂P over the smallest vertex of P.

By repeating this construction for each gut piece of M', we obtain a triangulation $\mathcal{R}$ of $\mathrm{ab}(X)$. Note that triangulations of different gut pieces are compatible along their common faces, since these faces are triangulated by coning over the same vertex. By construction, the triangulation $\mathcal{R}$ is in fact a pre-sutured triangulation, compatible with the pre-sutured manifold structure of $\mathrm{ab}(X)$. This triangulation induces a pre-sutured triangulation of X, compatible with the pre-sutured sub-3-complex structure of X; we will blur the distinction between the two, thinking of $\mathcal{R}$ as both a triangulation of $\mathrm{ab}(X)$ and a triangulation of X. The restriction of $\mathcal{R}$ to $\partial_h \mathrm{ab}(X)$ is easily seen to be flapless. Moreover, for each gut piece P, the vertical subset V_P of $\partial_v \mathrm{ab}(X)$ is simplicial and square-tiled.

We have already remarked how the number of gut pieces of M' is at most linear in t. The following elementary result gives more granular bounds concerning the triangulation of the guts.

Proposition 2.10 (Bounds on the triangulation of the guts)**.** *Let $\mathcal{T}$ be a triangulation of a compact oriented 3-manifold M with t tetrahedra. Let F be a transversely oriented normal surface in M, and let $M' = M \setminus\setminus F$. Then the pre-sutured triangulation $\mathcal{R}$ of the guts X of M' satisfies:*

(i) $|\mathcal{R}| \leqslant 50t$;

(ii) $\mathrm{area}(\partial_h X) \leqslant 32t$;

(iii) $\mathrm{area}(\partial_v X) \leqslant 36t$;

Proof. Fix an (abstract) tetrahedron T of $\mathcal{T}$. It is easy to see that $\partial_h X \cap T$ is a union of $n_t \leqslant 8$ normal triangles (0 or 2 per vertex of T) and $n_q \in \{0, 2\}$ normal quadrilaterals. As a consequence, the portion of $\partial_h X$ contained in T is triangulated with

$$3n_t + 4n_q \leqslant 32$$

triangles; this proves the second claim.

Let us now focus on a face R of T. Denote by r the intersection $\partial_h X \cap R$. By our previous remark on $\partial_h X \cap T$, we see that r is a union of $m \leqslant 8$ normal arcs in R. A

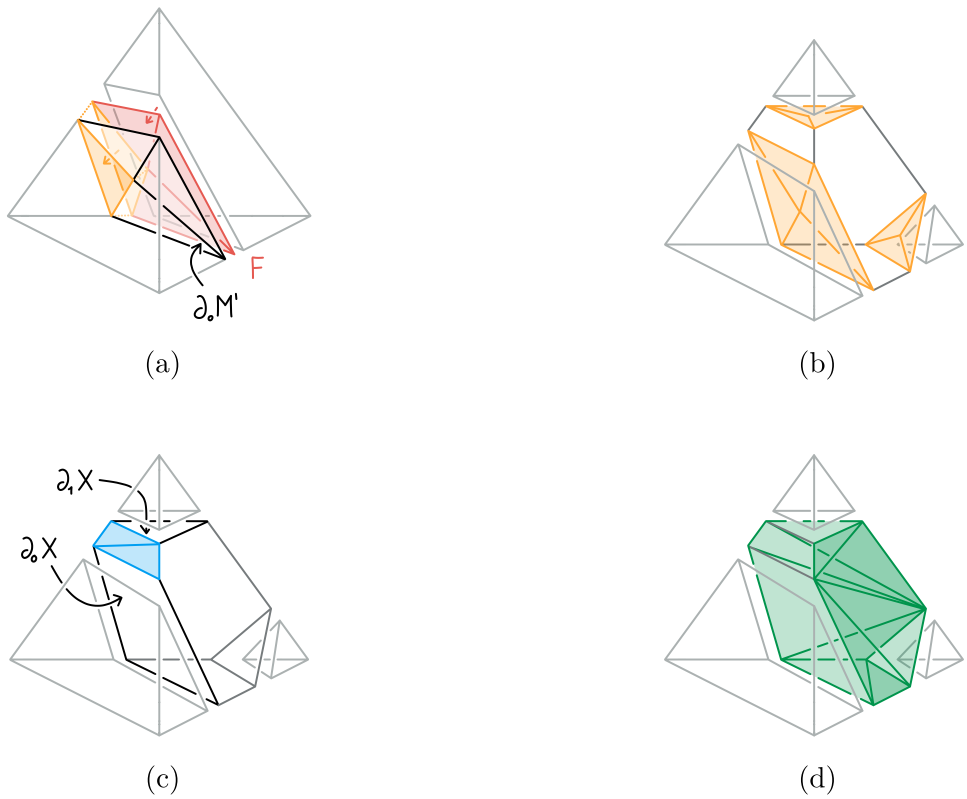

Figure 2.10. (a) The positive push-off of F maps each normal disc of F to a normal disc of $\partial_0 M'$ simplicially. The arrow on F represents the transverse orientation of F, while the arrow from a triangle of F to a triangle of $\partial_0 M'$ represents the positive push-off. (b) The faces of a gut piece that lie on $\partial_h X$ are triangulated according to the triangulation of $\partial_h M'$. (c) The faces of the gut piece that lie in Y and are vertical in $\partial_v X$ are triangulated as square-tiled subsets of $\partial_v X$. (d) The remaining faces of the gut piece are triangulated by coning over vertices.

simple counting argument reveals that $X \cap R$ has $1 + m/2$ polygonal components, and the total number of sides of these components is $3 + 2m$. Since each component of $X \cap R$ is triangulated by coning over a vertex, the total number of triangles in $X \cap R$ is

$$\text{area}(X \cap R) = 3 + 2m - 2(1 + m/2) = 1 + m \leqslant 9. \tag{2.1}$$

Since each tetrahedron has 4 faces, and $\partial_v X \subseteq \mathcal{T}^{(2)}$, this proves the third claim.

Finally, let us shift our attention back to the (arbitrary) tetrahedron T. A counting argument similar to the above shows that $X \cap T$ has $1 + n_t/2 + n_q/2$ polyhedral components – namely, the gut pieces of T. We have already seen that $\partial_h X \cap T$ is triangulated with $3n_t + 4n_q$ triangles. The intersection $\partial_h X \cap \partial T$ consists of $3n_t + 4n_q$ normal arcs. By applying (2.1) to each face of T, we see that $X \cap \partial T$ is triangulated with $4 + 3n_t + 4n_q$ triangles. Therefore, the total number of triangles in the triangulation of $\partial(X \cap T)$ is

$$(3n_t + 4n_q) + (4 + 3n_t + 4n_q) = 4 + 6n_t + 8n_q.$$

Since each vertex of $X \cap T$ has degree at least 3, and each component of $X \cap T$ is triangulated by coning the triangulation of its boundary over a vertex, we conclude that the number of tetrahedra of $X \cap T$ is at most

$$4 + 6n_t + 8n_q - 3\left(1 + \frac{1}{2}n_t + \frac{1}{2}n_q\right) = 1 + \frac{9}{2}n_t + \frac{13}{2}n_q \leqslant 1 + 36 + 13 = 50.$$

This concludes the proof of the first claim. ▢

We remark that the construction of the triangulation $\mathcal{R}$ of $\text{ab}(X)$ is very explicit. Given the bounds of Proposition 2.10, we can construct this triangulation algorithmically, in polynomial time in t and $\log w(F)$.

Finally, we introduce a notation for neighbourhoods of sub-2-complexes of F in M'. Let T be a triangle of F. If T^+ is contained in a parallelity piece of M', then we write $\mathcal{N}^+(T)$ for the unique parallelity piece containing T^+. If, instead, the positive push-off of T is contained in a gut piece of M', then we write $\mathcal{N}^+(T)$ for the unique gut component containing T^+. For a sub-2-complex G of F, we let

$$\mathcal{N}^+(G) = \bigcup_{T \subseteq G} \mathcal{N}^+(T),$$

where the union is taken over all triangles T of G. An analogous definition is given for $\mathcal{N}^-(G)$ in terms of the negative push-off. It is easy to see that $\mathcal{N}^+(G)$ and $\mathcal{N}^-(G)$ are sub-3-complexes of M'. Moreover, if G^+ is contained in the parallelity bundle of M', then $\mathcal{N}^+(G) \cong G \times [0,1]$ and $\mathrm{ab}(\mathcal{N}^+(G)) \cong \mathrm{ab}(G) \times [0,1]$; a similar statement holds for $\mathcal{N}^-(G)$.

2.4 Agol-Hass-Thurston orbit-counting algorithm

2.4.1 Orbit-counting algorithm

In [1], Agol, Hass, and Thurston introduce an algorithm to count the number of orbits of a given collection of pairings on an interval. More precisely, the setting is as follows. Let $N \geqslant 1$ be an integer. A *pairing* on $[1, N] \cap \mathbb{Z}$ is an order-preserving or order-reversing bijection

$$g\colon [a,b] \cap \mathbb{Z} \longrightarrow [c,d] \cap \mathbb{Z},$$

where $1 \leqslant a \leqslant b \leqslant N$ and $1 \leqslant c \leqslant d \leqslant N$ are integers. Let

$$\{g_i\colon [a_i,b_i] \cap \mathbb{Z} \longrightarrow [c_i,d_i] \cap \mathbb{Z}\}_{1 \leqslant i \leqslant k}$$

be a collection of pairings on $[1, N] \cap \mathbb{Z}$. These pairings induce an equivalence relation on $[1, N] \cap \mathbb{Z}$, generated by

$$\{x \sim g_i(x) : 1 \leqslant i \leqslant k, a_i \leqslant x \leqslant b_i\}.$$

We call the equivalence classes of this relation *orbits*. Note that a pairing $[a,b] \cap \mathbb{Z} \to [c,d] \cap \mathbb{Z}$ can simply be described by the four integers a, b, c, and d, together with a sign indicating whether the pairing is order-preserving or order-reversing. Therefore, the overall amount of information needed to describe the collection $g_1, \ldots, g_k$ is linear in $k \cdot (\log N + 1)$. The algorithm of Agol, Hass, and Thurston, in its most basic form, computes the number of orbits of a given collection of pairings.

Theorem 2.11 ([1, Theorem 12]). *There is an algorithm that takes as input an integer $N \geqslant 1$ and a collection $g_1, \ldots, g_k$ of pairings on $[1, N] \cap \mathbb{Z}$, and outputs the number of orbits. The running time of the algorithm is polynomial in $\log N$ and k.*

Even though this algorithm works in the full generality described above, its primary application – and the reason for its development – is computing properties of normal surfaces in triangulated 3-manifolds (or, similarly, of normal 1-manifolds in triangulated surfaces). For instance, the basic version of the algorithm of Agol, Hass, and Thurston stated above can be used to compute the number of components of a normal surface F in a triangulated 3-manifold M. We will now give an account of this classical result, already described in [1, Corollary 14], in order to familiarise the reader with the algorithm, and set up the stage for more complex applications.

Let N be the total number of normal discs of F; in other words, the integer N is the sum of the coordinates of the vector $\mathbf{v}_F$. We think of integers in $[1, N] \cap \mathbb{Z}$ as the normal discs of F, in such a way that two consecutive normal discs of the same type in the same tetrahedron are represented by consecutive integers. Therefore, the interval $[1, N] \cap \mathbb{Z}$ is split into $7t$ subintervals representing the different types of normal discs, where t is the number of tetrahedra of the triangulation of M. Let R be a triangle in the 2-skeleton of M, and suppose that R is adjacent to two (not necessarily distinct) tetrahedra T_1 and T_2. Each of the three types of normal arcs in R determines up to three pairings, pairing normal discs of T_1 with normal discs of T_2. These pairings are essentially the ones that determine the matching equations for normal surfaces; hence, we refer the reader to Figure 2.5b for a visual representation of the situation. If we repeat this procedure for each triangle R, we obtain a collection of $\mathcal{O}(t)$ pairings on $[1, N] \cap \mathbb{Z}$. By construction, the orbits of this collection are in natural bijection with the components of F; in fact, more precisely, an orbit consists exactly of the normal discs making up a component of F. Since $N = \mathcal{O}(w(F))$, we can then apply the algorithm of Theorem 2.11 to compute the number of components of F; the running time of the algorithm is polynomial in t and $\log w(F)$.

2.4.2 Weighted orbit-counting algorithm

The most general version of the algorithm of Agol, Hass, and Thurston allows the interval $[1, N] \cap \mathbb{Z}$ to be equipped with a weight function with values in $\mathbb{Z}^d$ for some $d \geqslant 1$. More precisely, to ensure efficiency, the weight function is given as a partition of $[1, N] \cap \mathbb{Z}$ into intervals $[p_j, q_j] \cap \mathbb{Z}$ for $1 \leqslant j \leqslant m$, together with vectors $\mathbf{z}_j \in \mathbb{Z}^d$ for $1 \leqslant j \leqslant m$. This data defines a weight function $z\colon [1, N] \cap \mathbb{Z} \to \mathbb{Z}^d$, by setting $z(x) = \mathbf{z}_j$ if $p_j \leqslant x \leqslant q_j$. For a subset $X \subseteq [1, N] \cap \mathbb{Z}$, we define the weight of X to be the sum of the weights of the elements of X. The general version of the algorithm of Agol, Hass, and Thurston operates in this setting, and computes the list of orbits of a given collection of pairings, together with the weight of each orbit. More precisely, the statement is as follows.

Theorem 2.12 ([1, Theorem 16]). *There is an algorithm that takes as input*

- *integers $N \geqslant 1$ and $d \geqslant 1$,*

- *a collection* $g_1, \dots, g_k$ *of pairings on* $[1, N] \cap \mathbb{Z}$,
- *a partition of* $[1, N] \cap \mathbb{Z}$ *into intervals* $[p_j, q_j] \cap \mathbb{Z}$ *for* $1 \leqslant j \leqslant m$, *and*
- *vectors* $\mathbf{z}_j \in \mathbb{Z}^d$ *for* $1 \leqslant j \leqslant m$,

and outputs the list of orbits with their weights. The running time of the algorithm is polynomial in k, m, d, $\log D$, *and* $\log N$, *where* D *is the maximum* ℓ^1*-norm of the vectors* $\mathbf{z}_1, \dots, \mathbf{z}_m$.

To be even more precise, the algorithm as described in [1] outputs a list of triples of the form $(r, s, \mathbf{w})$, signifying that each point in $[r, s] \cap \mathbb{Z}$ is a representative of a different orbit, all having the same weight $\mathbf{w} \in \mathbb{Z}^d$. However, for our purposes, we can assume that the output of the algorithm is simply the list of weights attained by the orbits, together with their multiplicities. We also remark that running the algorithm of Theorem 2.12 with $d = 1$ and a constant weight function $z(x) = 0$ for all x is equivalent to running the basic orbit-counting algorithm of Theorem 2.11.

The weighted orbit-counting algorithm can (and was developed to) be applied to the setting of normal surfaces in triangulated 3-manifolds, where it can be used to extract finer information than just the number of components. For instance, as the authors remark in [1, Corollary 17], one can apply the algorithm to compute the Euler characteristic of each component of a normal surface F in a triangulated 3-manifold M, in time that is polynomial in the size of the triangulation of M and $\log w(F)$. The procedure is a simple extension of what we have described above for counting the number of components. In particular, we take the same interval $[1, N] \cap \mathbb{Z}$ to be in bijection with the normal discs of F, with the same collection of pairings induced by the triangles of the 2-skeleton of M. Additionally, we endow the interval $[1, N] \cap \mathbb{Z}$ with a weight function $z \colon [1, N] \cap \mathbb{Z} \to \mathbb{Z}^{7t}$, where t is the number of tetrahedra of M. This weight function maps each normal disc of type q to the vector having a 1 in the q-th coordinate and 0 elsewhere (here, we think of types of normal discs as being indexed by the integers $1, \dots, 7t$). We can then run the algorithm of Theorem 2.12, that will output a list of components of F, each with its normal vector in $\mathbb{Z}^{7t}$. We conclude by noting that the Euler characteristic of a normal surface is an explicitly computable linear function of its normal vector.

2.4.3 Applications to normal curves and surfaces

We now collect a few results that concern the algorithmic computation of properties of normal curves in triangulated surfaces and normal surfaces in triangulated 3-manifolds, beyond the examples we have already discussed. These results are all well-known consequences of the algorithm of Agol, Hass, and Thurston; therefore, we will only provide sketches of the arguments, glossing over details such as the precise definitions of the pairings and weight functions.

Proposition 2.13 (Finding the components of a normal 1-manifold)**.** *There is an algorithm that takes as input a triangulation of a compact surface F with t triangles and a normal* 1*-manifold a in F, and outputs the list of components of a, each with its normal vector in $\mathbb{Z}^{3t}$. The running time of the algorithm is polynomial in t and* $\log w(a)$*.*

Proof. The algorithm is a 2-dimensional version of the one we described above for normal surfaces. We take an interval $[1, N] \cap \mathbb{Z}$ to be in bijection with the normal arcs of a, with pairings induced by the identifications between endpoints of normal arcs along the edges of the triangulation (see Figure 2.2b). We also use a weight function to keep track of the types of normal arcs. Running the algorithm of Theorem 2.12 will then output a list of components of a, each with its normal vector in $\mathbb{Z}^{3t}$. ▢

Proposition 2.14 (Detecting inessential normal 1-manifolds)**.** *There is an algorithm that takes as input a triangulation of a compact orientable surface F with t triangles and a connected normal* 1*-manifold a in F, and outputs whether a is a curve bounding a disc in F, whether a is a curve bounding a collared annulus in F, and whether a is a boundary-parallel arc in F. The running time of the algorithm is polynomial in t and* $\log w(a)$*.*

Proof. We can assume that F is connected and that a is non-empty. We take an interval $[1, N] \cap \mathbb{Z}$ to be in bijection with the components of $T \setminus\setminus a$, where T ranges over the triangles of F; we call these components *pieces*. For each triangle T, there are at most 7 "types" of pieces of $T \setminus\setminus a$. In fact, there is at most one piece adjacent to each vertex of T, at most one piece adjacent to all three edges of T, and three more types of "parallelity pieces", one for each type of normal arc in T. We then define a collection of pairings on $[1, N] \cap \mathbb{Z}$, representing the pairs of pieces that are glued together along parts of their boundary. Note that, if the points in $[1, N] \cap \mathbb{Z}$ are arranged so that consecutive parallelity pieces of the same type are adjacent, then the total number of pairings can be kept linear in t. We also use a weight function to keep track of the types of pieces. The algorithm of Theorem 2.12 will then output a list of components of $T \setminus\setminus a$, each with a vector in $\mathbb{Z}^{7t}$ counting the number of pieces of each type. It is not hard to prove that the Euler characteristic of a component of $T \setminus\setminus a$ is an explicitly computable linear function of the output vector. Therefore, we can obtain a list of the Euler characteristics of the components of $T \setminus\setminus a$. If there is only one such component, then the answer to all three questions in the statement is "no". If there are two components and one of them has Euler characteristic 1, then a is either a curve bounding a disc or a boundary-parallel arc; we can easily differentiate between the two cases by checking whether a has boundary or not. If there are two components and one of them has Euler characteristic 0, then a is a curve bounding a collared annulus. Finally, if there are two components and both have Euler characteristic at most -1, then the answer to all three questions is "no". ▢

Proposition 2.15 (Finding the components of a normal surface)**.** *There is an algorithm that takes as input a triangulation of a compact 3-manifold M with t tetrahedra and a normal surface F in M, and outputs the list of components of F, each with its normal vector in $\mathbb{Z}^{7t}$, its Euler characteristic, and the number of its boundary components. The running time of the algorithm is polynomial in t and $\log w(F)$.*

Proof. We have already discussed above how to compute the list of components of F, each with its normal vector and its Euler characteristic. For each component F' of F, we can then compute the normal vector of the multicurve $\partial F'$ in the triangulated surface ∂M, and use the algorithm of Proposition 2.13 to compute the number of components of $\partial F'$. ▢

Proposition 2.16 (Computing orientability and transverse orientation of a normal surface)**.** *There is an algorithm that takes as input*

- *a triangulation $\mathcal{T}$ of a compact orientable 3-manifold M with t tetrahedra,*
- *a connected normal surface F in M,*
- *two points $x, y \in F \cap \mathcal{T}^{(1)}$, and*
- *a transverse orientation at x,*

and outputs:

- *whether F is orientable or not;*
- *assuming F is orientable, the transverse orientation at y that is compatible with the given one at x.*

The running time of the algorithm is polynomial in t and $\log w(F)$.

Proof. The key fact is that the vector $2\mathbf{v}_F$ is the normal vector of a surface G in M which is the horizontal boundary of a regular neighbourhood of F in M. In particular, we see that G is connected if and only if F is not orientable. Therefore, we can apply the algorithm of Proposition 2.15 to the normal surface G and deduce whether F is orientable or not. Assuming it is, we now need to compute the transverse orientation at y that is compatible with the given one at x. Thinking of G as the unit normal bundle of F in M, a transverse orientation at x corresponds to a point x' of $G \cap \mathcal{T}^{(1)}$ – namely, one of the two points of G that lie above x – which can be readily computed. Finding the correct transverse orientation at y essentially amounts to deciding which of the two points of G that lie above y belongs to the same connected component of G as x'. To this aim, we describe an "orbit-tracking" trick that will also be useful later.

Pick one of the two points of G that lie above y and call it y'. In order to decide whether x' and y' lie in the same connected component of G, we run a slightly modified version of the algorithm of Proposition 2.15. We add two extra coordinates to the weight function used therein – call them the x' coordinate and the y' coordinate. The x' coordinate will be 1 on the normal discs of G that are adjacent to x', and 0 elsewhere. Similarly, the y' coordinate will be 1 on the normal discs of G that are adjacent to y', and 0 elsewhere. We then run the algorithm of Theorem 2.12 with this modified weight function. If the algorithm outputs an orbit where both the x' coordinate and the y' coordinate are positive, then x' and y' lie in the same connected component of G, and hence y' defines the correct transverse orientation at y. Otherwise, the points x' and y' lie in different connected components of G, and hence the transverse orientation at y is the opposite one. ▢

Recall that our combinatorial representation of a sub-2-complex G of a normal surface F in a triangulated 3-manifold M consists of two pieces of data: the edges of $\partial_F G$ and a list containing one triangle of F for each component of G. If we want to make use of this representation, we are faced with two tasks. Firstly, we need to be able to decide if a given representation is *valid*, in the sense that it comes from an actual sub-2-complex of F. Secondly, we need to be able to extract the relevant topological information about a sub-2-complex G from its combinatorial representation. This includes the Euler characteristic and number of boundary components of each component of G, but this is not enough: we also need to know how exactly the boundary components of $\mathrm{ab}(G)$ are glued to the rest of F. For this reason, we introduce the following definition. If F is oriented, then $\mathrm{ab}(G)$ inherits an orientation, and hence each boundary component b of $\mathrm{ab}(G)$ is canonically oriented. An *F-boundary sequence* for b is a sequence $e_1, \dots, e_k$, where each e_i is either an edge of $\partial_F G$ or the symbol ∂F, satisfying the following property: there exists a continuous orientation-preserving surjection $f\colon [0,k] \to b$ that is an embedding on $(0,k)$ and such that, for each $1 \leqslant i \leqslant k$, we have that

$$\begin{aligned} &\mathrm{emb}_G(f([i-1,i])) = e_i && \text{if } e_i \text{ is not } \partial F, \\ &\mathrm{emb}_G(f([i-1,i])) \subseteq \partial F && \text{if } e_i \text{ is } \partial F. \end{aligned}$$

Moreover, we ask that if $e_i = \partial F$ for some $1 \leqslant i \leqslant k$, then $e_{i+1} \neq \partial F$ (where indices are taken modulo k). Loosely speaking, one obtains an F-boundary sequence for b by walking along b (with the orientation induced by that of $\mathrm{ab}(C)$) and recording the sequence of edges of $\partial_F G$ that one encounters, interrupted by segments of ∂F. We remark that the F-boundary sequence for b is unique up to cyclic permutations.

Proposition 2.17 (Finding the components of a sub-2-complex of a normal surface)**.** *There is an algorithm that takes as input*

- *a triangulation of a compact oriented* 3*-manifold M with t tetrahedra,*
- *a connected transversely oriented normal surface F in M,*
- *a collection $\mathcal{E}$ of edges of F, and*
- *a collection $\mathcal{S}$ of triangles of F,*

and decides whether there exists a sub-2-complex G of F such that $\partial_F G$ consists precisely of the edges of $\mathcal{E}$, and $\mathcal{S}$ contains at least one triangle of each component of G. If this is the case, the algorithm also outputs, for each component C of G:

- *the list of triangles in $\mathcal{S}$ that are contained in C;*
- *the Euler characteristic of* $\mathrm{ab}(C)$*;*
- *the number of boundary components of* $\mathrm{ab}(C)$*;*
- *for each component b of $\partial\,\mathrm{ab}(C)$, an F-boundary sequence for b.*

The running time of the algorithm is polynomial in t, $\log w(F)$*, and the cardinalities of $\mathcal{E}$ and $\mathcal{S}$.*

Proof. Denote by $\mathcal{T}$ the triangulation of M, and by $\mathcal{R}$ the induced triangulation of F. Recall that the triangles of $\mathcal{R}$ are grouped into types, as detailed in Section 2.2.6. In particular, tetrahedron of $\mathcal{T}$ contains at most 16 types of triangles of $\mathcal{R}$: at most 4 for the unique type of normal quadrilaterals, and at most 3 per type of normal triangles. We can then arrange the triangles of $\mathcal{R}$ in an interval $[1, N] \cap \mathbb{Z}$, where triangles of the same type form contiguous blocks; within each block, the triangles are ordered so that triangles lying on consecutive normal discs are adjacent. It is easy to define a system of pairings on this interval, such that two triangles of $\mathcal{R}$ are paired if and only if they share an edge of $\mathcal{R}$ that is not in $\mathcal{E}$. We also define a weight function

$$z\colon [1, N] \cap \mathbb{Z} \longrightarrow \mathbb{Z}^d,$$

where the d coordinates encode several pieces of information for each triangle R of $\mathcal{R}$.

1. The first block of $16t$ coordinates keeps track of the type of R (here, we think of types as being indexed by the integers $1, \ldots, 16t$).

2. The second block of $|\mathcal{S}|$ coordinates keeps track of which triangle in $\mathcal{S}$ – if any – is equal to R.

3. The third block of $2|\mathcal{E}|$ coordinates keeps track of which edges in $\mathcal{E}$ are adjacent to R, and from which side.

4. There is also a fourth block of coordinates that keeps track, for each triangle R' of the same type as R that is adjacent to some edge in $\mathcal{E}$, whether R' appears before or after R, according to one of the two natural orders of the normal discs they lie on; we consider this order to be fixed once and for all. These coordinates will be used later to compute information about the boundary components of $\mathrm{ab}(G)$.

Run the algorithm of Theorem 2.12 on the interval $[1, N] \cap \mathbb{Z}$, the collection of pairings, and the weight function z we described above. We can now assess whether a sub-2-complex G of F as required in the statement exists. If any edge in $\mathcal{E}$ lies in ∂F, then the answer is "no". If the union of the edges in $\mathcal{E}$, seen as a graph, has a vertex of degree 1, then the answer is "no". Each orbit in the output of the algorithm of Agol, Hass, and Thurston either contains a triangle in $\mathcal{S}$, or it does not; this can be assessed by looking at the second block of coordinates of the weight of the given orbit. If some edge in $\mathcal{E}$ is contained in 0 or 2 orbits containing a triangle in $\mathcal{S}$ (which can be assessed by looking at the third block of coordinates), then the answer is "no". Otherwise, the answer is "yes", and each orbit containing a triangle in $\mathcal{S}$ corresponds to a component of $\mathrm{ab}(G)$; in other words, the sub-2-complex G is the union of the orbits that contain a triangle in $\mathcal{S}$.

Fix now a component C of $\mathrm{ab}(G)$. The weight $z(C)$ already contains the list of triangles in $\mathcal{S}$ that lie inside C. It is also not hard to see that the Euler characteristic of C is an explicitly computable linear function of the weight $z(C)$. Computing the number of boundary components of C and their F-boundary sequences is a bit more involved. First, we can use the data contained in $w(C)$ to construct an interval that represents the edges of ∂C, together with a collection of pairings encoding the adjacency of edges of ∂C. Note that we need to use the fourth block of coordinates here, to identify which edges of $\mathcal{E}$ lie on ∂C. Running the algorithm of Agol, Hass, and Thurston will then output the number of components of ∂C.

In order to compute the F-boundary sequences, start by picking an edge $e_1 \in \mathcal{E}$ that lies in ∂C. Note that we can recover which side of e_1 the component C lies on from $w(C)$. An application of Proposition 2.16 will allow us to deduce the correct orientation of the component b of ∂C containing e_1. We can then start walking along b according to this orientation, explicitly recording the sequence $e_1, \ldots, e_i$ of edges in $\mathcal{E}$ that we encounter. This process will terminate either when we get back to e_1 – in which case we have successfully computed the F-boundary sequence for b – or when we hit ∂F. If the latter happens, we set $e_{i+1} = \partial F$. In order to understand what the next edge $e_{i+2} \in \mathcal{E}$ is, we construct an interval representing the edges of $\partial C \cap \partial F$, together with a collection of pairings encoding the adjacency of said edges. Using the "orbit-tracking" trick described in the proof of Proposition 2.16, we can find the component of $\partial C \cap \partial F$ that comes after e_i in the F-boundary sequence of b, together with the following edge $e_{i+2} \in \mathcal{E}$. We can repeat this process until we get back to e_1, thus completing the F-boundary sequence for b. Using the same technique, we can compute the F-boundary sequences for all the components of ∂C that are not contained in ∂F. Once we have done this, since we know how many boundary components C has, we can deduce the number of components of ∂C that are contained in ∂F; the F-boundary sequence for these components is simply ∂F. ▢

Proposition 2.18 (Deciding containments of sub-2-complexes of normal surfaces)**.** *For a compact oriented 3-manifold M triangulated with t tetrahedra, a connected transversely oriented normal surface F in M, and a sub-2-complex G of F, the following questions can be answered by an algorithm with the specified running time.*

1. *Given a triangle T of F, decide whether T is contained in G, in polynomial time in t, $\log w(F)$, $|G|$, and $\ell(\partial_F G)$.*

2. *Given two triangles T_1 and T_2 in G, decide whether they belong to the same*

component of G, in polynomial time in t, $\log w(F)$, $|G|$, and $\ell(\partial_F G)$.

3. *Given a sub-2-complex G' of F, decide whether G' is contained in G, in polynomial time in t, $\log w(F)$, $|G|$, $\ell(\partial_F G)$, $|G'|$, and $\ell(\partial_F G')$.*

4. *Given a sub-2-complex G' of F, decide whether G' intersects G, and whether* $\mathrm{int}(G')$ *intersects* $\mathrm{int}(G)$*, in polynomial time in t, $\log w(F)$, $|G|$, $\ell(\partial_F G)$, $|G'|$, and $\ell(\partial_F G')$.*

Proof. The first two questions can be easily answered with the "orbit-tracking" trick described in the proof of Proposition 2.16. The third question reduces to the fourth one, since G' is contained in G if and only if $\mathrm{int}(G')$ is disjoint from $F \setminus G$. Therefore, we only need to show how to answer the fourth question. We have seen in the proof of Proposition 2.17 how to compute the list of components of $\mathrm{ab}(G)$ and $\mathrm{ab}(G')$; by checking all possible pairs of components, we can suppose, without loss of generality, that $\mathrm{ab}(G)$ and $\mathrm{ab}(G')$ are connected. If $\partial_F G$ and $\partial_F G'$ are both empty, then G and G' intersect if and only if they are contained in the same component of F, which can be decided with an "orbit-tracking" trick; the same holds for $\mathrm{int}(G)$ and $\mathrm{int}(G')$. Otherwise, it is not hard to see that $\mathrm{int}(G)$ and $\mathrm{int}(G')$ intersect if and only if there is a triangle of F that is contained in $G \cap G'$ and intersects $\partial_F G$ or $\partial_F G'$. Since the number of such triangles is bounded above by a polynomial in t, $\ell(\partial_F G)$, and $\ell(\partial_F G')$ (where t is the number of tetrahedra of M), we can check all possible triangles in polynomial time. Finally, if $\mathrm{int}(G)$ and $\mathrm{int}(G')$ are disjoint, then G and G' intersect if and only if $\partial_F G$ and $\partial_F G'$ intersect. This can be checked by direct inspection, thus providing an answer to the fourth question in the statement. ▢

2.4.4 Algorithms for curves in normal surfaces

The algorithm of Agol, Hass, and Thurston, and specifically the consequence we have stated in Proposition 2.17, can be used to efficiently operate on normal curves in normal surfaces. Recall that a normal curve a in a normal surface F (which, in turn, is properly embedded in a 3-manifold M triangulated with t tetrahedra) is encoded by listing the triangles of the sub-2-complex $\mathrm{supp}(a)$ of F, and describing a as a normal curve in $\mathrm{ab}(\mathrm{supp}(a))$. Many algorithmic problems about curves on surfaces – for instance, deciding whether two curves are isotopic – can be solved in polynomial time. However, one of the parameters in this polynomial dependence is (and has to be) the size of the triangulation of the surface; in our case, the number of triangles in the triangulation of F depends linearly on its weight. Therefore, if we were to naively apply these algorithms to the normal curve $a \subseteq F$, we would obtain a running time that is polynomial in $w(F)$. This is unacceptable for our purposes, since the normal surfaces we operate on will in general have weight that is exponential in t.

There is, however, a simple trick we can apply to circumvent this issue. The normal curves we will deal with have "small" support, meaning that they do not intersect the vast majority of the triangles of F. Therefore, we can take these large regions – disjoint from the support of our normal curves – and re-triangulate them, using the smallest possible number of triangles. This procedure will yield a new triangulation of F, that is unchanged on the support of the normal curves, but has a much smaller number of triangles. In particular, the size of this new triangulation will be linear in $|\chi(S)|$ and the area of the support of the normal curves – note that we cannot hope for a better dependence. This idea is described in detail in the following proposition.

Proposition 2.19 (Algorithmic retriangulation of normal surfaces)**.** *There is an algorithm that takes as input*

- *a triangulation of a compact orientable* 3*-manifold M with t tetrahedra,*
- *a connected orientable normal surface F in M, and*
- *a non-empty sub-*2*-complex F_0 of F,*

and outputs

- *a triangulation of a compact orientable surface G,*
- *a sub-*2*-complex G_0 of G, and*
- *a simplicial isomorphism $f\colon F_0 \to G_0$,*

such that the following hold:

1. $\operatorname{area}(G) \leqslant 22 \operatorname{area}(F_0) - 3\chi(F)$*;*
2. *f extends to a homeomorphism $F \to G$.*

The running time of the algorithm is polynomial in t, $|\chi(F)|$, $\log w(F)$, and $\operatorname{area}(F_0)$*.*

Proof. Fix an orientation of M, and a transverse orientation of F. Let F_1 be the closure of $F \setminus F_0$, which is a sub-2-complex of F. Note that $\ell(\partial_F F_1) = \ell(\partial_F F_0)$ is bounded above by in $3 \operatorname{area}(F_0)$; similarly, the number of components of F_1 is at most $3 \operatorname{area}(F_0)$. Therefore, we can apply the algorithm of Proposition 2.17 to the sub-2-complex F_1 of F; the running time will be polynomial in t, $\log w(F)$, and $\operatorname{area}(F_0)$. Then, for each component C of F_1, we obtain the Euler characteristic of $\operatorname{ab}(C)$, the number of boundary components of $\operatorname{ab}(C)$, and an F-boundary sequence for each of these boundary components.

Let C be an arbitrary component of F_1. We now show how to construct a triangulated oriented surface G_C such that there exists an orientation-preserving homeomorphism $g_C\colon \operatorname{ab}(C) \to G_C$ that restricts to a simplicial isomorphism on $C \cap F_0$. For each boundary component b of $\operatorname{ab}(C)$, construct an oriented circle S^1_b, endowed with the following cell structure. If $e_1, \ldots, e_k$ is the computed F-boundary sequence for b, then S^1_b has k edges, labelled in order $e_1, \ldots, e_k$. In particular, note that

$$\ell(S^1_b) = k \leqslant 2\ell(b \cap F_0) + 1. \tag{2.2}$$

Let

$$\partial G_C = \bigsqcup_{b \subseteq \partial \operatorname{ab}(C)} S^1_b,$$

where the disjoint union is taken over all boundary components of $\operatorname{ab}(C)$. By definition of F-boundary sequence, it is easy to see that there is an orientation-preserving homeomorphism $g'_C\colon \partial \operatorname{ab}(C) \to \partial G_C$ that restricts to a simplicial isomorphism on $C \cap F_0$. Moreover, it follows from (2.2) that

$$\ell(\partial G_C) \leqslant 2\ell(C \cap F_0) + m, \tag{2.3}$$

where m denotes the number of components of ∂G_C.

All we need to do now is to glue triangles to ∂G_C and to each other to realise an oriented surface G_C with the same genus as $\operatorname{ab}(C)$. We analyse several cases.

- If $\operatorname{ab}(C)$ is a disc with $\ell(\partial G_C) = 1$, then we can realise G_C with 1 triangle.
- If $\operatorname{ab}(C)$ is a disc with $\ell(\partial G_C) = 2$, then we can realise G_C with 2 triangles.
- If $\operatorname{ab}(C)$ is a disc with $\ell(\partial G_C) \geqslant 3$, then we can realise G_C with $\ell(\partial G_C) - 2$ triangles; these first three cases are depicted in Figure 2.11.
- If $\operatorname{ab}(C)$ is a planar surface with at least 2 boundary components, then we first add $m-1$ edges to ∂G_C to obtain a connected graph, and then we attach a triangulated disc that runs once over the edges of ∂G_C – consistently with the already defined orientation – and twice over the newly added edges; this is depicted in Figure 2.12a.

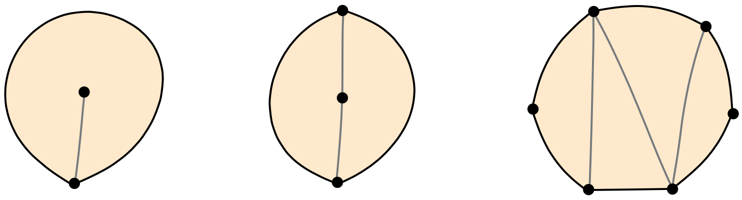

Figure 2.11. The construction of G_C when $\mathrm{ab}(C)$ is a disc. From left to right: the cases $\ell(\partial G_C) = 1$, $\ell(\partial G_C) = 2$, and $\ell(\partial G_C) \geqslant 3$.

An easy counting argument reveals that the resulting surface G_C is triangulated with

$$\mathrm{area}(G_C) = \ell(\partial G_C) + 2m - 4$$

triangles.

- If $\mathrm{ab}(C)$ is a surface of genus $g \geqslant 1$, then we first construct a $4g$-gon, and add m edges to connect it to all components of ∂G_C. We then attach a triangulated disc that runs twice over the edges of ∂G_C – consistently with the already defined orientation – and once over all the other edges. Finally, we pair the edges of the $4g$-gon in the standard commutator pattern to obtain a surface G_C of genus g; this is depicted in Figure 2.12b. The total number of triangles in G_C is

$$\mathrm{area}(G_C) = \ell(\partial G_C) + 2m + 4g - 2.$$

In all cases, we have constructed a triangulated oriented surface G_C that is orientation-preservingly homeomorphic to C. In particular, the homeomorphism $g'_C \colon \partial C \to \partial G_C$ extends to a homeomorphism $g_C \colon C \to G_C$. Moreover, if we denote by g the genus of $\mathrm{ab}(C)$, we have that

$$\begin{aligned} \mathrm{area}(G_C) &\leqslant \ell(\partial G_C) + 2m + 4g - 2 && \text{by case analysis} \\ &\leqslant 2\ell(C \cap F_0) + 3m + 4g - 2 && \text{by (2.3)} \\ &\leqslant 2\ell(C \cap F_0) - 3\chi(\mathrm{ab}(C)) + 4. && (2.4) \end{aligned}$$

We can then define

$$G = \left(F_0 \sqcup \bigsqcup_{C \subseteq F_1} G_C \right) \Big/ \{x \sim g_C(x) : C \text{ is a component of } F_1,\ x \in C \cap F_0\}.$$

(a) (b)

Figure 2.12. (a) The construction of G_C when C is a planar surface with at least 2 boundary components. There are $m-1$ edges (coloured in the picture) connecting the m boundary components of G_C, and a triangulated disc (shaded in the picture) attached to these edges and to ∂G_C. (b) The construction of G_C when C is a surface of genus $g \geqslant 1$. The edges of the $4g$-gon are coloured in the picture, with marks defining the gluing.

Since G is effectively constructed by gluing the triangulated surfaces G_C simplicially to the boundary of the 2-complex F_0, it is naturally equipped with a triangulation as follows. Recall that the edges of ∂G_C are labelled for each component C of F_1; the labels are either edges of $F_0 \cap F_1$ – in which case we glue the edges of ∂G_C to the edges of $C \cap F_0$ according to the labelling – or the symbol ∂F – in which case, no glueing is needed. We set G_0 to be the image of F_0 in G under the quotient map, and f to be the restriction of the quotient map to F_0; it is obvious from the construction that f extends to a homeomorphism $F \to G$. We remark that the construction we have described to produce the triangulated surface G is explicit, and can be realised in polynomial time in the number of triangles involved. As we are about to prove, this number is bounded above by a linear function of $\text{area}(F_0)$ and $|\chi(F)|$; this implies that the triangulation of G can be constructed in polynomial time in the size of the input.

Finally, we show how to obtain the desired upper bound on $\text{area}(G)$. Consider the following preliminary inequalities:

$$|F_1| \leqslant \ell(F_1 \cap F_0), \tag{2.5}$$

$$\chi(\text{ab}(F_0)) \leqslant \text{area}(F_0), \tag{2.6}$$

$$\chi(F) \leqslant \chi(\text{ab}(F_0)) + \chi(\text{ab}(F_1)). \tag{2.7}$$

Inequality (2.5) holds because the map sending each edge in $F_1 \cap F_0$ to the component of F_1 containing it is surjective. Inequality (2.6) follows from the fact that the only components of F_0 with positive Euler characteristic are discs, and each disc has area at least 1.

In order to prove (2.7), recall that $\mathrm{ab}(F_0)$ is homotopy equivalent to $\mathrm{thin}(F_0)$, and similarly for F_1. Therefore, if we write A for the closure in F of $F \setminus (\mathrm{thin}(F_0) \cup \mathrm{thin}(F_1))$, and a for $A \cap (\mathrm{thin}(F_0) \cup \mathrm{thin}(F_1))$, we have that

$$\chi(F) = \chi(\mathrm{ab}(F_0)) + \chi(\mathrm{ab}(F_1)) + \chi(A) - \chi(a). \tag{2.8}$$

Let $\Gamma = F_0 \cap F_1$, which we think of as a graph, with vertex set V and edge set E. Note that A deformation retracts onto Γ, via a retraction $r\colon A \to \Gamma$ that is an interval fibration except possibly at the vertices of Γ. When restricted to a, the map r is a 2-sheeted covering map, except possibly at the vertices of Γ; more precisely, each vertex $v \in V$ has $\deg_\Gamma(v) + |v \cap \partial F| \geqslant 2$ preimages in a, where $\deg_\Gamma(v)$ is the degree of v in Γ. By pulling back the graph structure of Γ via r, we obtain a graph structure on a, with exactly $2|E|$ edges and

$$\sum_{v \in V} (\deg_\Gamma(v) + |v \cap \partial F|) \geqslant 2 \max\{|V|, |E|\}$$

vertices. It is then easy to prove that $\chi(A) = \chi(\Gamma) \leqslant \chi(a)$; inequality (2.7) now follows from (2.8).

We can then compute

$$\begin{aligned}
\mathrm{area}(G) &= \mathrm{area}(F_0) + \sum_{C \subseteq F_1} \mathrm{area}(G_C) \\
&\leqslant \mathrm{area}(F_0) + \sum_{C \subseteq F_1} (2\ell(C \cap F_0) - 3\chi(\mathrm{ab}(C)) + 4) && \text{by (2.4)} \\
&= \mathrm{area}(F_0) + 2\ell(F_1 \cap F_0) - 3\chi(\mathrm{ab}(F_1)) + 4|F_1| \\
&\leqslant \mathrm{area}(F_0) + 6\ell(F_1 \cap F_0) - 3\chi(F) + 3\chi(\mathrm{ab}(F_0)) && \text{by (2.5) and (2.7)} \\
&\leqslant \mathrm{area}(F_0) + 6\ell(\partial F_0) - 3\chi(F) + 3\,\mathrm{area}(F_0) && \text{by (2.6)} \\
&\leqslant 22\,\mathrm{area}(F_0) - 3\chi(F) && \text{since } \ell(\partial F_0) \leqslant 3\,\mathrm{area}(F_0),
\end{aligned}$$

where the sums range over all components C of F_1. ▢

Thanks to Proposition 2.19, we can now translate any polynomial-time algorithm for normal curves on triangulated surfaces to a polynomial-time algorithm for normal curves on normal surfaces. Consider, for instance, [21, Theorem 1.2], which gives an algorithm to decide whether two normal 1-manifolds in a triangulated surface are isotopic. A straightforward application of Proposition 2.19 yields the following result.

Proposition 2.20 (Deciding isotopy of normal curves in normal surfaces)**.** *There is an algorithm that takes as input*

- *a triangulation of a compact orientable* 3*-manifold* M *with* t *tetrahedra,*
- *a connected orientable normal surface* F *in* M*, and*
- *two essential multicurves* a, b *in* F*,*

and decides whether a *and* b *are isotopic in* F*. The running time of the algorithm is polynomial in* t*,* $|\chi(F)|$*,* $\log w(F)$*,* $\mathrm{area}(\mathrm{supp}(a))$*,* $\mathrm{area}(\mathrm{supp}(b))$*,* $\log w(a)$*, and* $\log w(b)$*.*

Proof. Let $F_0 = \mathrm{supp}(a) \cup \mathrm{supp}(b)$. We can apply the algorithm of Proposition 2.19 to obtain a triangulated surface G, a sub-2-complex G_0 of G, and a simplicial isomorphism $f : F_0 \to G_0$ that extends to a homeomorphism $F \to G$. The algorithm of [21, Theorem 1.2] can then be used to determine whether $f(a)$ and $f(b)$ are isotopic in G; this is clearly equivalent to a and b being isotopic in F. This algorithm runs in polynomial time in $\mathrm{area}(G)$, $\log w(f(a))$, and $\log w(f(b))$. Since $\mathrm{area}(G)$ is linear in $|\chi(F)|$ and $\mathrm{area}(F_0)$, and $w(f(a)) = w(a)$ and $w(f(b)) = w(b)$, we see that the combination of these two steps can be performed with the polynomial running time declared in the statement. ▢

The same exact argument applied to [21, Theorem 1.1] – an algorithm to compute the intersection number of normal 1-manifolds in triangulated surfaces – yields the following result.

Proposition 2.21. *There is an algorithm that takes as input*

- *a triangulation of a compact orientable* 3*-manifold* M *with* t *tetrahedra,*
- *a connected orientable normal surface* F *in* M*, and*
- *two normal multicurves* a, b *in* F*,*

and outputs the intersection number $i(a, b)$*. The running time of the algorithm is polynomial in* t*,* $|\chi(F)|$*,* $\log w(F)$*,* $\mathrm{area}(\mathrm{supp}(a))$*,* $\mathrm{area}(\mathrm{supp}(b))$*,* $\log w(a)$*, and* $\log w(b)$*.*

Chapter 3

Certifying fibredness

3.1 Outline of the certificate

The aim of this chapter is to describe a certificate that can be used to verify that a given triangulated compact orientable 3-manifold M is fibred, in polynomial time in the size of the triangulation. In other words, we will prove that the problem of deciding whether a compact orientable 3-manifold is fibred is in NP. This is not an entirely new result: in [35, Corollary 1.4], Schleimer proves that fibredness of orientable irreducible atoroidal 3-manifolds can be certified in polynomial time. Despite this precedent in the literature, we will give a new proof of this fact. We do this for two reasons. The first is that our proof avoids the restriction to irreducible atoroidal 3-manifolds, and generalises beyond Schleimer's result to all fibred 3-manifolds. This is not directly relevant for our final goal – that is, certifying hyperbolicity – since all hyperbolic 3-manifolds are atoroidal, but it fills a gap in the literature, and it is therefore worth including. The second and most important reason is that our certificate is specifically designed to facilitate the recovery of the monodromy of the fibration, as we will see in Chapter 4. This is a crucial feature of our certificate, since our plan to verify hyperbolicity of a fibred 3-manifold is to check that its monodromy is pseudo-Anosov. The formal description of the certificate, given in Certificate 1, is very technical and involved; we aim to make it less overwhelming, by providing a high-level overview of the construction in the following paragraphs.

Suppose we are given a triangulated compact orientable 3-manifold M with t tetrahedra, that is guaranteed to be fibred. The first piece of information that we include in the certificate is, quite naturally, a normal fibre F of M. To keep the certificate polynomial in size, we must ensure that the logarithm of the weight of F is polynomial in t; the existence of such a fibre is proved in Propositions 3.13 and 3.16. One could imagine that simply providing the fibre would be sufficient to certify fibredness of M. This is not completely unreasonable, although the burden is then shifted to the verifier, who must check that the given normal surface F is indeed a fibre of M. In fact, the verifier could apply [22, Theorem 9.3] to cut M along F, obtaining a pre-sutured manifold $M' = M \setminus\setminus F$, and then check that M' is a product interval bundle. We do not know whether this last step be performed in polynomial time; however, it can be *certified* in polynomial time, thanks to [22, Theorem 12.1] (see the statement in Theorem 3.1). Therefore, a certificate consisting of a normal fibre F and an auxiliary certificate (as provided by [22, Theorem 12.1]) showing that $M \setminus\setminus F$ is a product interval bundle would prove that M is fibred, and could be verified in polynomial time.

This naive approach, however, has two drawbacks. The first one is that the size of this certificate would be polynomial in t *and* in the Euler characteristic of F, which itself is not in general polynomial in t. This dependence cannot be avoided, since the minimum number of tetrahedra needed to triangulate $M \setminus\setminus F$ is linear in $|\chi(F)|$. This would not obstruct our final goal, since the size of our certificate for hyperbolicity will depend on $|\chi(F)|$ anyway, but it would prevent us from proving that deciding fibredness is in NP. The second and more serious drawback is that there is no natural way to read off the monodromy of the fibration from this certificate: once we cut M along F, all the information about the monodromy is lost. In other words, the sutured manifold M' does not carry any information about how $\partial_0 M'$ is glued to $\partial_1 M'$.

There is a way to overcome these two drawbacks, at the cost of a more complicated certificate. For the sake of simplicity, suppose that the guts X and the parallelity bundle Y of M' are both submanifolds of M' (in the general case, we would be working with thin(X) and thick(Y) instead). Suppose, moreover, that M is closed; this assumption is not necessary, but it simplifies the exposition. One could optimistically imagine a situation in which the annuli $X \cap Y$ are all vertical in M'. This would imply that Y would be a product interval bundle, whose interval bundle structure agrees with that of M'. The guts X would also be homeomorphic as a pre-sutured manifold to a product interval bundle, whose interval bundle structure again agrees with that of M'. In this case, certifying that M is fibred – or, equivalently, that the pre-sutured manifold M' is a product interval bundle – would be exceedingly easy: one simply needs to provide a proof that X (with its pre-sutured manifold structure) is a product interval bundle, which can be done thanks to [22, Theorem 12.1] (again, see the statement in Theorem 3.1). The key observation to keep in mind here is that, even though the parallelity bundle Y could be huge, we do not need to deal with it, since it is guaranteed to be an interval bundle; conversely, we have no guarantee on the guts X, but the size of the triangulation of X is polynomial in t (see Proposition 2.10).

Unfortunately, the annuli $X \cap Y$ will not in general be vertical in M'. However, this optimistic situation forms the core of our certificate. One can show, as we do in Proposition 3.8, that if the fibre F is chosen to be least-weight, then each annulus in $X \cap Y$ is either vertical, or it bounds a "tube" in M'; the second case is depicted in Figures 3.5a and 3.5b. This means that we can decompose M' into three submanifolds:

- the union Y_v of the parallelity components that are vertical in M';
- the remaining parallelity components Y_t, that are contained in a union of tubes T;
- the guts X.

The reader can find a graphical representation this decomposition in Figure 3.4, where Y_v is denoted by $\mathcal{N}^+(F_p)$ and Y_t is denoted by $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$. The reason for the different notation is that working with subsurfaces of F rather than submanifolds of M' is more convenient for algorithmic purposes. The relevant subsurfaces are F_g, F_p, F_0, and D_0, representing the traces on $\partial_0 M' \cong F$ of X, Y_v, Y_t, and T respectively; the subsurfaces F'_g, F'_p, F_1, and D_1 play the same role for $\partial_1 M' \cong F$.

In order to show that M is fibred, we will then provide a certificate consisting of a least-weight fibre F and a decomposition of M' into the three submanifolds X, Y_v, and Y_t (as explained above, this decomposition is actually given in terms of subsurfaces of F). The submanifold Y_t could be very large, but – since it is a product interval bundle over planar surfaces – it can be retriangulated with a number of tetrahedra that is polynomial in t, as shown in Proposition 3.2. By attaching this triangulation to the standard triangulation of the guts X, we obtain a triangulated 3-manifold N that is homeomorphic to $X \cup Y_t$, such that the size of the triangulation of N is polynomial in t (recall that the triangulation of the guts is "small"). Like in the above simplified case, the parallelity components Y_v could be huge, but we do not need to deal with them. All the verifier needs to do is check that the 3-manifold N (with a suitable pre-sutured manifold structure) is a product interval bundle, which again can be done thanks to [22, Theorem 12.1]. This will guarantee that N and Y_v glue along their common boundary to give a product interval bundle structure to M'.

To (loosely) summarise, not all parallelity components are guaranteed to be vertical in M'. However, the ones that are not are contained in tubes. We retriangulate these tubes with few tetrahedra and attach them to the guts to obtain a 3-manifold $N \cong X \cup Y_t$. The triangulation of N is given as part of the certificate, together with an auxiliary certificate (provided by [22, Theorem 12.1]) showing that N is a product interval bundle. The verifier can check that N is indeed a product interval bundle, and this guarantees that $M' = (X \cup Y_t) \cup Y_v$ is also a product interval bundle – or, in other words, that F is a fibre of M. This certificate can also be used to recover the monodromy of the fibration; however, we defer the explanation to Section 4.1.

3.2 Topology and combinatorics of interval bundles

In this section, we aim to address a few questions about interval bundles that are algorithmic in nature, concerning the recognition of interval bundles, triangulations of interval bundles, and the detection of vertical surfaces in interval bundles. Let us clarify what we mean by "vertical surfaces". Let $M = F \times [0,1]$ for some compact orientable surface F; we will only be dealing with interval bundles of this form. A surface properly embedded in M is *vertical* if it can be isotoped preserving $\partial F \times \{0,1\}$ to a union of fibres of the form $\{x\} \times [0,1]$. In fact, the same definition applies to the more general setting of an arbitrary subset of M, that is said to be *vertical* if it can be isotoped preserving $\partial F \times \{0,1\}$ to a union of interval fibres. It is easy to see that every vertical surface in M is either an annulus or a *square*, by which we mean a disc A properly embedded in M such that ∂A intersects $\partial F \times \{0,1\}$ transversely in 4 points.

3.2.1 Recognition of interval bundles

The most basic algorithmic problem about interval bundles asks whether a given triangulated 3-manifold is a product interval bundle over some compact orientable surface. However, our 3-manifolds often come equipped with information about where the horizontal and vertical boundary of the candidate interval bundle should be – in other words, with a pre-sutured triangulation. The question we are interested in is then the following.

Problem (SUTURED INTERVAL BUNDLE RECOGNITION).

Input: a pre-sutured triangulation $\mathcal{T}$ of a compact 3-manifold.

Output: whether $\mathcal{T}$ is a suitable pre-sutured triangulation of $F \times [0,1]$ for some compact orientable surface F.

The size of the input is measured by the number of tetrahedra in $\mathcal{T}$. $\times$

It turns out that a positive answer to this question can be certified in polynomial time.

Theorem 3.1 (Sutured interval bundle recognition is in NP)**.** *The problem Sutured interval bundle recognition is in* NP*.*

This result appears as Theorem 12.1 in [22], stated in greater generality for handle structures instead of pre-sutured triangulations. In fact, to be precise, since we are given as input a pre-sutured triangulation instead of a *sutured triangulation*, a pre-processing step is required, to verify that the given pre-sutured triangulation is indeed a sutured triangulation. In practice, for a given pre-sutured triangulation of a 3-manifold M, one simply needs to check that $\partial_v M$ is a union of annuli, each of which intersects both $\partial_0 M$ and $\partial_1 M$; this can be easily carried out in polynomial time in the size of the triangulation. More crucially, the reader should be aware that the published proof of [22, Theorem 12.1] contains a gap, in that it only works for sutured triangulations with non-empty sutures – in other words, when the surface F has non-empty boundary. However, this issue can easily be fixed by cutting the 3-manifold along a vertical normal annulus with bounded weight, which can be provided as part of the certificate; this is discussed at the beginning of [24, Section 11].

3.2.2 Triangulations of interval bundles

It should come as no surprise that the minimum number of tetrahedra needed to triangulate an interval bundle over a planar surface F is linear in the number of boundary components of F. For our fibredness certificate, we will need a more refined version of this estimate, that asks how many tetrahedra are needed to extend a given triangulation of a vertical subset of $\partial F \times [0,1]$ to a triangulation of $F \times [0,1]$. The answer turns out to be linear in the number of triangles of the prescribed triangulation, provided that the vertical subset touches every component of $\partial F \times [0,1]$.

Proposition 3.2 (Triangulating interval bundles over planar surfaces efficiently)**.** *Let $M = F \times [0,1]$ for some compact orientable planar surface F. Let $\Omega \subseteq \partial F$ be a union of finitely many circles or closed intervals with non-empty interiors; suppose that Ω intersects every component of ∂F. Let $\mathcal{T}_0$ be a triangulation of $\Omega \times [0,1] \subseteq \partial_v M$ with t_0 triangles, such that every point in $\partial\Omega \times \{0,1\}$ is a vertex of $\mathcal{T}_0$. Then there is a triangulation $\mathcal{T}$ of M such that:*

1. *the restriction of $\mathcal{T}$ to $\Omega \times [0,1]$ is $\mathcal{T}_0$;*

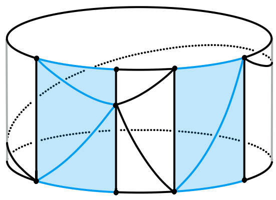

Figure 3.1. The triangulation $\mathcal{T}_1$ restricted to $b_i \times [0,1]$. The edges and triangles of $\mathcal{T}_1$ that are not in $\mathcal{T}_0$ are highlighted.

2. *the restriction of $\mathcal{T}$ to $\partial_h M$ is flapless;*

3. *the number of tetrahedra of $\mathcal{T}$ is at most $26t_0$.*

Proof. It is not restrictive to assume that M is connected. We start by finding a triangulation $\mathcal{T}_1$ of $\partial_v M$ that agrees with $\mathcal{T}_0$ on $\Omega \times [0,1]$. Denote by $b_1, \dots, b_m$ the components of ∂F; for each $1 \leqslant i \leqslant m$, we construct a triangulation of $b_i \times [0,1]$. Let $\Omega_i = \Omega \cap b_i$, and let s_i be the number of triangles of $\mathcal{T}_0$ contained in $\Omega_i \times [0,1]$. If $b_i \subseteq \Omega$, then we simply triangulate $b_1 \times [0,1]$ by restricting $\mathcal{T}_0$. Otherwise, denote by k the number of components of Ω_i, and by h the number of edges of $\mathcal{T}_0$ that are contained in $\partial\Omega_i \times [0,1]$; a simple counting argument shows that $s_i \geqslant h$. The closure of the complement of $\Omega_i \times [0,1]$ in $b_i \times [0,1]$ is a union of topological discs; each of these discs has a natural cell structure of its boundary, with one edge on $b_i \times \{0\}$, one on $b_i \times \{1\}$, and the remaining edges on $\partial\Omega_i \times [0,1]$ coming from $\mathcal{T}_0$. The number of these discs is k, and the combined simplicial length of their perimeters is exactly $2k + h$. Therefore, their union can be triangulated with exactly h triangles, as shown in Figure 3.1.

In conclusion, we have constructed a triangulation of $b_i \times [0,1]$ with at most $2s_i$ triangles (recall that $h \leqslant s_i$) that agrees with $\mathcal{T}_0$ on $\Omega_i \times [0,1]$. Repeating this construction for every boundary component b_i yields a triangulation $\mathcal{T}_1$ of $\partial_v M$ that agrees with $\mathcal{T}_0$ on $\Omega \times [0,1]$ and has

$$t_1 \leqslant 2s_1 + \cdots + 2s_m = 2t_0 \qquad (3.1)$$

triangles.

The next step is to find a triangulation of a union S of squares in M that agrees with $\mathcal{T}_1$ on $S \cap \partial_v M$, such that $M \setminus\setminus S$ is a 3-ball. To this aim, for each $1 \leqslant i \leqslant m$, fix a simplicial arc $a_i \subseteq b_i \times [0,1]$ that is vertical in M. Denote by ℓ_i the simplicial length of a_i, and – without loss of generality – assume that $\ell_1 \leqslant \ell_i$ for every $2 \leqslant i \leqslant m$. Find squares $S_2, \dots, S_m$ in M such that $S_i \cap \partial_v M = a_1 \cup a_i$ for each $2 \leqslant i \leqslant m$, and

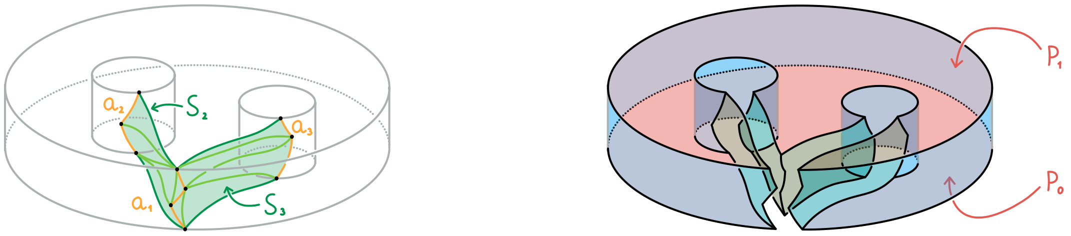

(a) (b)

Figure 3.2. (a) For each $2 \leqslant i \leqslant m$, the square S_i can be triangulated with $\ell_1 + \ell_i$ triangles. (b) The topological 3-ball $B = M \setminus\!\setminus S$ has a natural cell structure, consisting of triangles coming from $\mathcal{T}_1$ and (two copies of) $\mathcal{T}_2$, and two polygons P_0 and P_1.

$S_i \cap S_j = a_1$ for all $1 \leqslant i < j \leqslant m$. The boundary of each S_i has a natural cell structure with exactly $\ell_1 + \ell_i + 2$ edges; therefore, we can triangulate S_i with $\ell_1 + \ell_i$ triangles. Let $\mathcal{T}_2$ be the triangulation of $S = S_2 \cup \cdots \cup S_m$ obtained from this construction, as shown in Figure 3.2a. The number t_2 of triangles of $\mathcal{T}_2$ is exactly

$$t_2 = (m-1)\ell_1 + \ell_2 + \cdots + \ell_m. \tag{3.2}$$

Finally, we construct a triangulation of M, by exploiting the fact that $B = M \setminus\!\setminus S$ is a topological 3-ball. The triangulations $\mathcal{T}_1$ and $\mathcal{T}_2$ induce a cell structure on the boundary of B, consisting of $t_1 + 2t_2$ triangles and two polygons P_0 and P_1, where $P_i = B \cap \partial_i M$ for $i \in \{0, 1\}$; this is depicted in Figure 3.2b. Note that, for each $i \in \{0, 1\}$, the number of edges of P_i is exactly $r_i + 2m - 2$, where r_i is the number of edges of $\mathcal{T}_1$ that are contained in $\partial_i M$. We can then triangulate each P_i with $r_i + 2m - 2$ triangles, by coning over a point in the interior of P_i; this construction guarantees that the triangulation of each P_i is flapless, and it produces a triangulation of ∂B with exactly

$$t_1 + 2t_2 + r_0 + r_1 + 4m - 4$$

triangles. By coning over a vertex of B, we obtain a triangulation $\mathcal{T}_3$ of B whose number t_3 of tetrahedra is at most

$$t_3 \leqslant t_1 + 2t_2 + r_0 + r_1 + 4m - 5. \tag{3.3}$$

We conclude by gluing this triangulation of B along S to obtain a triangulation $\mathcal{T}$ of M.

By construction, the restriction of $\mathcal{T}$ to $\Omega \times [0,1]$ is exactly $\mathcal{T}_0$, and the restriction of $\mathcal{T}$ to $\partial_h M$ is flapless, so all that is left is bounding the number of tetrahedra of $\mathcal{T}$ – or, equivalently, of $\mathcal{T}_3$. Let s be the number of edges of $\mathcal{T}_1$; note that, trivially, we have the bound $s \leqslant 3t_1$. We also remark that, equally trivially, we have that $r_i \geqslant m$ for $i \in \{0,1\}$. Elementary counting arguments show that

$$\ell_1 + \cdots + \ell_m + r_0 + r_1 \leqslant s \leqslant 3t_1, \tag{3.4}$$

$$m\ell_1 \leqslant \ell_1 + \cdots + \ell_m \leqslant 3t_1 - 2m. \tag{3.5}$$

We can finally bound the number of tetrahedra of $\mathcal{T}_3$ as follows:

$$\begin{aligned}
t_3 &\leqslant t_1 + 2t_2 + r_0 + r_1 + 4m - 5 && \text{by (3.3)}\\
&= t_1 + 2(m-1)\ell_1 + 2\ell_2 + \cdots + 2\ell_m + r_0 + r_1 + 4m - 5 && \text{by (3.2)}\\
&\leqslant 4t_1 + (2m-3)\ell_1 + \ell_2 + \cdots + \ell_m + 4m - 5 && \text{by (3.4)}\\
&\leqslant 7t_1 + 2(m-2)\ell_1 + 2m - 5 && \text{by (3.5)}\\
&\leqslant 13t_1 - 4\ell_1 - 2m - 5 && \text{by (3.5)}\\
&< 13t_1 && \text{since } \ell_1 \geqslant 0 \text{ and } m \geqslant 0\\
&\leqslant 26t_0 && \text{by (3.1).} \quad \square
\end{aligned}$$

3.2.3 Detection of vertical surfaces in interval bundles

In order to give an algorithm that detects (essential) vertical surfaces, we need the two following propositions. The first one (Proposition 3.3) states that incompressible boundary-incompressible surfaces in interval bundles are vertical, provided that they satisfy some additional conditions. The second one (Proposition 3.4) states that squares and annuli in interval bundles are incompressible and boundary-incompressible, provided that they satisfy some additional conditions. Together, these two propositions give a list of simple criteria for deciding whether a surface in an interval bundle is vertical and "essential", in the sense that it is admissibly isotopic to $a \times [0,1]$ for some 1-manifold $a \subseteq F$ without "trivial" components.

Proposition 3.3 (Essential surfaces in interval bundles are vertical)**.** *Let $M = F \times [0,1]$ for some compact orientable surface F, and consider the boundary pattern $\Gamma = \partial F \times \{0,1\}$. Let A be a surface properly embedded in (M,Γ) that is incompressible and boundary-incompressible. Suppose that no component of A is a surface without boundary* *or a disc intersecting* Γ *at most twice, and that no component of* ∂A *is contained in* $\partial_v M$*. Then* A *is vertical in* M*.*

Proof. We can assume without loss of generality that A is connected, and moreover that M is connected. If any component of $A \cap \partial_v M$ is an arc with both endpoints on $\partial_0 M$ or $\partial_1 M$, then we immediately see that A is boundary-compressible. Therefore, every component of $A \cap \partial_v M$ is a vertical arc.

We now wish to invoke [18, Proposition 5.6] to conclude that A must be vertical. To this aim, we must first translate the definition of "essential" from [18, Definition 3.1] to our setting. We say that a surface S properly embedded in (M, Γ) is *essential in the sense of [18]* if

- S is incompressible in M, and
- for any (not necessarily clean) boundary-compressing disc D for S in M that intersects Γ at most once, the arc $D \cap S \subseteq S$ cuts a disc off of S that intersects Γ at most once.

Then [18, Proposition 5.6] implies the following: if F is not a sphere nor a disc, then any surface S properly embedded in (M, Γ) that is essential in the sense of [18] is horizontal or vertical in M, provided that no component of S is a sphere or a disc intersecting Γ at most three times; in fact, "three" can be replaced by "two" here, since S will necessarily intersect Γ an even number of times. Note that, if F were a sphere or a disc, then A would necessarily be compressible or boundary-compressible. Moreover, the surface A cannot be horizontal, because ∂A intersects $\partial_h M$ by assumption.

Therefore, all we need to show is that A is essential in the sense of [18]. We already know that A is incompressible in M. Let D be a boundary-compressing disc for A in M that intersects Γ at most once. If D is clean, then by boundary-incompressibility of A we know that $D \cap A$ cuts a clean disc off of A. Otherwise, the boundary of D is naturally split into three non-empty arcs: let $a = D \cap A$, $b_v = D \cap \partial_v M$, and $b_0 = D \cap \partial_h M$. Without loss of generality, assume that $b_0 \subseteq \partial_0 M$. Let c be the subarc of $A \cap \partial_v M$ that connects an endpoint of b_v to $\partial_0 M$; the existence of such a subarc is guaranteed by our previous observation that every component of $A \cap \partial_v M$ is a vertical arc. This situation is depicted in Figure 3.3a.

The arc $b_v \cup c$ cuts a disc D_v off of $\partial_v M$. The disc $D \cup D_v$ can be isotoped slightly off of $\partial_v M$ to a clean boundary-compressing disc D' for A. Since A is boundary-incompressible, we deduce that the arc $D' \cap A$ cuts a clean disc off of A. It is then easy to conclude (see

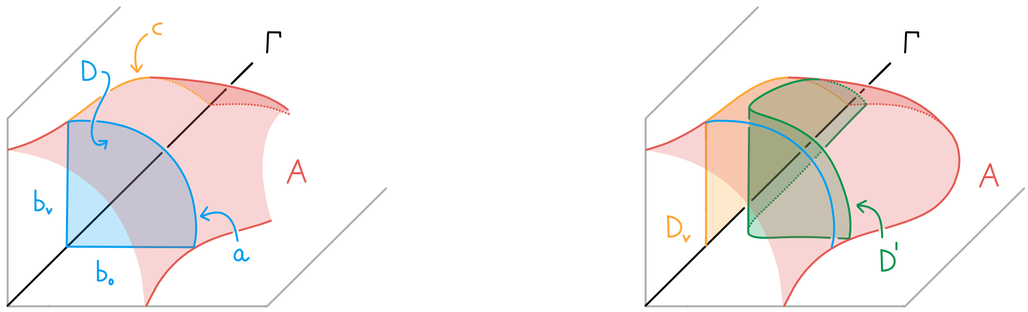

(a) (b)

Figure 3.3. (a) The boundary-compressing disc D for A intersects Γ at most once. (b) The disc $D \cup D_v$ can be isotoped to a clean boundary-compressing disc D' for A.

Figure 3.3b) that a cuts a disc off of A that intersects Γ exactly once – in particular, the point of intersection is $\partial c_v \setminus b_v$. This concludes the proof, since A is essential in the sense of [18] and therefore vertical in M. ▢

Proposition 3.4 (Vertical annuli and squares in interval bundles). *Let $M = F \times [0,1]$ for some compact orientable surface F, and consider the boundary pattern $\Gamma = \partial F \times \{0,1\}$. Let A be a surface properly embedded in (M,Γ). Suppose that each component of A is either a square intersecting both $\partial_0 M$ and $\partial_1 M$, or an annulus with one boundary component contained in $\partial_0 M$ and the other in $\partial_1 M$. Additionally, suppose that no component of $A \cap \partial_0 M$ is a curve bounding a disc in $\partial_0 M$ or a boundary-parallel arc in $\partial_0 M$. Then A is incompressible and boundary-incompressible in (M,Γ). Moreover, the surface A is vertical in M.*

Proof. We can assume without loss of generality that A is connected. If A is an annulus, then $A \cap \partial_0 M$ represents a non-trivial element in $\pi_1 F$. Since the projection of M onto F is a homotopy equivalence, we deduce that A is π_1-injective in M, and hence incompressible. To prove boundary-incompressibility, note that any arc in A with two endpoints on the same component of ∂A bounds a clean disc in A. On the other hand, if an arc in A has one endpoint in $F \times \{0\}$ and one endpoint in $F \times \{1\}$, then the two endpoints cannot be connected by a clean arc in ∂M. We conclude that A is boundary-incompressible in (M,Γ).

If A is a square, then it is incompressible since it is simply connected. Let D be a clean compressing disc for A, and let $b = D \cap A$. If b has both endpoints on the same component of $\partial A \setminus \Gamma$, then it bounds a clean disc in A. Otherwise, the only way in which the two endpoints of b can lie on the same component of $\partial M \setminus \Gamma$ is if b connects the two components of $\partial A \cap \partial_v M$. Let D' be the union of D with the closure of the component of $A \setminus b$ that intersects $\partial_0 M$. Then the projection of D' to $\partial_0 M$ gives a homotopy of $A \cap \partial_0 M$ into $\partial\partial_0 M$ fixing $\partial(A \cap \partial_0 M)$, contradicting the assumption that $A \cap \partial_0 M$ is not a boundary-parallel arc. Hence, we conclude that A is boundary-incompressible in (M, Γ).

In both cases, verticality of A follows from Proposition 3.3. ▢

As anticipated, the two propositions above combine to give an algorithm to decide whether a normal surface A in a triangulated interval bundle M is vertical and essential. The running time of the algorithm will be polynomial in the size of the triangulation of M and the logarithm of the weight of A.

Proposition 3.5 (Detecting essential vertical surfaces in interval bundles)**.** *There is an algorithm that takes as input a suitable pre-sutured triangulation of a compact orientable 3-manifold M and a normal surface A in M, and decides whether the two following conditions are simultaneously satisfied:*

(i) the surface A is vertical in M;

(ii) no component of A intersects $\partial_0 M$ in a curve bounding a disc in $\partial_0 M$ or a boundary-parallel arc in $\partial_0 M$.

The running time of the algorithm is polynomial in the size of the triangulation of M and the logarithm of the weight of A.

Proof. Given M and A, we claim that the output is "yes" if and only if the following two conditions hold:

- each component of A is either a square intersecting both $\partial_0 M$ and $\partial_1 M$, or an annulus with one boundary component contained in $\partial_0 M$ and the other in $\partial_1 M$;
- no component of $A \cap \partial_0 M$ is a curve bounding a disc in $\partial_0 M$ or a boundary-parallel arc in $\partial_0 M$.

If the two above conditions hold, then the output is "yes" by Proposition 3.4. Conversely, if A is vertical then the first condition above clearly holds. Therefore, the two conditions above are equivalent to those in the statement, and we can focus on checking them.

We start by running the algorithm of Proposition 2.15 to obtain a list of the (normal isotopy classes of the) components of A, together with their topological types. By checking on one component at a time, we can assume without loss of generality that A is connected. If A is not a disc or an annulus, then we output "no". If A is an annulus, then we check that it is disjoint from $\partial_v M$; if A is a disc, we can verify that it is a square by counting how many times it intersects $\partial\partial_h M$. Either way, we then check that A intersects both $\partial_0 M$ and $\partial_1 M$; this guarantees that the first condition above is satisfied. Finally, we can run the algorithm of Proposition 2.14 applied to the connected normal 1-manifold $A \cap \partial_0 M$ in $\partial_0 M$ to verify that the second condition above also holds. ▢

3.3 A certificate for fibredness

We now describe in detail our certificate for fibredness. We split our discussion into four parts, establishing a pattern that we will follow for the other main certificates in this thesis – namely, the monodromy certificate in Section 4.3 and the hyperbolicity certificates in Section 5.5.

- First, we define the certificate.
- Then, we prove the *existence* of the certificate: every fibred 3-manifold admits a valid certificate.
- Then, we prove the *correctness* of the certificate: if a 3-manifold admits a valid certificate, then it is fibred.
- Finally, we address the *verification* of the certificate: we show how to decide whether a given certificate is valid or not in polynomial time.

This certificate is quite long and really technical; we strongly encourage the reader to follow Section 3.1 for a more informal overview of the certificate, before diving into the details. We also recommend consulting Figure 3.4 while reading the definition and proofs of the certificate, as it illustrates the key components involved in the constructions.

Certificate 1 (Certificate for fibredness)**.** Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M with t tetrahedra, and let F be a transversely oriented connected normal surface in M. Let $M' = M \setminus\setminus F$, and denote by X the guts of M'. We say that a certificate Σ lies in $\mathfrak{S}_{\text{fib}}(M, F)$ if it consists of:

- sub-2-complexes F_0, F_1, D_0 of F;
- a pre-sutured triangulation $\mathcal{R}$ of a compact oriented 3-manifold N;

- a sub-2-complex D_1 of $\partial_1 N$;
- sub-3-complexes N' and N'' of N;
- sub-2-complexes G''_0 and G''_1 of $\partial_h N$;
- a simplicial map $f_g \colon N' \to X$,
- two additional certificates Σ' and Σ''.

We define the *size* of Σ to be the number

$$
\begin{aligned}
|\Sigma| = t + \log(w(F)+1) + |\mathcal{R}| + |\Sigma'| + |\Sigma''| \\
+ \ell(\partial_F F_0) + \ell(\partial_F F_1) + \ell(\partial_F D_0),
\end{aligned}
$$

where $|\Sigma'|$ and $|\Sigma''|$ are the sizes of the certificates Σ' and Σ'' respectively. We remark that the certificate Σ can be described with a number of binary digits that is polynomial in $|\Sigma|$.

Let F_g and F'_g be the sub-2-complexes of F such that $F_g^+ = \partial_0 X$ and $(F'_g)^- = \partial_1 X$. Denote by F_p the closure of $F \setminus (F_g \cup F_0)$ in F, and by F'_p the closure of $F \setminus (F'_g \cup F_1)$ in F. We say that Σ is *valid*, and write $\Sigma \in \mathfrak{S}^*_{\mathrm{fib}}(M, F)$, if the following conditions are satisfied:

(a.1) F_g and F_0 have disjoint interiors;

(a.2) F_0 and F_p are disjoint;

(a.3) F'_g and F_1 have disjoint interiors;

(a.4) F_1 and F'_p are disjoint;

(a.5) $\mathcal{N}^+(F_p) = \mathcal{N}^-(F'_p)$;

(b.1) $N \cong G \times [0,1]$ for some compact orientable surface G;

(b.2) $\mathcal{R}$ is a suitable pre-sutured triangulation of N;

(b.3) the restriction of $\mathcal{R}$ to $\partial_h N$ is flapless;

(b.4) $N = N' \cup N''$, and N' and N'' have disjoint interiors;

(b.5) G''_0 and G''_1 are disjoint;

(b.6) $G''_0 \cup G''_1 = \partial_h N \cap N''$;

(b.7) $\mathrm{ab}(N'') \cong G'' \times [0,1]$ for some compact orientable surface G'';

(b.8) $\mathcal{R}''$ is a suitable pre-sutured triangulation of N'', where $\mathcal{R}''$ is the pre-sutured triangulation of N'' that, as a triangulation, is simply the restriction of $\mathcal{R}$ to N'', and such that $\partial_i N'' = G_i''$ for $i \in \{0, 1\}$;

(c.1) f_g restricts to an orientation-preserving homeomorphism $\text{int}(N') \to \text{int}(X)$;

(c.2) $f_g(\partial_i N \cap N') = \partial_i X$ for $i \in \{0, 1\}$;

(c.3) $f_g(N' \cap N'') = (\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)) \cap X$;

(c.4) f_g restricts to a simplicial isomorphism

$$N' \cap \partial_h N'' \longrightarrow (\partial_F F_0)^+ \cup (\partial_F F_1)^-,$$

and this restriction extends to an orientation-preserving homeomorphism $\partial_h N'' \to F_0^+ \cup F_1^-$;

(d.1) each component of $\text{thick}(D_0)$ is a disc intersecting ∂F in a single (possibly empty) arc;

(d.2) $F_0 \subseteq D_0$;

(d.3) D_0 is disjoint from F_p;

(d.4) each component of $\text{thick}(D_1)$ is a disc intersecting $\partial\partial_1 N$ in a single (possibly empty) arc;

(d.5) $\partial_1 N \cap N'' \subseteq D_1$;

(d.6) $f_g(D_1 \cap N')$ is disjoint from F_p;

(e.1) $\max\{\ell(\partial_F F_0), \ell(\partial_F F_1), \ell(\partial_F D_0), \ell(\partial D_1)\} \leqslant 32t$;

(e.2) $\mathcal{R}$ has at most $936t$ tetrahedra;

(f.1) Σ' certifies that $\mathcal{R}$ is a suitable pre-sutured triangulation of N (according to the algorithm of Theorem 3.1), and its size is polynomial in $|\mathcal{R}|$;

(f.2) Σ'' certifies that the pull-back of $\mathcal{R}''$ under $\text{emb}_{N''}$ is a suitable pre-sutured triangulation of $\text{ab}(N'')$ (according to the algorithm of Theorem 3.1), and its size is polynomial in $|\mathcal{R}''|$. ×

Proposition 3.6 (Basic properties of Certificate 1)**.** *In the setting of Certificate 1, if* $\Sigma \in \mathfrak{S}^*_{\text{fib}}(M, F)$*, then the following hold:*

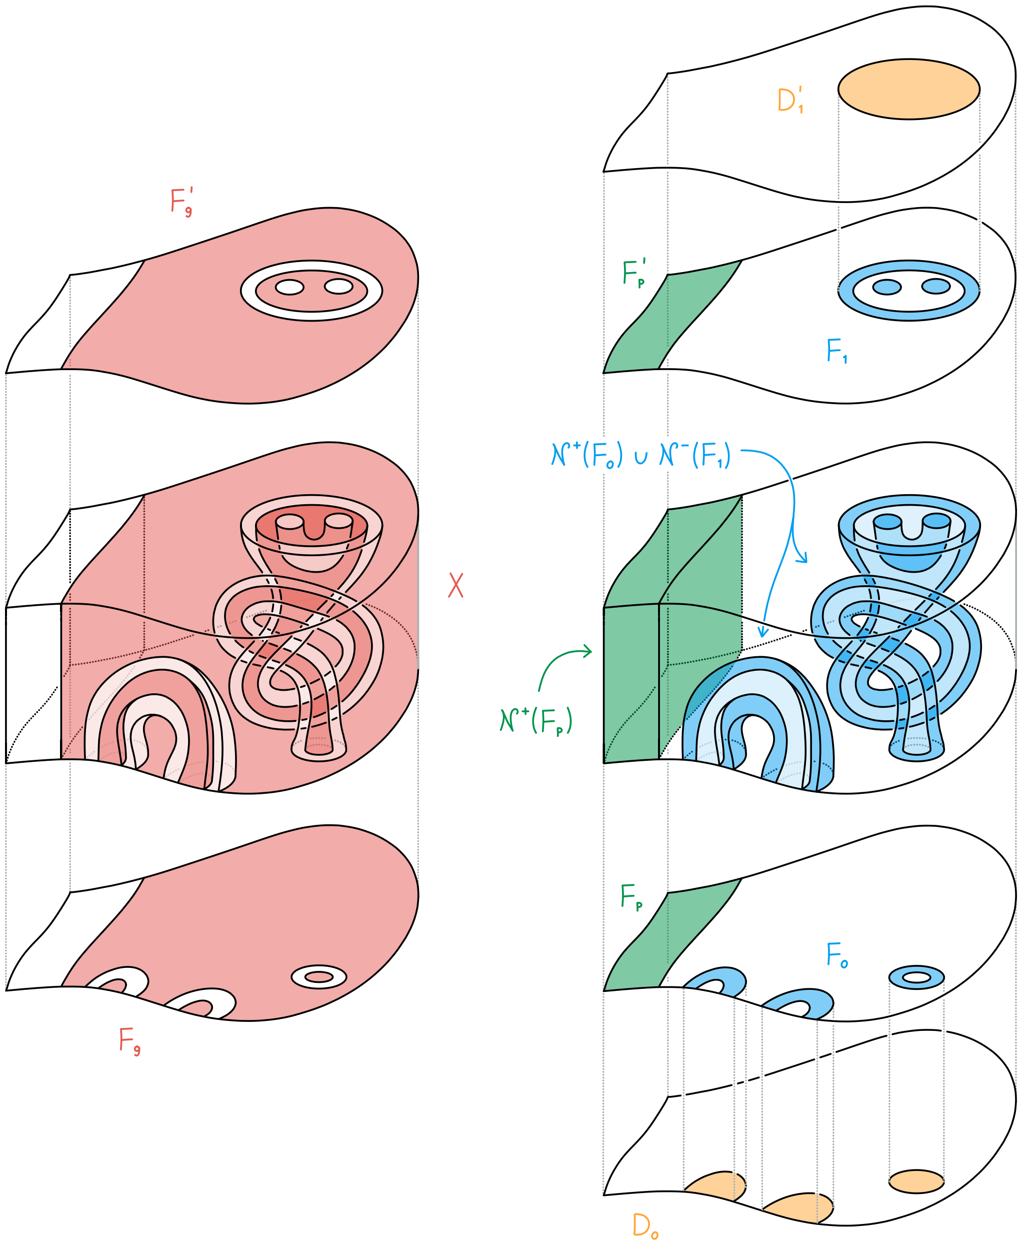

Figure 3.4. Some objects that are part of Certificate 1. The top and bottom portions of the figure represent sub-2-complexes of F, namely F_p, F_0, F_1, F_g, F'_p, F'_g, and D'_1. Note that D'_1 is not part of Certificate 1, but it is used in the proof of Proposition 3.9. The middle portion of the figure depicts a portion of M'. Alluding to the fact that, by Proposition 3.11, if the certificate is valid then F is a fibre of M, we have drawn M' as an interval bundle, with $\partial_0 M' = F^+$ at the bottom and $\partial_1 M' = F^-$ at the top. On the left, we show the guts X of M'. On the right, the parallelity bundle Y of M' is shown, partitioned into the sub-3-complexes $\mathcal{N}^+(F_p)$ and $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$.

1. *F_g, F'_g, F_p, F'_p, F_0, and F_1 are unions of normal discs;*

2. *$\mathcal{N}^+(F_p)$, $\mathcal{N}^+(F_0)$, and $\mathcal{N}^-(F_1)$ are unions of parallelity components of M';*

3. *the parallelity bundle of M' is the union of $\mathcal{N}^+(F_p)$ and $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$;*

4. *$(\mathcal{N}^+(F_0) \cup \mathcal{N}^+(F_1)) \cap \partial_h M' = F_0^+ \cup F_1^-$;*

5. *$\mathcal{N}^+(F_p)$ and $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$ are disjoint;*

6. *f_g restricts to a homeomorphism $f_g^{-1}(X') \to X'$, where X' is the complement in X of the singular locus of X;*

7. *$\mathrm{thick}(F_0)$ and $\mathrm{thick}(F_1)$ are planar surfaces;*

8. *$\mathrm{area}(F_g) \leqslant 32t$ and $\mathrm{area}(F'_g) \leqslant 32t$;*

9. *$\mathrm{area}(\partial_h N) \leqslant 2808t$.*

Proof. We prove each statement separately.

1. By definition, the surface F_g is a union of normal discs. Properties (a.1) and (a.2) imply that F_0 and F_p are unions of components of $\mathrm{clos}(F \setminus F_g)$. In particular, they must both be unions of normal discs. The same argument applies to F'_g, F'_p, and F_1.

2. This follows immediately from the fact that F_p and F_0 are unions of components of $\mathrm{clos}(F \setminus F_g)$, and similarly F_1 is a union of components of $\mathrm{clos}(F \setminus F'_g)$.

3. This is a consequence of the fact that $F = F_g \cup F_0 \cup F_p$ and $F = F'_g \cup F_1 \cup F'_p$.

4. We only prove the statement for $\partial_0 M'$. Let Z be a component of $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$. Each component C of $Z \cap \partial_0 M'$ is a component of F_0^+ or of F_p^+. Suppose, for a contradiction, that C is a component of F_p^+. Property (a.5) implies that $C = Z \cap \partial_0 M'$. But C must be disjoint from F_0^+ by property (a.2), and $Z \cap \partial_1 M' \subseteq (F'_p)^-$ must be disjoint from F_1^- by property (a.3). However, this contradicts the fact that Z is a component of $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$.

5. Two parallelity components intersect if and only if their horizontal boundaries intersect. However, property (a.5) implies that $\mathcal{N}^+(F_p) \cap \partial_h M' = F_p^+ \cup (F'_p)^-$. We readily see that the horizontal boundaries of $\mathcal{N}^+(F_p)$ and $\mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$ are disjoint.

6. If x is a point in $X \setminus X'$, then it admits a neighbourhood U in X such that $U \setminus \partial X$ is connected. Since, by property (c.1), the restriction of f_g to $\mathrm{int}(N')$ is a homeomorphism onto $\mathrm{int}(X)$, we see that $f_g^{-1}(U \setminus \partial X)$ is connected. Therefore, there can be at most one point in N' mapping to x under f_g.

7. Properties (d.1) and (d.2) immediately imply that $\mathrm{thick}(F_0)$ is a planar surface. To see that $\mathrm{thick}(F_1)$ is planar, note first that properties (c.1) and (c.4) combine to give a homeomorphism $\partial_1 N \to (F'_g \cup F_1)^-$ sending $\partial_1 N \cap N''$ to F_1^-. From properties (d.4) and (d.5), we conclude that $\mathrm{thick}(F_1)$ is contained in a disjoint union of discs, and hence it is planar.

8. The sub-2-complex F_g of F is simplicially isomorphic to $\partial_0 X$, hence by Proposition 2.10 we have that $\mathrm{area}(F_g) \leqslant 32t$. The argument for F'_g is completely analogous.

9. By property (e.2), there are at most $936t$ tetrahedra in the triangulation $\mathcal{R}$ of N. Any such tetrahedron has at most 3 faces in ∂N (note that $\mathcal{R}$ is a suitable pre-sutured triangulation of an interval bundle by property (b.2), hence no component of N can consist of a single tetrahedron). Therefore, we find that

$$\mathrm{area}(\partial_h N) \leqslant \mathrm{area}(\partial N) \leqslant 3 \cdot 936t = 2808t. \qquad \square$$

3.3.1 Existence

We prove that every least-weight normal fibre F of a triangulated compact connected orientable 3-manifold M admits a valid Certificate 1. As outlined in Section 3.1, the crucial step is to show that the components of the intersection between the guts X and the parallelity bundle Y of $M \setminus\setminus F$ – or, more precisely, between $\mathrm{thin}(X)$ and $\mathrm{thick}(Y)$ – are either vertical or contained in "tubes". The precise definition of "being contained in a tube" is given below.

Definition 3.7 (Cutting a $D^2 \times [0,1]$ off of an interval bundle)**.** Let $M = F \times [0,1]$ for some compact orientable surface F. We say that an annulus or square A properly embedded in M *cuts a* $D^2 \times [0,1]$ *off of* M if there is an embedding $f\colon D^2 \times [0,1] \to M$ such that

$$f(D^2 \times \{0,1\}) \subseteq \partial_h M \quad \text{and} \quad A \subseteq f(\partial D^2 \times [0,1]) \subseteq A \cup \partial_v M,$$

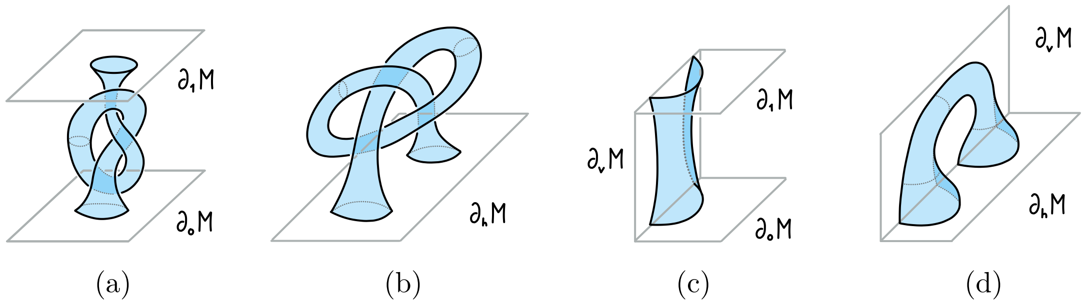

Figure 3.5. (a) An annulus cutting a (possibly knotted) $D^2 \times [0,1]$ off of M that connects $\partial_0 M$ and $\partial_1 M$. (b) An annulus cutting a (possibly knotted) $D^2 \times [0,1]$ off of M that connects $\partial_0 M$ (or $\partial_1 M$) to itself. (c) A square cutting a $D^2 \times [0,1]$ off of M that connects $\partial_0 M$ and $\partial_1 M$. (d) A square cutting a $D^2 \times [0,1]$ off of M that connects $\partial_0 M$ (or $\partial_1 M$) to itself.

and moreover $f^{-1}(A)$ is vertical in $D^2 \times [0,1]$. We refer to the image of f as a $D^2 \times [0,1]$ cut off of M by A. ×

There are essentially four ways in which an annulus or square can cut a $D^2 \times [0,1]$ off of M, as shown in Figure 3.5.

Proposition 3.8 (Parallelity bundle of a least-weight fibre). *Let $\mathcal{T}$ be a triangulation of a compact connected orientable 3-manifold M that fibres over the circle, and let F be a least-weight normal fibre of M. Denote by M' the 3-manifold $M \setminus\!\setminus F \cong F \times [0,1]$, by X the guts of M', and by Y the parallelity bundle of M'. Let A be a component of* $\mathrm{thin}(X) \cap \mathrm{thick}(Y)$. *Then A is vertical in M' or cuts a $D^2 \times [0,1]$ off of M'.*

Proof. Let $\Gamma = \partial\partial_h M'$. We analyse two cases, depending on whether A is an annulus or a square.

When A is an annulus. If A is not vertical, then by Proposition 3.3 it must be compressible or boundary-compressible in (M', Γ). Suppose that A is compressible in M', and let D be a non-trivial compressing disc for A. Compressing A along D yields two disjoint discs A_1 and A_2, with ∂A_1 and ∂A_2 contained in $\partial_h M'$. Since $\partial_h M'$ is incompressible in M', the boundary of A_1 bounds a disc D_1 in $\partial_h M'$; by irreducibility of M', we deduce that the sphere $A_1 \cup D_1$ bounds a 3-ball B_1 in M' (when F is a sphere, the 3-manifold M' is not irreducible, but this statement is still true up to replacing D_1 with the other disc in $\partial_h M'$ having the same boundary). Similarly, we find a 3-ball B_2 in

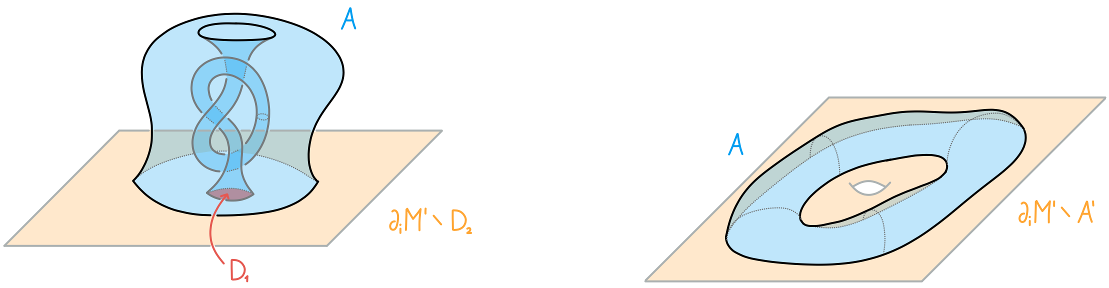

(a) (b)

Figure 3.6. (a) In a trivial annular simplification, the surface F' is the union of three pieces: the surface $\partial_i M' \setminus D_2$, the annulus A, and the disc D_1. (b) In an essential annular simplification, the surface F' is the union of two pieces: the surface $\partial_i M' \setminus A'$ and the annulus A.

M' whose boundary is the union of A_2 and a disc $D_2 \subseteq \partial_h M'$. If B_1 and B_2 are disjoint, then A cuts a $D^2 \times [0,1]$ off of M', namely $B_1 \cup B_2 \cup \mathcal{N}(D)$; this situation is depicted in Figure 3.5a or 3.5b. If, instead, one of the two 3-balls is contained in the other – say, $B_1 \subseteq B_2$, then we consider the surface

$$F' = (\partial_i M' \setminus D_2) \cup A \cup D_1,$$

where $\partial_i M'$ is the component of $\partial_h M'$ containing D_1 and D_2; see Figure 3.6a.

This construction is called *trivial annular simplification*, and is described in [23, §6.4]. Note that F' is isotopic to $\partial_i M'$ in M, and hence to F. The annulus A does not intersect the edges of $\mathcal{T}$; therefore, we see that

$$w(\partial_i M') - w(F') = |(D_2 \setminus D_1) \cap \mathcal{T}^{(1)}|,$$

where weights are taken with respect to $\mathcal{T}$. However, since $w(\partial_i M') = w(F)$, by minimality of F we deduce that $D_2 \setminus D_1$ must be disjoint from the 1-skeleton of $\mathcal{T}$, and moreover that $w(F') = w(F)$. Recall that A is a component of $\mathrm{thin}(X) \cap \mathrm{thick}(Y)$. If we let π be the natural projection of $\partial_v \,\mathrm{thick}(Y)$ onto $\partial_v Y$, induced by the retraction $\mathrm{thick}(Y) \to Y$, then $\pi(A)$ is a vertical subset of $\partial_v X \cap \partial_v Y$; in particular, it will contain a component e of the intersection $M' \cap \mathcal{T}^{(1)}$. Note that, since $D_2 \setminus D_1$ cannot intersect e, we necessarily have that e has an endpoint on D_1 and the other on $\partial_i M' \setminus D_2$. Let a be a vertical arc in A such that $\pi(a) = e$. There is an isotopy of M' that takes a to e and

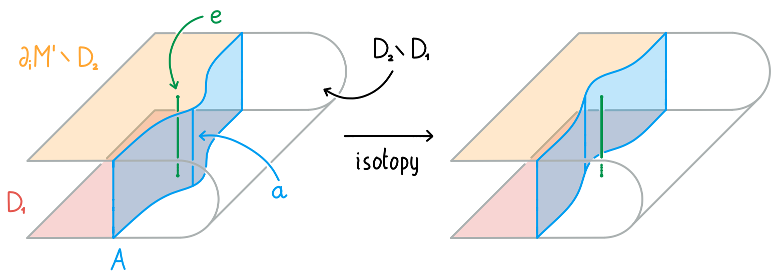

Figure 3.7. Dragging a to e and then past it produces an isotopy of M' that sends the surface $F' = (\partial_i M' \setminus D_2) \cup A \cup D_1$ to a surface F'' that does not intersect e and, hence, has weight strictly less than that of F'.

then pushes it slightly past e; this isotopy takes F' to a surface F'' such that

$$F'' \cap \mathcal{T}^{(1)} = (F' \cap \mathcal{T}^{(1)}) \setminus e.$$

See Figure 3.7 for a depiction of this isotopy. The surface F'' is isotopic to F in M, and has weight strictly less than that of F, contradicting the minimality of F.

Suppose now that A is incompressible and boundary-compressible in (M', Γ), and let D be a non-trivial boundary-compressing disc for A. Compressing A along D yields a disc A_1, with ∂A_1 contained in $\partial_h M'$. Like above, we find that A_1 cobounds a 3-ball B_1 with $\partial_h M'$. We see that D cannot be contained in B_1, since A is incompressible in M'. Therefore, the annulus A cobounds a solid torus (namely, $B_1 \cup \mathcal{N}(D)$) with an annulus A' in $\partial_h M'$. Denote by $\partial_i M'$ the component of $\partial_h M'$ containing A', and consider the surface

$$F' = (\partial_i M' \setminus A') \cup A;$$

see Figure 3.6b. This is an instance of *essential annular simplification* as described in [23, §6.4]. Note that F' is isotopic to $\partial_i M'$ in M, and hence to F. Moreover, the same argument as above shows that F' can be isotoped to a general position surface with weight strictly less than that of F, contradicting the minimality of F.

When A is a square. If A is not vertical, then by Proposition 3.3 it must be boundary-compressible in (M', Γ). Let D be a non-trivial boundary-compressing disc for A. Compressing A along D yields two disjoint discs A_1 and A_2, with ∂A_1 and ∂A_2 each intersecting Γ twice. It is easy to see that A_1 and A_2 cobound 3-balls – say, respectively, B_1 and B_2 – with $\partial M'$.

Suppose first that the boundary-compressing disc D intersects $\partial_h M'$, and is therefore disjoint from $\partial_v M'$. If one of the two 3-balls is contained in the other, say $B_1 \subseteq B_2$, then A cuts a $D^2 \times [0,1]$ off of M', namely $B_2 \setminus (B_1 \cup \mathcal{N}(D))$; this situation is depicted in Figure 3.5d. Otherwise, the two 3-balls are disjoint. Let B be the closure of the component of $M' \setminus A$ intersecting B_1 and B_2; then B is itself a 3-ball, and it is the union of B_1, B_2, and a neighbourhood of D. Let $A' = B \cap \partial_h M'$, and let $\partial_i M'$ be the component of $\partial_h M'$ containing A'. Consider the surface

$$F' = (\partial_i M' \setminus A') \cup A.$$

Like above, we can argue that F' is isotopic to F in M, and it can be isotoped to a general position surface with weight strictly less than that of F, contradicting the minimality of F.

Suppose now that the boundary-compressing disc D intersects $\partial_v M'$ – and is therefore disjoint from $\partial_h M'$. If $B_1 \subseteq B_2$ or $B_2 \subseteq B_1$, then in fact A also admits a boundary-compressing disc that intersects $\partial_h M'$, and we reduce to the previous case. Otherwise, the two 3-balls are disjoint, and it is then easy to see that A cuts a $D^2 \times [0,1]$ off of M', namely $B_1 \cup B_2 \cup \mathcal{N}(D)$; this situation is depicted in Figure 3.5c or 3.5d. ▢

Proposition 3.9 (Existence of Certificate 1)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented fibred 3-manifold M, and let F be a transversely oriented least-weight normal fibre of M. Then $\mathfrak{S}^*_{\text{fib}}(M, F)$ is non-empty.*

Proof. Once again, we refer the reader to Figure 3.4 for a depiction of the objects involved in the construction. We employ the notation of Certificate 1 concerning M', X, Y, F_g, and F'_g.

Constructing F_p, F_0, and F_1. Denote by Y the parallelity bundle of M'. Let A be the union of the components of $\text{thin}(X) \cap \text{thick}(Y)$ that are vertical in M', and let A' be the union of the other components; by Proposition 3.8, every component of A' cuts a $D^2 \times [0,1]$ off of M'. Note that if a component of $\text{thick}(Y)$ intersects A, then the interval bundle structure of that component agrees with the interval bundle structure of M'; in particular, said component does not intersect A'. We let F_p be the largest sub-2-complex of F such that $\text{thick}(\mathcal{N}^+(F_p))$ is the union of the components of $\text{thick}(Y)$ that intersect A; by the above remark, we see that $\text{thick}(\mathcal{N}^+(F_p))$ is disjoint from A'. We then let F_0 and F_1 be the sub-2-complexes of F such that the intersection of the remaining components of thick(Y) – that is, those that intersect A' – with $\partial_h M'$ is thick($F_0^+ \cup F_1^-$). These choices guarantee that properties (a.1) to (a.5) are satisfied.

D_0 **and** D_1'**.** Consider a component C of thick($\mathcal{N}^+(F_0)$), endowed with its intrinsic interval bundle structure ; in particular, we have that $\partial_v C \subseteq A' \cup \partial_v M'$. Recall that every component of A' cuts a $D^2 \times [0,1]$ off of M'. There must be a component of $A' \cap C$ that cuts a $D^2 \times [0,1]$ containing C off of M'; otherwise, the interval bundle structure of C would agree with that of M', contradicting the fact that C intersects A'. Let A'' be such a component of $A' \cap C$, and denote by B a $D^2 \times [0,1]$ cut off by A'' that contains C. It is clear from Figure 3.5 that the intersection $\partial_0 M' \cap B$ is thick(E^+) for a sub-2-complex E of F such that thick(E) is a union of 0, 1, or 2 disjoint discs, each of which intersects ∂F in at most one arc.

Let $\mathcal{D}$ be the set of all the sub-2-complexes obtained this way; more precisely, set

$$\mathcal{D} = \left\{ E : \begin{array}{c} C \text{ is a component of thick}(\mathcal{N}^+(F_0)),\ A'' \text{ is a component of } C \cap A',\ B \text{ is a} \\ D^2 \times [0,1] \text{ cut off of } M' \text{ by } A'' \text{ such that } C \subseteq B,\ E \text{ is a sub-2-complex} \\ \text{of } F \text{ such that thick}(E^+) \text{ is a component of } \partial_0 M' \cap B \end{array} \right\}.$$

We claim that it is impossible for two sub-2-complexes E_1 and E_2 in $\mathcal{D}$ to intersect but not be contained one in the other. Suppose for a contradiction that this is the case, and let $B_i = $ thick($\mathcal{N}^+(E_i)$) for $i \in \{1,2\}$, with the interval bundle structure induced by the fact that each of them is a $D^2 \times [0,1]$ cut off by some non-vertical annulus or square, say A_1 and A_2 respectively. Note that thick(E_1) and thick(E_2) are two discs in F that intersect, are not contained one in the other, and such that ∂_F thick(E_1) and ∂_F thick(E_2) are disjoint. Therefore, it follows that $F = $ thick(E_1) $\cup$ thick(E_2) – in particular, the surface F is a sphere or a disc – and hence that $M' = B_1 \cup B_2$. Moreover, the interval bundle structures of B_1 and B_2 agree on their intersection, and therefore extend to an interval bundle structure on M', which has to agree with the original one on M', since it has the same horizontal boundary. However, this contradicts the fact that A_1 and A_2 are not vertical in M'.

We now let D_0 be the union of all the sub-2-complexes in $\mathcal{D}$. It follows from the definition of $\mathcal{D}$ that property (d.2) is satisfied. By the claim above, any two maximal sub-2-complexes in $\mathcal{D}$ are disjoint, and therefore property (d.1) is satisfied. Finally, to show that D_0 is disjoint from F_p, let E be an element of $\mathcal{D}$, and let $B = \mathcal{N}^+(E)$. Note that every component of thick(Y) is either disjoint from or contained in B. If, for the sake of contradiction, some component C of thick(F_p) is contained in B, then B contains a vertical arc. It is then easy to see that this contradicts the assumption that B is a $D^2 \times [0,1]$ cut off by a non-vertical annulus or square; we refer the reader to Figure 3.5 once again. This shows that property (d.3) is satisfied.

Repeating the same exact constructions for components of $\mathrm{thick}(\mathcal{N}^-(F_1))$, we obtain a sub-2-complex D_1' of F such that $F_1 \subseteq D_1'$ and D_1' is a union of disjoint discs, each of which intersects ∂F in a single arc; moreover, the sub-2-complex $(D_1')^-$ of $\partial_1 M$ is disjoint from $\mathcal{N}^+(F_p)$. We remark that, by constructing D_0 and D_1', we have also indirectly proved that $\mathrm{thick}(F_0)$ and $\mathrm{thick}(F_1)$ are planar surfaces.

Constructing N, N', N'', G_0'', G_1'', f_g, D_1, $\mathcal{R}$, Σ', and Σ''. Let $N_0' = X$, $N_0'' = \mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$, $N_0 = N_0' \cup N_0''$. We take N to be $\mathrm{ab}(N_0)$. We also let $N' = \mathrm{emb}_{N_0}^{-1}(N_0')$ and $N'' = \mathrm{emb}_{N_0}^{-1}(N_0'')$, noting that property (b.4) is satisfied. Moreover, we see that $\mathrm{thin}(N_0)$ is the closure in M' of a union of components of $M' \setminus A$. Since A is vertical in M', we see that $\mathrm{thin}(N_0)$ is a product interval bundle over a compact orientable surface – namely $\mathrm{thin}(F_g \cup F_0)$; this implies that N is a product interval bundle over $G = \mathrm{ab}(F_g \cup F_0)$, and in particular that property (b.1) is satisfied. This interval bundle structure has $\partial_i N = \mathrm{emb}_{N_0}^{-1}(\partial_i M')$ for $i \in \{0,1\}$.

Let N_0'' inherit the intrinsic pre-sutured subcomplex structure from the parallelity bundle of M'. This induces a pre-sutured subcomplex structure on N'', with

$$\partial_h N'' = \mathrm{emb}_{N_0}^{-1}(\partial_h N_0'') = \mathrm{emb}_{N_0}^{-1}(F_0^+ \cup F_1^-) \quad \text{and} \quad \partial_v N'' = \mathrm{emb}_{N_0}^{-1}((X \cup \partial_v M') \cap N_0'').$$

For $i \in \{0,1\}$, let $G_i'' = \mathrm{emb}_{N_0}^{-1}(\partial_i N_0'')$; this assignment depends on the choice of intrinsic interval bundle structure, but will always satisfy properties (b.5) and (b.6). Property (b.7) is also satisfied, with $G'' = \mathrm{ab}(G_0'')$.

We let $f_g \colon N' \to X$ be the restriction of emb_{N_0} to N'. This is a simplicial map, that can be realised by gluing together the edges of N' that are identified in M' under emb_{N_0}; in particular, the map f_g restricts to a homeomorphism $\mathrm{int}(N') \to \mathrm{int}(X)$, as property (c.1) requires. Properties (c.2) to (c.4) easily follow from the definitions of N, N', and N'', since emb_{N_0} is essentially the identity up to some edge identifications, and these edges lie on the intersection $X \cap \mathcal{N}^+(F_p)$. As far as D_1 is concerned, we simply set $D_1 = \mathrm{emb}_{N_0}^{-1}((D_1')^-)$. Property (d.4) is satisfied because the same statement holds for $D_1' \subseteq F$, while property (d.5) is a consequence of the containment

$$\partial_1 N \cap N'' = \mathrm{emb}_{N_0}^{-1}(F_1^-) \subseteq \mathrm{emb}_{N_0}^{-1}((D_1')^-) = D_1;$$

Property (d.6) follows from the fact that, as remarked previously, the sub-2-complex $(D_1')^-$ of $\partial_1 M'$ is disjoint from $\mathcal{N}^+(F_p)$.

Finally, we show how to construct the triangulation $\mathcal{R}$. Let $\mathcal{R}'$ be the pull-back of the triangulation of the guts X under emb_{N_0}; this defines a triangulation of the sub-3-complex N' of N. Note that every component of $\partial_v N''$ intersects N', since if – say – F_0 contained a boundary component of F, then it would be impossible for D_0 to contain F_0. Since G'' is planar, we can apply Proposition 3.2 to obtain a suitable pre-sutured triangulation $\mathcal{R}''$ of N'' that agrees with $\mathcal{R}'$ on $N' \cap N''$, and such that the restriction of $\mathcal{R}''$ to $\partial_h N''$ is flapless. The two triangulations $\mathcal{R}'$ and $\mathcal{R}''$ glue together to form a triangulation $\mathcal{R}$ of N; since $\partial_h N$ is simplicial in this triangulation, we can endow $\mathcal{R}$ with the structure of a suitable pre-sutured triangulation of N. Properties (b.2) and (b.8) are essentially tautological, while property (b.3) is a consequence of the fact that the restrictions of $\mathcal{R}'$ to $\partial_h N \cap N'$ and of $\mathcal{R}''$ to $\partial_h N''$ are both flapless. The existence of certificates Σ' and Σ'' satisfying properties (f.1) and (f.2) follows from Theorem 3.1.

Quantitative bounds. Recall that F_g is simplicially isomorphic to $\partial_0 X$; Proposition 2.10 the implies that $\mathrm{area}(F_g) \leqslant 32t$. Moreover, since the triangulation of F_g is flapless, we have that $\ell(\partial F_g) \leqslant \mathrm{area}(F_g) \leqslant 32t$. Since $\partial_F F_0$ is contained in ∂F_g, the bound on its length follows immediately; the same argument used for F_0 would also show that $\ell(\partial_F F_1) \leqslant 32t$. By construction, we have that $\partial_F D_0 \subseteq \partial_F F_0$, and hence $\ell(\partial_F D_0) \leqslant 32t$ as well. Finally, we recall that $D_1 = \mathrm{emb}_{N_0}^{-1}((D_1')^-)$, and that $\partial_F D_1' \subseteq \partial_F F_1$, so we deduce that $\ell(\partial D_1) = \ell(\partial_F D_1') \leqslant 32t$; this concludes the proof of property (e.1).

As far as property (e.2) is concerned, we only need to bound the sizes of $\mathcal{R}'$ and $\mathcal{R}''$. The former is addressed by Proposition 2.10, giving that the number of tetrahedra of $\mathcal{R}'$ is at most $50t$. A bound on the latter comes from Proposition 3.2, provided that we can estimate the number of triangles in $N' \cap N''$. Note that $N' \cap N''$ is a union of triangles in $\mathrm{emb}_{N_0}^{-1}(\partial_v X)$, so Proposition 2.10 bounds its area by $36t$. Proposition 3.2 then gives a bound of $936t$ for the size of $\mathcal{R}''$, and hence a bound of $986t$ for the size of $\mathcal{R}$. ▢

3.3.2 Correctness

We prove that, whenever a connected normal surface F in a triangulated compact connected oriented 3-manifold M admits a certificate in $\mathfrak{S}^*_{\mathrm{fib}}(M, F)$, then M is fibred and F is a fibre of M. We start with a preliminary lemma.

Lemma 3.10 (Extending homeomorphisms of the boundary of interval bundles)**.** *Let M*

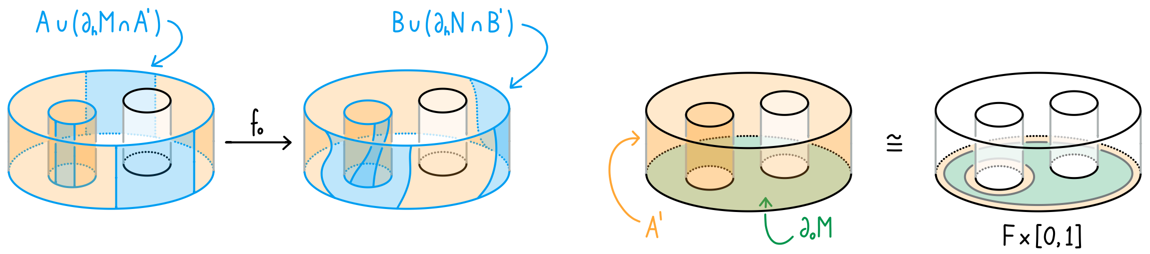

(a) (b)

Figure 3.8. (a) The homeomorphism f_0 restricts to a homeomorphism from $A \cup (\partial_h M \cap A')$ to $B \cup (\partial_h N \cap B')$, and this restriction can be extended to a homeomorphism $A' \to B'$. The union of the shaded regions is, respectively, A' on the left and B' on the right. (b) The pair (M, A'') is homeomorphic to $(F \times [0,1], F \times \{0\})$; this can be seen by "flattening" A' on the horizontal boundary of $F \times [0,1]$.

and N be pre-sutured manifolds that are both homeomorphic to $F \times [0,1]$ for some compact orientable surface F, and let A be a vertical subset of $\partial_v M$. Let $f_0 \colon \partial_h M \cup A \to \partial N$ be an embedding such that $f_0(\partial_i M) = \partial_i N$ for $i \in \{0,1\}$. Then there is a homeomorphism $f \colon M \to N$ of pre-sutured manifolds that coincides with f_0 on A.

Proof. Firstly, note that $B = f_0(A)$ is a vertical subset of $\partial_v N$. Let A' be the union of all the components of $\partial_v M$ that intersect A. Consider the restriction of f_0 to $A \cup (\partial_h M \cap A')$; this is a homeomorphism to $B \cup (\partial_h N \cap B')$ for some union B' of components of $\partial_v N$, and it can be extended to a homeomorphism $f' \colon A' \to B'$ (see Figure 3.8a). Let $A'' = \partial_0 M \cup A'$ and $B'' = \partial_0 N \cup B'$. Since $\partial_0 M$ and $\partial_0 N$ are homeomorphic, we can extend f' to a homeomorphism $f'' \colon A'' \to B''$. Finally, note that (M, A'') and (N, B''), as pairs of topological spaces, are both homeomorphic to $(F \times [0,1], F \times \{0\})$; this is shown in Figure 3.8b. Every homeomorphism $F \times \{0\} \to F \times \{0\}$ extends to a homeomorphism $F \times [0,1] \to F \times [0,1]$; hence, the homeomorphism $f'' \colon A'' \to B''$ extends to a homeomorphism $f \colon M \to N$. This can be chosen so that $f(\partial_i M) = \partial_i N$ for $i \in \{0,1\}$, and by construction we have that f agrees with f_0 on A, as required. ▢

Proposition 3.11 (Correctness of Certificate 1)**.** *Suppose that M is a triangulated compact connected oriented 3-manifold, and F is a transversely oriented normal surface in M. Let $\Sigma \in \mathfrak{S}^*_{\text{fib}}(M, F)$ be a valid certificate. Then the following hold.*

1. *Let $M_N = X \cup \mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$, and endow M_N with the structure of a pre-sutured*

subcomplex given by

$$\partial_i M_N = \partial_i M' \cap M_N \qquad \text{for } i \in \{0,1\},$$
$$\partial_v M_N = (\partial_v M' \cup \mathcal{N}^+(F_p)) \cap M_N.$$

Then there exists a map $f\colon N \to M_N$ *such that* $f(\partial_i N) = \partial_i M_N$ *for* $i \in \{0,1\}$*; moreover, the map* f *restricts to* f_g *on* N'*, and to a homeomorphism* $N'' \to \mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$*.*

2. *The pre-sutured manifold* $\mathrm{ab}(M_N)$ *is a product interval bundle over a compact orientable surface; moreover, it is homeomorphic to* N *as a pre-sutured manifold.*

3. *The pre-sutured manifold* M' *is a product interval bundle over* F*, and the interval bundle structure agrees with that of* M_N*, as well as with the intrinsic one on* $\mathcal{N}^+(F_p)$*.*

4. *The properly embedded surface* $\mathrm{thin}(X) \cap \mathrm{thick}(\mathcal{N}^+(F_p))$ *is vertical in* M'*.*

Proof. Let $M'' = \mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1)$, so that $M_N = X \cup M''$. Endow M'' with its intrinsic interval bundle structure. Let $A = M'' \cap X$ and $B = N'' \cap N'$; by property (c.3), the image of B under f_g is precisely A. By property (c.2), there is a homeomorphism $f_0\colon \partial_h N'' \to \partial_h M''$ whose restriction to $B \cap \partial_h N$ is a simplicial isomorphism $B \cap \partial_h N \to A \cap \partial_h M'$ that agrees with f_g. This implies that the restriction of f_g to B is actually a homeomorphism $B \to A$. Denote by f_1 the homeomorphism $\partial_h N'' \cup B \to \partial_h M'' \cup A$ that agrees with f_0 on $\partial_h N''$ and with f_g on B. Note that each component of B intersects two distinct components of $\partial_h N''$, as N'' is a product interval bundle by property (b.7). Therefore, we see that each component of A intersects two distinct components of $\partial_h M''$. It follows that M'' is also a product interval bundle, and that A is a vertical subset of $\partial_v M''$. We can then apply Lemma 3.10 to extend f_1 to a homeomorphism $f_2\colon N'' \to M''$ of pre-sutured subcomplexes that agrees with f_1 on B. For a completely rigorous justification of this step, one should in fact apply Lemma 3.10 to $\mathrm{ab}(N'')$ and $\mathrm{ab}(M'')$ instead of N'' and M'', and exploit the fact that the singular locus of N'' is contained in B to obtain the desired homeomorphism f_2.

Let $f\colon N \to M_N$ be the map that coincides with f_g on N' and with f_2 on N''. By construction of f_2 and property (c.2), we have that $f(\partial_i N) = \partial_i M_N$ for $i \in \{0,1\}$. Moreover, since f_2 is a homeomorphism and f_g restricts to a homeomorphism on the interior of M', we see that f restricts to a homeomorphism $\mathrm{int}(N) \to \mathrm{int}(M_N)$. It follows that f lifts to a homeomorphism of pre-sutured manifolds $N \to \text{ab}(M_N)$, thus proving that $\text{ab}(M_N)$ is a product interval bundle over a compact orientable surface, namely $\text{ab}(F_g \cup F_0)$.

Consider now the intrinsic interval bundle structure on $\mathcal{N}^+(F_p)$. The intersection $\mathcal{N}^+(F_p) \cap M_N$ is a vertical subset of $\partial_v \mathcal{N}^+(F_p)$, since it is the intersection of the guts of M' with the union of product parallelity components $\mathcal{N}^+(F_p)$. It is not hard to see that $\mathcal{N}^+(F_p) \cap M_N$ is also a vertical subset $\partial_v M_N$; in fact, by property (a.5), this intersection is a union of squares in $\partial_v M_N$, each of which intersects both $\partial_0 M_N$ and $\partial_1 M_N$. We deduce that the interval bundle structures on $\mathcal{N}^+(F_p)$ and M_N glue together to give a product interval bundle structure on the pre-sutured manifold M'. Finally, since $\text{thin}(X) \cap \text{thick}(\mathcal{N}^+(F_p))$ is a union of boundary components of a regular neighbourhood of $\mathcal{N}^+(F_p) \cap M_N$, it is a vertical surface in M'. ▢

3.3.3 Verification

We prove that the validity of a certificate in $\mathfrak{S}_{\text{fib}}(M, F)$ can be verified in polynomial time in the size of the certificate.

Proposition 3.12 (Verification of Certificate 1)**.** *There is an algorithm that takes as input*

- *a triangulation of a compact connected oriented* 3*-manifold* M*,*
- *a transversely oriented connected normal surface* F *in* M*, and*
- *a certificate* $\Sigma \in \mathfrak{S}_{\text{fib}}(M, F)$*,*

and decides whether $\Sigma \in \mathfrak{S}^*_{\text{fib}}(M, F)$*. The running time of the algorithm is polynomial in* $|\Sigma|$.

Proof. Many of the required properties can be easily verified, either by direct inspection, or by a straightforward application of the algorithm of Agol, Hass, and Thurston or one of its corollaries listed in Section 2.4. For example, one can check that $\mathcal{R}$ is a pre-sutured triangulation of N by direct verification, and that F_p is a sub-2-complex of F using the algorithm of Proposition 2.17. We will not discuss other instances of this kind of verification, and instead focus on the steps that require a non-trivial algorithmic approach.

Properties (a.). The sub-2-complexes F_g and F_g' of F can be constructed explicitly, using Proposition 2.16 to compute transverse orientations. Proposition 2.18 immediately tells us whether property (a.1) holds. We can apply Proposition 2.17 to find the components of $\mathrm{clos}(F \setminus F_g)$ and of F_0, and check that each component of the latter is also a component of the former. This allows us to compute F_p, and verify that property (a.2) holds. We perform an analogous test for F_1 and F_g', making sure that properties (a.3) and (a.4) are satisfied. Finally, to verify property (a.5), we pick a triangle T in each component of F_p, and we check that the transverse orientations of F at T and $\triangle(T)$ agree using Proposition 2.16; we also verify that $\triangle(T) \subseteq F_p'$ using Proposition 2.18. This guarantees that $\mathcal{N}^+(F_p) \subseteq \mathcal{N}^-(F_p')$; the verification of the other containment is perfectly symmetric.

Properties (b.). Properties (b.1) and (b.2) can be verified with the aid of the certificate Σ'. Similarly, properties (b.7) and (b.8) can be verified with the aid of the certificate Σ''. The other properties can be checked by direct inspection in polynomial time in the size of the triangulation $\mathcal{R}$.

Properties (c.). The only property whose verification is non-trivial is property (c.4). We can directly check that f_g restricts to a simplicial isomorphism

$$N' \cap \partial_h N'' \longrightarrow (\partial_F F_0)^+ \cup (\partial_F F_1)^-.$$

We are then in the following situation: we have a homeomorphism from a subset of the boundary of $\partial_h N''$ to a subset of the boundary of $F_0^+ \cup F_1^-$, and we want to extend it to a homeomorphism $\partial_h N'' \to F_0^+ \cup F_1^-$. We can answer this question by analysing each component C of $\partial_h N''$ separately. If C is disjoint from N', then there is no such extension, because every component of F_0 and F_1 intersects X. Otherwise, the subset $f_g(C \cap N')$ of $(\partial_F F_0)^+ \cup (\partial_F F_1)^-$ contains at least one edge. If more than one component of $F_0^+ \cup F_1^-$ intersects $f_g(C \cap N')$ in at least one edge, then there is no extension. Otherwise, suppose without loss of generality that there is a unique component C' of F_0 such that $f_g(C \cap N') \subseteq (C')^+$. If $f_g(C \cap N') \neq (\partial_F C')^+$, then there is no extension. Otherwise, let $h = ((-)^+)^{-1} \circ f_g|_{C \cap N'}$; this is a simplicial isomorphism from $C \cap N'$ to $\partial_F C'$. We only need to check that h extends to an orientation-preserving homeomorphism $C \to C'$. A first necessary condition is that $\mathrm{ab}(C)$ and $\mathrm{ab}(C')$ are homeomorphic; this can be verified thanks to Proposition 2.17. Secondly, we need the $\partial_h N''$-boundary sequences of C and the F-boundary sequences of C' to be compatible, in the following sense. There must exist a bijection between components of $\partial\,\mathrm{ab}(C)$ and components of $\partial\,\mathrm{ab}(C')$ such that, if a component b of $\partial\,\mathrm{ab}(C)$ corresponds to a component b' of $\partial\,\mathrm{ab}(C')$ under this bijection, and a $\partial_h N''$-boundary sequence of C is $e_1, \ldots, e_k$, then $e_1', \ldots, e_k'$ is an

F-boundary sequence of C', where

$$\begin{cases} (e_i')^+ = f_g(e_i) & \text{if } e_i \text{ is an edge of } C \cap N', \\ e_i' = \partial F & \text{if } e_i = \partial\partial_h N'' \end{cases} \qquad \text{for each } i \in \{1, \dots, k\}.$$

Boundary sequences can be computed thanks to Proposition 2.17; the existence of a bijection with the desired property is easy to verify, since each edge of C' appears in at most one F-boundary sequence. If these conditions hold for each component C of $\partial_h N''$, then property (c.4) is satisfied, otherwise it is not.

Properties (d.). Properties (d.2) and (d.3) can be verified using Proposition 2.17, and properties (d.5) and (d.6) can be inspected directly. As far as property (d.1) is concerned, we can use Proposition 2.17 again to retrieve information about the components of D_0. We can then manually combine components of D_0 into components of thick(D_0), since the number of singular points of D_0 is bounded above by $\ell(\partial_F D_0)$. In particular, we can obtain the list of components of thick(D_0), together with their Euler characteristics and F-boundary sequences. This is enough to check whether property (d.1) is satisfied.

Properties (e.). These conditions are just inequalities, which can be verified directly.

Properties (f.). These properties can be verified thanks to the algorithm of Theorem 3.1. ▢

3.4 Algorithmic fibredness detection

So far, we have developed the tools to certify that a given normal surface in a triangulated 3-manifold is the fibre of a fibration over the circle. Crucially, and unavoidably, the size of such a certificate depends on the "size" – that is, the logarithm of the weight – of the normal surface. Therefore, the only issue we need to address is whether every triangulated fibred 3-manifold admits a normal fibre whose size is polynomial in the size of the triangulation. This question is answered in the affirmative by Proposition 3.13 below, for fibred 3-manifolds admitting a fibre of negative Euler characteristic. We remark that this argument is not unknown to experts: in fact, it appears almost verbatim in [36], and it is anyway a straightforward application of the results in [39]. However, we include it here for the sake of completeness, especially since the statement of [36, Theorem 6.3.3] requires the 3-manifold to be atoroidal, and does not provide a bound on the Euler characteristic.

Proposition 3.13 (A fibre of small weight and negative Euler characteristic)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected orientable 3-manifold M with t tetrahedra. Suppose that M fibres over the circle with fibre F_0 of negative Euler characteristic. Then there is a normal fibre F of M with $\chi(F) \geqslant \chi(F_0)$ and*

$$w(F) \leqslant t^3 \cdot 2^{7t+6}.$$

As remarked above, the proof of Proposition 3.13 relies heavily on the tools developed in [39], of which we give a very brief summary. Let $\mathcal{T}$ be a triangulation of a compact connected oriented irreducible 3-manifold M with (possibly empty) toroidal boundary, and denote by t the number of tetrahedra of $\mathcal{T}$. The matching equations of $\mathcal{T}$ (see Section 2.2.5) define a system of linear equations in $\mathbb{R}^{7t}$; the non-negative solutions of this system form a convex polyhedral cone $\mathcal{P}_{\mathcal{T}} \subseteq \mathbb{R}^{7t}_{\geqslant 0}$. Denote by $\overline{\mathcal{P}}_{\mathcal{T}}$ the set of points $\mathbf{x} \in \mathcal{P}_{\mathcal{T}}$ such that $x_1 + \ldots + x_{7t} = 1$, which is a compact convex polyhedron. In particular, we can talk about the *faces* of $\overline{\mathcal{P}}_{\mathcal{T}}$; more precisely, a face of $\overline{\mathcal{P}}_{\mathcal{T}}$ is the intersection of $\overline{\mathcal{P}}_{\mathcal{T}}$ with a hyperplane of the form $\{\mathbf{x} \in \mathbb{R}^{7t} : x_i = 0 \text{ for } i \in \mathcal{I}\}$ for some subset $\mathcal{I} \subsetneq \{1, \ldots, 7t\}$. If C is a face of $\overline{\mathcal{P}}_{\mathcal{T}}$, we denote by $\operatorname{cone}(C)$ the set of positive scalar multiples of points in C. More generally, we use the same notation for any subset C of some Euclidean space $\mathbb{R}^n$.

For every normal surface F in M, its normal vector $\mathbf{v}_F$ is a point in $\mathcal{P}_{\mathcal{T}} \cap \mathbb{Z}^{7t}$. Conversely, every point $\mathbf{w} \in \mathcal{P}_{\mathcal{T}} \cap \mathbb{Z}^{7t}$ that also satisfies the consistency equations of $\mathcal{T}$ defines a normal surface $F_{\mathbf{w}}$ in M. For a non-empty normal surface F, let $\overline{\mathbf{v}_F}$ be the projection of $\mathbf{v}_F$ onto $\overline{\mathcal{P}}_{\mathcal{T}}$ – that is, the unique point in $\overline{\mathcal{P}}_{\mathcal{T}}$ that is a scalar multiple of $\mathbf{v}_F$; this will necessarily be a rational point. We say that a face C of $\overline{\mathcal{P}}_{\mathcal{T}}$ *carries* a normal surface F if $\overline{\mathbf{v}_F} \in C$. For every normal surface F, there is a unique minimal face C of $\overline{\mathcal{P}}_{\mathcal{T}}$ that carries F, which we denote by C_F. Note that if a face C of $\overline{\mathcal{P}}_{\mathcal{T}}$ carries a normal surface, then every integral point in $\operatorname{cone}(C)$ will satisfy the consistency equations, and hence represent a normal surface in M.

For a compact orientable surface F, define

$$\chi_-(F) = -\chi(F \setminus \{\text{sphere and disc components of } F\}).$$

The *Thurston seminorm* of a homology class $\alpha \in H_2(M, \partial M; \mathbb{R})$ is defined as the minimum value of $\chi_-(F)$ over all oriented surfaces F properly embedded in M that represent α. This can be extended to a seminorm $x \colon H_2(M, \partial M; \mathbb{R}) \to \mathbb{R}$ by linearity, convexity, and continuity (see [37] for details). Tollefson and Wang call an oriented properly embedded incompressible boundary-incompressible surface $F \subseteq M$ *taut* if its homology class $[F] \in H_2(M, \partial M; \mathbb{R})$ is non-zero, no union of components of F is homologically trivial, and $x([F]) = \chi_-(F)$. If, moreover, a taut surface is in general position with respect to $\mathcal{T}$ and has minimal weight amongst all taut general position surfaces in its homology class, then it is called *lw-taut*; note that this notion of "lw" does not coincide with our notion of "least-weight".

The following is one of the main results of [39].

Theorem 3.14 ([39, Theorem 3.3 and Corollary 3.4])**.** *Let F be an oriented lw-taut normal surface in M. Then every normal surface carried by C_F is lw-taut. Moreover, there is a way to assign orientations to normal surfaces carried by C_F such that, for every pair G, H of normal surfaces carried by C_F, the following equalities hold:*

$$[G + H] = [G] + [H] \quad \textit{and} \quad x([G + H]) = x([G]) + x([H]).$$

We are now ready to prove Proposition 3.13.

Proof of Proposition 3.13. Let G be an oriented least-weight normal fibre of M that is isotopic to F_0. Note that, as discussed at the beginning of [37, §3], every incompressible surface in M that is homologous to F_0 is actually isotopic to it; in particular, this implies that G is lw-taut. From now on, all the surfaces we consider will be carried by C_G; we will implicitly endow them with the orientation given by Theorem 3.14. The same theorem then implies that every normal surface carried by C_G is lw-taut, and that the operation of taking homology representatives is linear with respect to the normal sum on C_G; similarly, the Thurston seminorm x is linear with respect to the normal sum on C_G.

Let C' be the minimal face of the x-unit ball in $H_2(M, \partial M; \mathbb{R})$ such that $[G]$ lies in $\text{cone}(C')$. Crucially, by the linearity properties mentioned above, the homology classes of all the surfaces carried by C_G will also lie in the closure of $\text{cone}(C')$. Since G is a fibre of M, Theorems 3 and 5 of [37] imply that C' is a top-dimensional face of the x-unit ball, and every integral homology class in the interior of $\text{cone}(C')$ is represented by a union of parallel fibres of M; in fact, every oriented incompressible surface with no sphere and disc components and whose homology class lies in the interior of $\text{cone}(C')$ is a union of parallel fibres of M.

Let now $F_1, \ldots, F_m$ be fundamental normal surfaces in M such that $G = F_1 + \ldots + F_m$; it is clear that these are all carried by C_G. Up to reordering, we can assume that $\{[F_1], \ldots, [F_n]\}$ is a basis of the subspace of $H_2(M, \partial M; \mathbb{R})$ spanned by $[F_1], \ldots, [F_m]$ for some $1 \leqslant n \leqslant m$. Note that the integer n is bounded above by

$$n \leqslant \dim H_2(M, \partial M; \mathbb{R}) \leqslant 2t,$$

where the second inequality follows from the fact that there are at most $2t$ triangles of $\mathcal{T}$ that do not lie on ∂M. Let $F' = F_1 + \ldots + F_n$. We know that there is at least one linear combination of $[F_1], \ldots, [F_n]$ that lies in the interior of cone(C') – namely, $[G]$. Therefore, the class $[F']$ must also lie in the interior of cone(C'). In particular, since F' is lw-taut, this implies that each component of F' must be a fibre of M.

Concerning the weight of F', Proposition 2.6 gives the inequality

$$w(F_i) \leqslant t^2 \cdot 2^{7t+5} \qquad \text{for every } 1 \leqslant i \leqslant n,$$

from which we get the bound

$$w(F') = w(F_1) + \ldots + w(F_n) \leqslant n \cdot t^2 \cdot 2^{7t+5} \leqslant t^3 \cdot 2^{7t+6}.$$

The Euler characteristic of F' satisfies

$$\begin{aligned}
-\chi(F') &= x([F']) && \text{since } F' \text{ is lw-taut} \\
&= x([F_1] + \ldots + [F_n]) \\
&\leqslant x([F_1]) + \ldots + x([F_m]) && \text{by convexity of } x \\
&= x([G]) && \text{by linearity of } x \\
&= -\chi(G) && \text{since } G \text{ is lw-taut.}
\end{aligned}$$

We conclude the proof by taking F to be a component of F'. As previously observed, the surface F is a fibre of M; its weight and Euler characteristic satisfy

$$w(F) \leqslant w(F') \leqslant t^3 \cdot 2^{7t+6} \qquad \text{and} \qquad \chi(F) \geqslant \chi(F') \geqslant \chi(G) = \chi(F_0),$$

as desired. ▢

As we will see in the proof of Theorem 3.17, recognising 3-manifolds that fibre over the circle with fibre a sphere, disc, or annulus is already known to be in NP. Since Proposition 3.13 covers 3-manifolds with fibre of negative Euler characteristic, the only remaining case is when the fibre is a torus. In this setting, the argument of Proposition 3.13 does not work, since it heavily relies on the Thurston seminorm, which vanishes everywhere for these 3-manifolds. Fortunately, torus bundles are simple enough that we can completely classify their incompressible surfaces – or, at least, the ones we care about.

Let M be a compact oriented 3-manifold that fibres over the circle with fibre a torus, and let $F_0 \subseteq M$ a transversally oriented toroidal fibre of M. Our goal is to show that all connected orientable incompressible surfaces in M that are not homologically trivial are fibres of *some* fibration of M (that is, not necessarily the one defined by F_0). Upon fixing a basis of $H_1(F_0; \mathbb{Z})$, the monodromy φ of F_0 is represented by an element of $\mathrm{SL}(2, \mathbb{Z})$. It is well-known that if an element of $\mathrm{SL}(2, \mathbb{Z})$ has an integral eigenvector, then it is conjugate to

$$\mathbf{L}_n^+ = \begin{pmatrix} 1 & 0 \\ n & 1 \end{pmatrix} \qquad \text{or} \qquad \mathbf{L}_n^- = \begin{pmatrix} -1 & 0 \\ n & -1 \end{pmatrix}$$

for some $n \in \mathbb{Z}$.

Proposition 3.15 (Incompressible surfaces in torus bundles)**.** *In the setting above, let F be a closed orientable incompressible surface embedded in M that is not null-homologous. Then F is a toroidal fibre of some fibration of M.*

Proof. Since F_0 is a fibre of M, the 3-manifold $M' = M \setminus\!\setminus F_0$ is homeomorphic to $F_0 \times [0, 1]$. In fact, it is useful to have a concrete model for M' as $M' = \mathbb{R}^2/\mathbb{Z}^2 \times [0, 1]$, and for M as

$$M = {\mathbb{R}^2/\mathbb{Z}^2 \times [0, 1]} \Big/ {\{(x, 1) \sim (\varphi(x), 0) : x \in \mathbb{R}^2/\mathbb{Z}^2\}},$$

where $F_0 = \mathbb{R}^2/\mathbb{Z}^2 \times \{0\}$. Isotope F so that it intersects F_0 transversely in the smallest possible number of curves. Denote by F' the intersection of F with M'. We now employ a very classical argument – sketched, for instance, in the proof of [26, Proposition 11.4.12] – to show that F' must be a collection of vertical annuli

First, we show that no component of $F \cap F_0$ can bound a disc in F_0. For a contradiction, suppose that there is a disc D embedded in F_0 such that $D \cap F = \partial D$ (up to taking an innermost disc, this is not restrictive). Since F is incompressible, there must be a disc D' embedded in F such that $\partial D' = \partial D$. By irreducibility of M, the sphere $D \cup D'$ must bound a 3-ball B in M; note that $F \cap B = D'$. The 3-ball B can be used to isotope D' to D, and then push it slightly past D, thus removing ∂D from the intersection $F \cap F_0$; this contradicts the minimality of the intersection. By swapping the roles of F and F_0 – note that we only used the fact that F is incompressible – we conclude that no component of $F \cap F_0$ can bound a disc in F_0. In particular, we have shown that $F \cap F_0$ is a collection of parallel curves in the torus F_0.

Next, we show that F' is incompressible and boundary-incompressible in M'. Let D be a compression disc for F' in M'. Since F is incompressible in M, there is a disc D' embedded in F such that $\partial D' = \partial D$. If the disc D' is contained in M', then we see that D is trivial. Otherwise, we find a component of $F \cap F_0$ that bounds a disc in F_0, the existence of which we have already ruled out; this proves that F' is incompressible. Let D be a boundary-compression disc for F' in M'. Up to taking an innermost disc, we can assume that $D \cap F' =\subseteq \partial D$. Let a_1 and a_2 the two components of $\partial F'$ intersecting D; these are two parallel curves on a component of $\partial_h M'$ – importantly, they are distinct because F' is orientable. Compress F' along D to obtain a new surface F''. Denote by a the boundary component of F'' that is obtained by surgery on $a_1 \cup a_2$ along $D \cap \partial_h M'$. Since a_1 and a_2 are parallel, we see that a bounds a disc D' in $\partial_h M'$. It is not hard to show, using incompressibility of F' and irreducibility of M', that the component of F'' containing a cobounds a 3-ball in M' with D'. This means that the component of F' containing a_1 and a_2 is an annulus cobounding a solid torus in M' with an annulus in $\partial_h M'$. This solid torus can be used to isotope said component of F' onto F_0, and then slightly past it, thus removing a_1 and a_2 from the intersection $F \cap F_0$; this contradicts the minimality of the intersection.

In conclusion, we found that F' is an orientable incompressible boundary-incompressible surface in M'. It is well-known (see [26, Proposition 9.3.18]) that any such surface is isotopic in M' either to a union of *horizontal* surfaces – that is, surfaces of the form $F_0 \times \{t\}$ for some $t \in [0,1]$ – or to a union of *vertical* surfaces – that is, surfaces of the form $a \times [0,1]$ for some essential curve a in F_0. If F' is horizontal, we deduce that F is isotopic to F_0 in M, and we are done. If F' is vertical, then it is a union of annuli; in particular, this implies that there is a slope on F_0 that is preserved by φ. In other words, the monodromy φ has an integral eigenvector, and it is therefore conjugate to $\mathbf{L}_n^+$ or $\mathbf{L}_n^-$ for some integer n. In fact, it is not restrictive to assume that $\varphi = \mathbf{L}_n^+$ or $\varphi = \mathbf{L}_n^-$. In particular, this implies that each component of $F \cap F_0$ is a curve isotopic in F_0 to $\{0\} \times \mathbb{R}/\mathbb{Z} \times \{0\}$. By "straightening" F', we can assume that each component of F' is a flat annulus of the form

$$\{(x,y,z) \in \mathbb{R}^2 \times [0,1] : \alpha \cdot x + \beta \cdot z = \gamma\} \subseteq \mathbb{R}^2/\mathbb{Z}^2 \times [0,1],$$

for constants $\alpha, \beta, \gamma \in \mathbb{R}$.

When $\varphi = \mathbf{L}_n^+$. This case is depicted in Figure 3.9a. Let T be the torus $\mathbb{R}^2/\mathbb{Z}^2$, and consider the projection $p\colon M \to T$ defined by

$$p(x, y, z) = (x, z).$$

It is easy to see that this projection defines a circle bundle structure over T on M. We see that $F_0 = p^{-1}(\mathbb{R}/\mathbb{Z} \times \{0\})$, while $F = p^{-1}(a)$ for some essential curve $a \subseteq T$. Let $\psi\colon T \to T$ be a homeomorphism that sends a to $\mathbb{R}/\mathbb{Z} \times \{0\}$. This homeomorphism lifts to a homeomorphism of M sending F to F_0; since F_0 is a toroidal fibre of M, we conclude that the same holds for F.

When $\varphi = \mathbf{L}_n^-$. In this case, depicted in Figure 3.9b, the 3-manifold M is a circle bundle over the Klein bottle K. More precisely, if we model K as the quotient of the Euclidean plane $\mathbb{R}^2$ by the group of isometries generated by $(x, z) \mapsto (x + 1, z)$ and $(x, z) \mapsto (-x, z + 1)$, then the projection map $p\colon M \to K$ is given by

$$p(x, y, z) = (x, z).$$

Like in the torus bundle case, we have that $F = p^{-1}(a)$ for some two-sided essential curve $a \subseteq K$. However, unlike in the torus bundle case, there are only two such curves up to isotopy. One is the meridian $m = \{(x, 0) \in K\}$, and its preimage $p^{-1}(m)$ in M is precisely F_0. The other two-sided essential curve is

$$b = \{(x, z) \in K : x = 1/4 \text{ or } x = 3/4\},$$

and it is separating in K; therefore, its preimage $p^{-1}(b)$ separates M. Since we are assuming that F is not null-homologous, we conclude that F must be isotopic to F_0 in M; in particular, it is a fibre of M. ▢

As a consequence of the previous proposition, we find that triangulated torus bundles admit a normal fibre of bounded weight.

Proposition 3.16 (A toroidal fibre of small weight)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected orientable 3-manifold M with t tetrahedra. Suppose that M fibres over the circle with fibre a torus. Then there is a toroidal normal fibre F of M with*

$$w(F) \leqslant t^2 \cdot 2^{7t+5}.$$

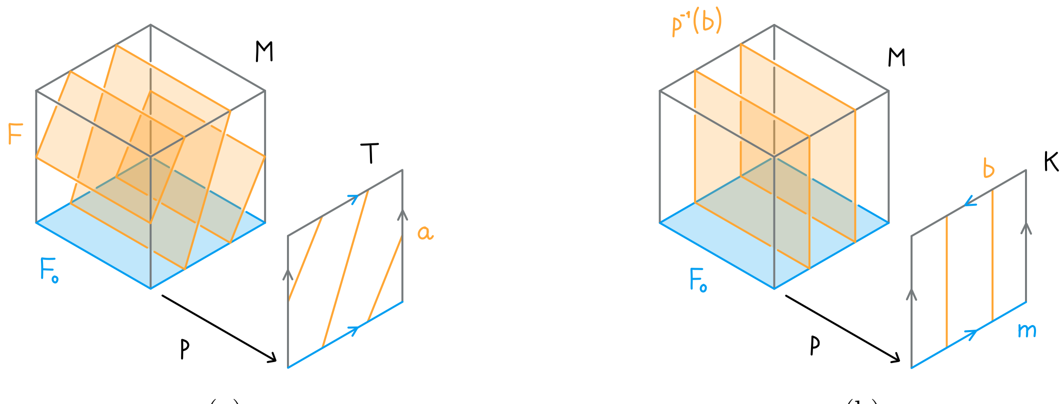

Figure 3.9. (a) When the monodromy φ is isotopic to $\mathbf{L}_n^+$ for some $n \in \mathbb{Z}$, the 3-manifold M is a circle bundle over the torus T; the surface F is related to F_0 by a homeomorphism of M. (b) When the monodromy φ is isotopic to $\mathbf{L}_n^-$ for some $n \in \mathbb{Z}$, the 3-manifold M is a circle bundle over the Klein bottle K; the surface F is isotopic to F_0, since $p^{-1}(b)$ is separating.

Proof. Let F be a least-weight toroidal normal fibre of M. We will prove that F is fundamental; by Proposition 2.6, this will imply the bound in the statement. Suppose, for the sake of contradiction, that F can be presented as a normal sum $F = G_1 + G_2$, with G_1 and G_2 non-empty normal surfaces. We can assume that G_1 and G_2 minimise the number of components of $G_1 \cap G_2$; by [28, Lemma 3.3.30], this guarantees that G_1 and G_2 are connected.

As we already observed at the start of the proof of Proposition 3.13, the fibre F is lw-taut. Tollefson and Wang's Theorem 3.14 then implies that every normal surface carried by C_F is lw-taut – in particular, this holds for G_1 and G_2. Therefore, the surfaces G_1 and G_2 are orientable, incompressible, and homologically non-trivial. By Proposition 3.15, we deduce that G_1 and G_2 are toroidal fibres of M; however, this contradicts the minimality of F. ▢

By combining Propositions 3.13 and 3.16, we obtain that if a triangulated orientable 3-manifold fibres over the circle, and the fibre is a torus or has negative Euler characteristic, then said 3-manifold admits a fibre of exponential weight in the number of tetrahedra. Since the size of our fibredness certificate depends logarithmically on the weight of the fibre, this is enough to certify fibredness in polynomial time. Formally, we consider the following decision problem.

Problem (Fibredness detection).

Input: a triangulation of a compact connected oriented 3-manifold M.

Output: whether M fibres over the circle.

The size of the input is measured by the number of tetrahedra in the triangulation of M.

×

Theorem 3.17 (Fibredness detection is in NP). *The problem* Fibredness detection *is in* NP*.*

Proof. The problems of recognising compact orientable 3-manifolds that fibre over the circle with fibre of non-negative Euler characteristic and not a torus have been already addressed in the literature. Specifically, the following recognition problems are known to be in NP:

- recognising $S^2 \times S^1$, by [16, Theorem 3];
- recognising $D^2 \times S^1$, by [16, Theorem 3];
- recognising $S^1 \times [0,1] \times S^1$, by Theorem 3.1 (that is, [22, Theorem 12.1], with the caveat that, since we are trying to recognise an interval bundle over the closed surface $S^1 \times S^1$, the discussion at the beginning of [24, Section 11] is necessary).

Since these are only a finite number of cases, they can each be handled by their respective verification algorithms, and we can therefore focus only on the case where M fibres over the circle with fibre of negative Euler characteristic or a torus.

The verification algorithm takes as input a triangulation $\mathcal{T}$ of a compact connected oriented 3-manifold M, and a certificate consisting of a transversely oriented normal surface F in M and a certificate $\Sigma \in \mathfrak{S}_{\text{fib}}(M, F)$. The algorithm then verifies that F is connected and orientable, using Propositions 2.15 and 2.16, and that $\Sigma \in \mathfrak{S}^*_{\text{fib}}(M, F)$, using the algorithm of Proposition 3.12. These verifications can be performed in polynomial time in $|\mathcal{T}|$ and $|\Sigma|$. If these checks are successful, then the third statement of Proposition 3.11 implies that M fibres over the circle with fibre F. Moreover, Propositions 3.9, 3.13, and 3.16 guarantee that, if M fibres over the circle with fibre of negative Euler characteristic or a torus, then there exist a transversely oriented connected normal surface F in M and a certificate $\Sigma \in \mathfrak{S}^*_{\text{fib}}(M, F)$ such that $|\Sigma|$ is bounded above by a polynomial in $|\mathcal{T}|$. ▢

Chapter 4

Certifying the monodromy

4.1 Outline of the certificate

In this chapter, we will describe a way to certify the image of a curve under the monodromy of a fibration of a 3-manifold. More precisely, the setting is the following. Let M be a compact oriented 3-manifold, triangulated with t tetrahedra. Suppose that M fibres over the circle with fibre $F \subseteq M$ and monodromy $\varphi \in \mathrm{Mcg}(S)$. Let a_0 be a normal curve in F. We will show how to construct a normal curve a_1 in F that is isotopic to $\varphi(a_0)$; at the same time, we will provide a certificate that can be used to verify that a_1 is indeed the image of a_0 under φ. The relevant parameters here are the number t of tetrahedra of M, the logarithm of the weight of F, the Euler characteristic of F, and the size of a_0 – which, as we recall from Section 2.2.6, is measured in terms of $\mathrm{area}(\mathrm{supp}(a_0))$ and $\log(w(a_0))$. Our certificate can be verified in polynomial time in these parameters.

The complete description of the monodromy certificate can be found in Certificate 2. In the following paragraphs, we will outline the main ideas behind the construction of the certificate, ignoring the many technical details that will be expanded on in Section 4.3. Our monodromy certificate builds upon the fibredness certificate we defined in Chapter 3. As a brief reminder, for a least-weight normal fibre F, we have decomposed the parallelity bundle Y of $M' = M \setminus\setminus F$ into two parts: a union Y_v of parallelity components that are vertical in $M' \cong F \times [0,1]$, and the union Y_t of the remaining parallelity components, which is contained in a union of "tubes" in M'. By gluing Y_t to the guts X of M', we obtain a 3-manifold N, which can be triangulated with a number of tetrahedra that is linear in t; here, we are assuming for the sake of simplicity that X, Y_v, and Y_t are all sub-3-manifolds of M' (instead of sub-3-complexes, as they would be in general). The 3-manifold N is naturally homeomorphic to a product interval bundle; by gluing N and Y_v along vertical subsets of their boundaries, we recover the product interval bundle M'.

Recall that the positive and negative push-offs provide simplicial isomorphisms

$$(-)^+ \colon F \longrightarrow \partial_0 M' \qquad \text{and} \qquad (-)^- \colon F \longrightarrow \partial_1 M'.$$

Therefore, in order to recover the monodromy φ, we simply need to understand the map $\partial_0 M' \to \partial_1 M'$ induced by the product interval bundle structure of M'. Since Y_v and N are both vertical in M', this problem reduces to understanding the maps $\partial_0 Y_v \to \partial_1 Y_v$ and $\partial_0 N \to \partial_1 N$. The idea is that Y_v is a union of parallelity components, and therefore its interval bundle structure is fully encoded in the combinatorics of the normal surface F. We have no such guarantee on N, which is the union of the guts and non-vertical parallelity components of M'. However, since N is triangulated with few tetrahedra, we can recover the map $\partial_0 N \to \partial_1 N$ directly; note that this would be impossible for Y_v, since the size of the triangulation of $\partial_h Y_v$ is linear in $w(F)$.

We now expand on this intuitive idea, and describe how to construct the certificate. Recall from our notation in Certificate 1 that F_p, F_g F'_g, F_0, and F_1 are subsurfaces of F such that

$$\mathcal{N}^+(F_p) = Y_v, \quad \mathcal{N}^+(F_g) = \mathcal{N}^-(F'_g) = X, \quad \text{and } \mathcal{N}^+(F_0) \cup \mathcal{N}^-(F_1) = Y_t.$$

As mentioned above, the sub-3-manifold Y_v of M' is a union of parallelity components whose intrinsic interval bundle structure agrees with that of M'. Therefore, the restriction of the monodromy φ to F_p can be understood locally, inside each tetrahedron of M, simply by shifting each normal disc of F_p to the next normal disc, according to the transverse orientation of F. But this is precisely the effect of the transfer map $\triangle$ defined in Section 2.3.5; in other words, we have that $\varphi|_{F_p} = \triangle|_{F_p}$. Since the transfer map can be readily computed for any given normal disc of F, this takes care of the monodromy on F_p.

Let a_0 be an essential curve (or multi-curve) in F. Denote – as we will do in the proof of Proposition 4.11 – by F' the subsurface $F_g \cup F_0$. We can decompose a_0 into the two 1-manifolds $a_0 \cap F_p$ and $a_0 \cap F'$. As discussed in the previous paragraph, we can easily compute $\varphi(a_0 \cap F_p)$. Therefore, the only thing we need to certify is the image of $a_0 \cap F'$ under φ. Let $b_0 = (a_0 \cap F')^+$, which is a 1-manifold in $\partial_0 N$. The idea is then to provide as a certificate a normal surface A in N that is vertical and such that $A \cap \partial_0 N = b_0$. The 1-manifold $b_1 = A \cap \partial_1 N \subseteq \partial_1 N$ is then the image of b_0 under the map $\partial_0 N \to \partial_1 N$; if we let c be the 1-manifold in F such that $c^- = b_1$, we can define

$$a_1 = \triangle(a_0 \cap F_p) \cup c,$$

and this curve will be isotopic to $\varphi(a_0)$.

Of course, this description is overly simplistic, and it sweeps some concrete issues under the rug. Without delving into the full technical details – this is reserved for the next sections – let us highlight these complications, and explain how we deal with them; in fact, the final blueprint for the monodromy certificate will look something like Figure 4.7. The first and most fundamental issue is that the vertical normal surface A should have bounded weight. In particular, we need a result that guarantees the existence, in a product interval bundle N, of a vertical normal surface whose intersection with $\partial_0 N$ is a prescribed normal 1-manifold b_0, and whose weight is bounded in terms of reasonable parameters. This is precisely what we prove in Proposition 4.10, where we show that such a surface A exists, with weight bounded above by a function that is linear in $w(b_0)$ and exponential in the size of the triangulation of N; this is sufficient for our setting, since $w(b_0) \leqslant w(a)$ and the number of tetrahedra of N is linear in t.

The second issue, which we have ignored so far, is that the parallelity components of Y_t had been retriangulated in order to keep the number of tetrahedra of N linear in t. Concretely, this means that we cannot freely transport normal 1-manifolds from F' to $\partial_0 N$ and expect them to remain normal. In fact, the map $F' \to \partial_0 N$ is simplicial on F_g (since the guts X did not undergo any retriangulation), but not on F_0. This is where the fact that Y_t is contained in a union of "tubes" comes into play. Recall from Certificate 1 that there is a subsurface D_0 of F' containing F_0 that is a union of disjoint discs (assuming, for the sake of simplicity, that M is closed); this subsurface is essentially the intersection of $\partial_0 M'$ with the "tubes". Obviously, the curve a_0 can be isotoped to be disjoint from D_0; in Proposition 4.9, we show that this can be done without increasing the weight of a_0 too much. Once a_0 is disjoint from D_0 – and, hence, from F_0 – the 1-manifold $b_0 = (a_0 \cap F')^+$ is guaranteed to be normal in $\partial_0 N$. A similar trick is required to isotope b_1 away from the intersection of the "tubes" with $\partial_1 N$; this ensures that c is a normal 1-manifold in F'_g.

Finally, we did not mention that Proposition 4.10 guarantees the existence of a vertical normal surface with prescribed intersection b_0 with $\partial_0 N$ *only* when no component of b_0 is a curve bounding a disc or a boundary-parallel arc. Since a_0 is essential, no component of b_0 will be a curve bounding a disc, but it is entirely possible that b_0 contains boundary-parallel arcs. To address this problem, we remove the boundary-parallel arcs from b_0 and then add them back to b_1. The feasibility of this procedure relies on the fact that a set of boundary parallel arcs in $\partial_0 N$ can be encoded by very little information, namely its intersection with $\partial\partial_0 N$ and some additional combinatorial data. It is then possible to recover the image of these boundary-parallel arcs under the map $\partial_0 N \to \partial_1 N$ simply by looking at the combinatorial data. This construction is formalised in Section 4.2, where we introduce the notion of a *compressed system of boundary-parallel arcs* and study their combinatorial and topological properties. The subsequent Section 4.3 is devoted to the detailed construction of the monodromy certificate which we have outlined here.

4.2 Compressed systems of boundary-parallel arcs

Definition 4.1 (Compressed system of boundary-parallel arcs)**.** Let $\mathcal{T}$ be a triangulation of a compact orientable surface F. A *compressed system of boundary-parallel arcs* σ in F consists of:

- an orientation of each component of ∂F; this defines, for each edge e of $\mathcal{T}$ contained in ∂F, a preferred assignment of the symbols $e^{(+)}$ and $e^{(-)}$ to the two endpoints of e, such that moving along e from $e^{(-)}$ to $e^{(+)}$ is consistent with the given orientation of ∂F;
- non-negative integers $\sigma^+(e)$ and $\sigma^-(e)$ for each edge e of $\mathcal{T}$ that is contained in ∂F.

We require that, for each component b of ∂F, the equality

$$\sum_{e \subseteq b} \sigma^+(e) = \sum_{e \subseteq b} \sigma^-(e) \qquad (4.1)$$

holds, where the sums range over all edges e of $\mathcal{T}$ that are contained in b. ×

Definition 4.2 (Weight of a compressed system of boundary-parallel arcs)**.** Let $\mathcal{T}$ be a triangulation of a compact orientable surface F, and let σ be a compressed system of boundary-parallel arcs in F. The *weight* of σ is the integer

$$w(\sigma) = \sum_{e \subseteq \partial F} (\sigma^+(e) + \sigma^-(e)). \qquad ×$$

We will use compressed systems of boundary parallel arcs to encode unions of disjoint boundary-parallel arcs in a triangulated surface F. This correspondence will be detailed in Section 4.2.2, but the core idea is simple. After fixing an orientation of ∂F, we can imagine that each component of a system a of boundary-parallel arcs in F runs along ∂F according to its orientation, starting right before the positive endpoint of an edge, and terminating right after the negative endpoint of an edge. The system a is then completely determined by recording, for each edge e of ∂F, two numbers:

- the number $\sigma^+(e)$ of intersections between a and e that occur right before the positive endpoint $e^{(+)}$ of e;
- the number $\sigma^-(e)$ of intersections between a and e that occur right after the negative endpoint $e^{(-)}$ of e.

Condition (4.1) is then a consequence of the fact that each arc of a starts and ends on the same component of ∂F. Importantly, this correspondence only works if no component of a cobounds a disc with a subarc of an edge of ∂F, otherwise our assumption that each component of a starts before a positive endpoint and ends after a negative endpoint would not hold.

We also remark that the definition we gave of $w(\sigma)$ does not exactly reflect the weight of the system of boundary-parallel arcs a; in fact, the weight of σ only counts the number of intersections of a with ∂F. We chose this definition of weight because it directly reflects the amount of information that is required to encode σ; however, the discrepancy between the two weights is bounded in terms of the size of the triangulation of F, as we state below in Remark 4.3.

4.2.1 Normal vector

Let $\mathcal{T}$ be a flapless triangulation of a compact orientable surface F, and let σ be a compressed system of boundary-parallel arcs in F. We can associate a normal vector $\mathbf{v}_\sigma$ to σ as follows. There is a unique way to assign a non-negative integer $\mu_\sigma(e)$ to each edge $e \subseteq \partial F$ such that $\mu_\sigma(e) = 0$ for at least one edge e in each component of ∂F, and moreover

$$\mu_\sigma(e_1) + \sigma^+(e_1) = \mu_\sigma(e_2) + \sigma^-(e_2) \tag{4.2}$$

whenever e_1 and e_2 are two edges in ∂F with $e_1^{(+)} = e_2^{(-)}$.

For each edge $e \subseteq \partial F$, let

$$\rho_\sigma(e^{(+)}) = \mu_\sigma(e) + \sigma^+(e). \tag{4.3}$$

Note that, by (4.2), we also have that

$$\rho_\sigma(e^{(-)}) = \mu_\sigma(e) + \sigma^-(e). \tag{4.4}$$

Extend the function ρ_σ to assign 0 to each vertex v of $\mathcal{T}$ that does not lie on ∂F. We can then define the vector $\mathbf{v}_\sigma$ by setting

$$\mathbf{v}_\sigma(e) = \begin{cases} \sigma^+(e) + \sigma^-(e) & \text{if } e \subseteq \partial F, \\ \rho_\sigma(v_1) + \rho_\sigma(v_2) & \text{if } e \not\subseteq \partial F \text{ and } \partial e = \{v_1, v_2\} \end{cases}$$

for each edge e of $\mathcal{T}$. The vector $\mathbf{v}_\sigma$ is easily seen to verify the matching equations; this can be proved algebraically or, alternatively, by showing it is the normal vector of a normal 1-manifold, as described below in Section 4.2.2.

4.2.2 Topological construction

The construction of $\mathbf{v}_\sigma$ can also be described in a purely topological fashion; this description is conveniently explained using the framework of *train tracks* (see [30]). Construct a train track by drawing, for each vertex v of $\mathcal{T}$ that lies on ∂F, a small properly embedded arc around v, that we call the *link* of v. Then, for each edge e of $\mathcal{T}$ that is contained in ∂F, add a close parallel copy of e as a branch. The result is a trivalent train track: the link of each vertex on ∂F is now split into three branches, and there is one additional branch for each edge in ∂F.

This train track comes equipped with a natural measure induced by σ. The weight of a branch associated to an edge $e \subseteq \partial F$ is $\mu_\sigma(e)$. For each vertex $v \in \partial F$, if we let e_1 and e_2 be the two edges contained in ∂F such that $e_1^{(+)} = e_2^{(-)} = v$, then the weights of the three branches in the link of v are defined as follows (see also Figure 4.1a):

- the branch that does not intersect ∂F has weight $\rho_\sigma(v)$;
- the branch that intersects e_1 near its positive endpoint $e_1^{(+)}$ has weight $\sigma^+(e_1)$;
- the branch that intersects e_2 near its negative endpoint $e_2^{(-)}$ has weight $\sigma^-(e_2)$.

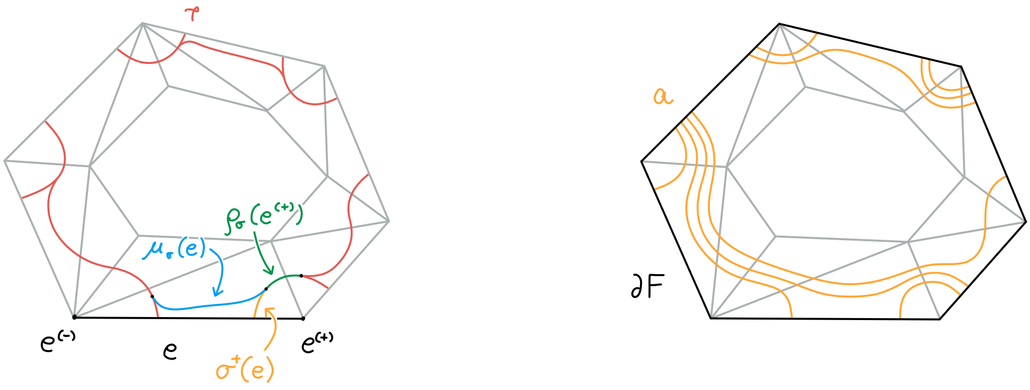

(a) (b)

Figure 4.1. (a) The train track τ used in the topological construction of a topological representative of σ, endowed with the induced measure. (b) The normal 1-manifold a is carried by τ.

It is easy to check that the switch equations are satisfied, thanks to (4.3) and (4.4). Finally, remove all the branches of weight zero to obtain a train track τ in F; this measured train track is depicted in Figure 4.1a.

The measure we described above yields a 1-manifold a carried by τ and properly embedded in F. Since $\mathcal{T}$ is flapless, the train track τ does not form any bigons with the edges of $\mathcal{T}$, and therefore the 1-manifold a is normal, as depicted in Figure 4.1b. Moreover, it is immediate to compute the normal vector of a, and to check that $\mathbf{v}_a = \mathbf{v}_\sigma$; we call such a 1-manifold a *topological representative* of σ. In particular, note that $|a \cap \partial F| = w(\sigma)$.

We remark that the 1-manifold a can be chosen to lie in an arbitrarily small neighbourhood of the union of the edges in ∂F on which μ_σ is non-zero; this will be useful later in the definition of augmentation.

Remark 4.3. From this topological construction, it is easy to derive a bound on the weight of the normal 1-manifold a. In fact, each component of a is a normal arc intersecting each edge of $\mathcal{T}$ at most twice. If we denote by t the number of triangles of $\mathcal{T}$, there are at most $2t$ edges in $\mathcal{T}$ – recall that $\mathcal{T}$ is flapless – and exactly $w(\sigma)/2$ components of a; therefore, we immediately find that

$$w(a) \leqslant 2t \cdot w(\sigma).$$

×

4.2.3 Augmentation

A compressed system of boundary-parallel arcs σ can be used to augment a normal 1-manifold, provided that said 1-manifold satisfies a compatibility condition with σ.

Definition 4.4 (Compatibility of a compressed system of boundary-parallel arcs with a normal 1-manifold)**.** Let $\mathcal{T}$ be a triangulation of a compact orientable surface F, and let a be a normal 1-manifold in F. A compressed system of boundary-parallel arcs σ in F is *compatible* with a if, for each edge e of $\mathcal{T}$ contained in ∂F, at most one of $|a \cap e|$ and $\mu_\sigma(e)$ is non-zero. ×

Definition 4.5 (Augmenting a normal 1-manifold by a compressed system of boundary-parallel arcs)**.** Let $\mathcal{T}$ be a flapless triangulation of a compact orientable surface F, and let a be a normal 1-manifold in F. Let σ be a compressed system of boundary-parallel arcs in F that is compatible with a. An *augmentation* of a by σ is a normal 1-manifold a' in F such that

$$\mathbf{v}_{a'} = \mathbf{v}_a + \mathbf{v}_\sigma.$$ ×

Note that the augmentation is only defined up to normal isotopy in F. However, since we will mostly only care about the normal isotopy class of the result, we will write "the" augmentation to refer to any normal isotopy representative of the augmentation.

Thanks to the topological construction detailed in Section 4.2.2, the augmentation can also be described topologically. Note that, since σ is compatible with a, for every edge $e \subseteq \partial F$ intersecting a we have that $\mu_\sigma(e) = 0$. In particular, as remarked above, this implies that the train track used in the construction of a topological representative of σ – and, as a consequence, the topological representative itself – can be chosen to be disjoint from a. Therefore, the augmentation a' can be realised by taking the union of a and a disjoint topological representative a'' of σ, as depicted in Figure 4.2; since $\mathbf{v}_{a''} = \mathbf{v}_\sigma$ and a'' is disjoint from a, it is clear that $\mathbf{v}_{a'} = \mathbf{v}_a + \mathbf{v}_\sigma$.

Note that, if b is a normal 1-manifold in F that is isotopic to a fixing ∂F, then a topological representative a'' of σ can be chosen to be disjoint from both a and b. As a consequence, the augmentations of a and b by σ can be chosen to be isotopic fixing ∂F. Therefore, we have proved the following.

Proposition 4.6 (Isotopy class relative boundary of augmentations)**.** *Let $\mathcal{T}$ be a flapless triangulation of a compact orientable surface F, and let a and b be normal 1-manifolds in F that are isotopic fixing ∂F. Let σ be a compressed system of boundary-parallel arcs in F that is compatible with a. Let a' be an augmentation of a by σ. Then b is compatible with σ, and there exists an augmentation b' of b by σ that is isotopic to a' fixing ∂F.*

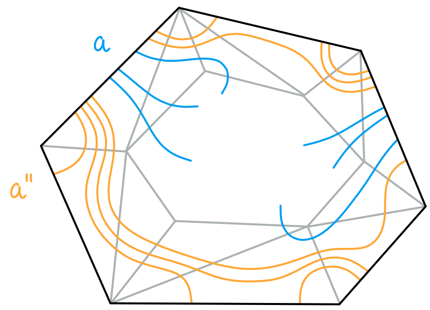

Figure 4.2. The augmentation of a by σ can be realised as the disjoint union of a and a topological representative a'' of σ.

The following two proposition complete our two-step programme to circumvent the issue that boundary-parallel arcs are not allowed in the prescribed boundary of Proposition 4.10. First, we show in Proposition 4.7 that any normal 1-manifold a can be decomposed into an "essential" part, that has no boundary-parallel components, and a boundary-parallel part, that is encoded by a compressed system of boundary-parallel arcs σ. If the 1-manifold a sits in $\partial_0 M$ for some product interval bundle M, then Proposition 4.10 yields a vertical surface A in M with $A \cap \partial_0 M$ being the essential part of a. Then, in Proposition 4.8, we show that the combinatorial data contained in σ can be used to easily compute the image of the boundary-parallel part of a under the map $\partial_0 M \to \partial_1 M$, without the aid of a vertical surface.

Proposition 4.7 (Decomposition of a normal 1-manifold into essential and boundary-parallel components)**.** *Let $\mathcal{T}$ be a flapless triangulation of a compact orientable surface F, and let a be a normal 1-manifold in F. Suppose that no component of a is a curve bounding a disc, or an arc cobounding a disc with a subarc of an edge in ∂F. Then there exist a normal 1-manifold b in F and a compressed system of boundary-parallel arcs σ in F such that:*

(i) no component of b is a curve bounding a disc or a boundary-parallel arc;

(ii) σ is compatible with b;

(iii) there is an augmentation b' of b by σ that is isotopic to a in F fixing ∂F;

(iv) $w(b) \leqslant w(a)$ and $w(\sigma) \leqslant |\partial a|$.

Proof. In the context of this proof, we will say that a disc $D \subseteq F$ is *thin* if it satisfies the following conditions:

- ∂D is the union of two arcs r and s with $\partial r = \partial s$, such that r is a normal arc in F and $s \subseteq \partial F$;
- $D \cap \mathcal{T}^{(0)} \subseteq \partial F$.

Note that every component of a topological representative of a compressed system of boundary-parallel arcs cobounds a thin disc with ∂F; this can be inferred from the construction detailed in Section 4.2.2.

Fix an orientation on each boundary component of F. We now describe a procedure that, at the i-th step (for $i \geqslant 1$), constructs normal 1-manifolds b_i and c_i in F, and a compressed system of boundary-parallel arcs σ_i in F, satisfying the following properties:

(i) $b_i \subseteq a$;

(ii) c_i is a topological representative of σ_i;

(iii) σ_i is compatible with b_i;

(iv) b_i and c_i are disjoint;

(v) $b_i \cup c_i$ is isotopic to a fixing ∂F;

(vi) $w(\sigma_i) + |\partial b_i| = |\partial a|$.

We start by setting $b_0 = a$, $c_0 = \varnothing$, and σ_0 to be the compressed system of boundary-parallel arcs in F such that $\sigma_0^+(e) = \sigma_0^-(e) = 0$ for every edge $e \subseteq \partial F$; the properties listed above are clearly satisfied for $i = 0$. Suppose now that b_i, c_i, and σ_i have been inductively constructed. If no component of b_i is a boundary-parallel arc, then we terminate the procedure. Otherwise, let r be an outermost boundary-parallel arc component of b_i – that is, a component r of b_i that cobounds a disc D with ∂F such that $D \cap b_i = r$. Set $b_{i+1} = b_i \setminus r$.

If some component of c_i that is contained in D cobounds a thin disc E with ∂F such that $r \subseteq E$, then we let $D' = \operatorname{clos}(F \setminus D)$ and $r' = r$, noting that r is disjoint from c_i and b_{i+1}, and that D' is a thin disc. Otherwise, there is an ambient isotopy of F that fixes ∂F, c_i, and b_{i+1}, and that sends D to a thin disc D'; this isotopy is depicted in Figure 4.3. In this case, let $r' = \operatorname{clos}(\partial D' \setminus \partial F)$.

Either way, we have constructed an ambient isotopy of F that fixes ∂F, c_i, and b_{i+1}, and that sends r to a normal arc r' that cobounds a thin disc D' with some subarc s of ∂F. Note that, by assumption, the arc s is not a subarc of an edge of $\mathcal{T}$. There is a unique edge e_1 of $\mathcal{T}$ contained in ∂F such that s does not contain e_1, but it contain $e_1^{(+)}$

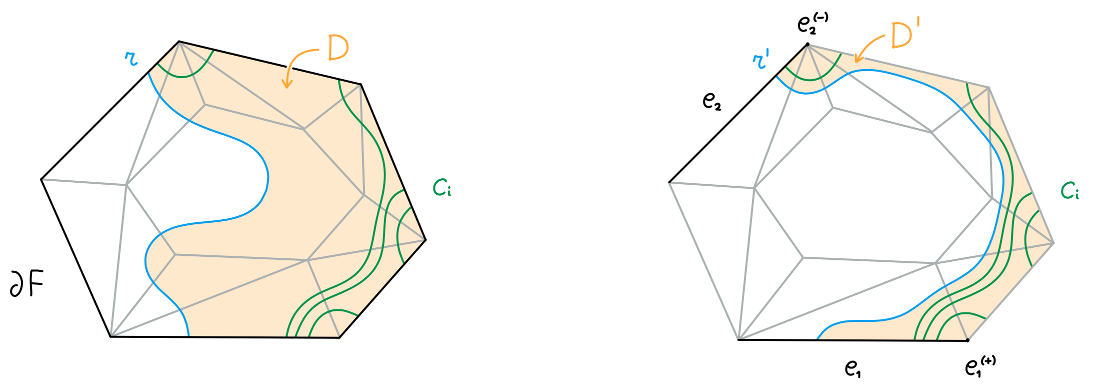

(a) (b)

Figure 4.3. (a) An outermost component r of b_i cobounds a disc D with ∂F. (b) There is an ambient isotopy of F that fixes ∂F, b_{i+1}, and c_i, and that sends r to a normal arc r' that cobounds a thin disc D' with ∂F. Also, the two edges e_1 and e_2 involved in the definition of ν.

(see Figure 4.3b) Similarly, there is a unique edge e_2 of $\mathcal{T}$ contained in ∂F such that s does not contain e_2, but it contains $e_2^{(-)}$.

Let ν be the compressed system of boundary-parallel arcs in F that is constantly 0 except for

$$\nu^+(e_1) = \nu^-(e_2) = 1.$$

Since r' cobounds a thin disc with s, it is easy to see that r' is a topological representative of ν. Therefore, if we let $c_{i+1} = c_i \cup r'$ and $\sigma_{i+1} = \sigma_i + \nu$ (where addition is performed component-wise), then we have that c_{i+1} is a topological representative of σ_{i+1}. The other required properties are easy to check, recalling that b_{i+1} is disjoint from D'.

The procedure terminates after finitely many steps – say n – since the number of components of b_i decreases at each step. Let $b = b_n$, $c = c_n$, and $\sigma = \sigma_n$. The only way for the procedure to terminate is if no component of b is a boundary-parallel arc; moreover, since $b \subseteq a$, we see that no component of b is a curve bounding a disc. We have that σ is compatible with b by construction, and that $b \cup c$ is an augmentation of b by σ, which is also isotopic to a fixing ∂F. Finally, the condition on $w(\sigma)$ follows from

$$w(\sigma) = |\partial a| - |\partial b| \leqslant |\partial a|. \qquad \square$$

Proposition 4.8 (Augmenting a normal surface by a compressed system of boundary-parallel arcs). *Let $\mathcal{T}$ be a suitable pre-sutured triangulation of $M = F \times [0,1]$ for some compact oriented surface F, and let Ω be a simplicial sub-1-manifold of $\partial F \times \{0\}$. Suppose that $\mathcal{T}$ restricts to a flapless triangulation of $\partial_h M$, and to a square-tiled triangulation of $\Omega \times [0,1]$; let $\uparrow\colon \Omega \times \{0\} \to \Omega \times \{1\}$ be the upward shift map induced by this square-tiled triangulation. Let A be a vertical normal surface in M, and denote by a_i the intersection of A with $\partial_i M$ for $i \in \{0,1\}$. Suppose that $\partial a_i \subseteq \Omega \times \{i\}$ for $i \in \{0,1\}$, that $\uparrow(\partial a_0) = \partial a_1$, and that x and $\uparrow(x)$ are endpoints of a component of $A \cap \partial_v M$ for each point $x \in \partial a_0$.*

Let σ_0 be a compressed system of boundary-parallel arcs in $\partial_0 M$ that is compatible with a_0. Suppose that $\sigma_0^+(e) = \sigma_0^-(e) = 0$ for every edge e of $\mathcal{T}$ contained in the closure of $(\partial F \setminus \Omega) \times \{0\}$. Recall that σ_0 comes with a preferred orientation on $\partial F \times \{0\}$; fix on $\partial F \times \{1\}$ the orientation induced by that on $\partial F \times \{0\}$ under the map $(x,0) \mapsto (x,1)$. For each edge e of $\mathcal{T}$ contained in $\Omega \times \{0\}$, let $\sigma_1^+(\uparrow(e)) = \sigma_0^+(e)$ and $\sigma_1^-(\uparrow(e)) = \sigma_0^-(e)$. Let $\sigma_1^+(e) = \sigma_1^-(e) = 0$ for every edge e of $\mathcal{T}$ contained in the closure of $(\partial F \setminus \Omega) \times \{1\}$. Then the following hold.

1. *The datum σ_1 is a compressed system of boundary-parallel arcs in $\partial_1 M$ that is compatible with a_1 and such that $w(\sigma_1) = w(\sigma_0)$.*

2. *Let a_0' be an augmentation of a_0 by σ_0. Then there exist an augmentation a_1' of a_1 by σ_1 such that $\uparrow(\partial a_0') = \partial a_1'$, and a vertical surface A' properly embedded in M such that $A' \cap \partial_i M = a_i'$ for $i \in \{0,1\}$, and x and $\uparrow(x)$ are endpoints of a component of $A' \cap \partial_v M$ for each point $x \in \partial a_0'$.*

Proof. It is clear from the definition that σ_1 is a compressed system of boundary-parallel arcs in $\partial_1 M$ with $w(\sigma_1) = w(\sigma_0)$; it is compatible with a_1 because, if $a_1 \cap e \neq \varnothing$ for some edge e' of $\mathcal{T}$ contained in $\partial F \times \{1\}$, then $e' \subseteq \Omega \times \{1\}$ and

$$a_1 \cap e' = \uparrow(a_0 \cap e) \neq \varnothing,$$

where e is the edge of $\mathcal{T}$ contained in $\partial F \times \{0\}$ such that $e' = \uparrow(e)$; it follows that $\sigma_0(e) = 0$ – since σ_0 is compatible with a_0 – and therefore $\sigma_1(e) = 0$.

Let a_0' be an augmentation of a_0 by σ_0. Up to normal isotopy, we can assume that a_0' is obtained as in Figure 4.2. More precisely, the 1-manifold a_0' is the disjoint union of a_0 and a normal 1-manifold a_0'' that is carried by a train track τ_0, as explained in Section 4.2.2. The same construction applied to a_1 and σ_1 yields an augmentation a_1' of a_1 by σ_1 that is the disjoint union of a_1 and a normal 1-manifold a_1'' carried by a train

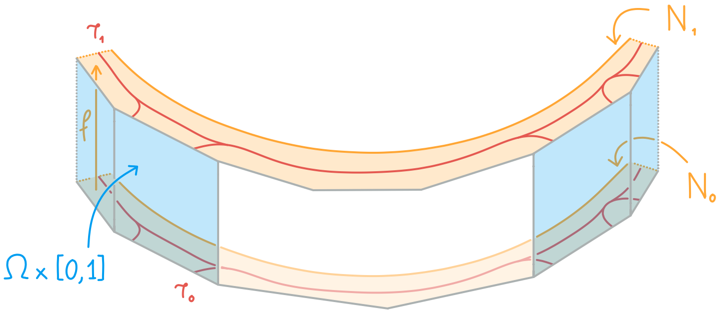

Figure 4.4. There is a homeomorphism $f\colon N_0 \to N_1$ that extends $\uparrow\colon \Omega\times\{0\} \to \Omega\times\{1\}$ and such that $f(\tau_0) = \tau_1$.

track τ_1. It is immediate from the definition of augmentation and of σ_1 that

$$\begin{aligned} |a_1' \cap \uparrow(e)| &= |a_1 \cap \uparrow(e)| + \sigma_1^+(\uparrow(e)) + \sigma_1^-(\uparrow(e)) \\ &= |a_0 \cap e| + \sigma_0^+(e) + \sigma_0^-(e) \\ &= |a_0' \cap e| \end{aligned}$$

for each edge e of $\mathcal{T}$ contained in $\Omega \times \{0\}$. Therefore, up to normal isotopy of a_1', we can assume that $\uparrow(\partial a_0') = \partial a_1'$.

Similarly, since for every edge $e \subseteq \partial F \times \{0\}$, the intersection $\tau_0 \cap e$ only depends on whether $\sigma_0^+(e)$ and $\sigma_0^-(e)$ are zero or not, up to normal isotopy of τ_1 we can assume that

$$\tau_1 \cap (\partial F \times \{1\}) = \uparrow(\tau_0 \cap (\partial F \times \{0\})).$$

In fact, it is not hard to construct neighbourhoods N_i of $\partial F \times \{i\}$ in $F \times \{i\}$ with $\tau_i \subseteq N_i$ for $i \in \{0, 1\}$, and a homeomorphism $f\colon N_0 \to N_1$ that extends $\uparrow$ and such that $f(\tau_0) = \tau_1$; see Figure 4.4. It follows that we can find a vertical branched surface B in M, isotopic to $\tau_0 \times [0, 1]$, such that $B \cap \partial_i M = \tau_i$ for $i \in \{0, 1\}$, and x and $\uparrow(x)$ are endpoints of a component of $B \cap \partial_v M$ for each point $x \in \tau_0 \cap (\partial F \times \{0\})$. This branched surface, depicted in Figure 4.5a, can be chosen to be disjoint from A, precisely because components of $A \cap \partial_v M$ are vertical arcs, each of which connects a point $x \in \partial a_0$ to $\uparrow(x)$.

Recall that τ_0 and τ_1 are endowed with measures induced by σ_0 and σ_1 respectively. Note that, for $i \in \{0, 1\}$, the following is true: for every edge $e \subseteq \Omega \times \{i\}$, the weights of the branches of τ_i that intersect e only depend on $\sigma_i^+(e)$ and $\sigma_i^-(e)$; moreover, the weights of the branches of τ_i that intersect $\Omega \times \{i\}$ completely determine the measure

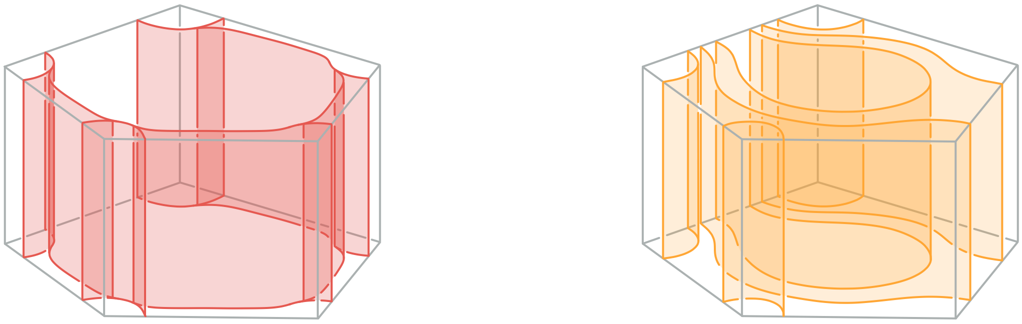

(a)

(b)

Figure 4.5. (a) The vertical branched surface B. (b) The vertical surface A'' carried by B.

on τ_i. Therefore, we see that the homeomorphism f induces a measure-preserving homeomorphism from τ_0 to τ_1. This implies that there is a measure on B that agrees with the measures on τ_0 and τ_1. This measure yields a vertical surface A'' carried by B, that we can assume to be disjoint from A (since B is). Moreover, we can assume that $A'' \cap \partial_i M = a_i''$ for $i \in \{0, 1\}$, and that each component of $A'' \cap \partial_v M$ connects a point $x \in \partial a_0''$ to $\uparrow(x)$; see Figure 4.5b for a depiction of A''. The required surface A' is then simply the union of A and A''. ▢

4.3 A certificate for the monodromy

In this section, we describe in detail our certificate for the image of a curve under the monodromy. Since this certificate is quite technical and involved, we refer the reader to Section 4.1 for an outline of the construction.

Certificate 2 (Certificate for the image of a curve under the monodromy)**.** Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M with t tetrahedra, let F be a transversely oriented connected normal surface in M, and let a_0 and a_1 be normal multicurves in F. Suppose that $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ is given; from now on, the notation and objects defined in Certificate 1 will refer to those of Σ_{fib}. We say that a certificate Σ lies in $\mathfrak{S}_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$ if it consists of:

- a normal multi curve a_0' in F;
- a compressed system of boundary-parallel arcs σ_0 in $\partial_0 N$;

- a normal surface A in N.

We define the *size* of Σ to be the number

$$|\Sigma| = t + \log(w(F)+1) + |\Sigma_{\text{fib}}| + \log(w(a_0)+1) + \operatorname{area}(\operatorname{supp}(a_0)) + \log(w(a_1)+1) + \operatorname{area}(\operatorname{supp}(a_1)) + \log(w(a_0')+1) + \operatorname{area}(\operatorname{supp}(a_0')) + \log(w(\sigma_0)+1) + \log(w(A)+1).$$

Let:

- $\Omega_{[0,1]} = f_g^{-1}(\mathcal{N}^+(F_p))$;
- $\Omega_i = \Omega_{[0,1]} \cap \partial_i N$ for $i \in \{0,1\}$;
- $\uparrow\colon \Omega_0 \to \Omega_1$ be the upward shift map induced by the fact that $\mathcal{R}$ restricts to a square-tiled triangulation of $\Omega_{[0,1]}$;
- $\partial\partial_1 N$ be endowed with the orientation induced by that of σ_0 on $\partial\partial_0 N$, as described in the statement of Proposition 4.8;
- σ_1 be the compressed system of boundary-parallel arcs in $\partial_1 N$ such that

  $$\sigma_1(\uparrow(e)^{(+)}) = \sigma_0(e^{(+)}) \quad \text{and} \quad \sigma_1(\uparrow(e)^{(-)}) = \sigma_0(e^{(-)})$$

  for each edge e of $\mathcal{R}$ contained in Ω_0, and $\sigma_1(e^{(+)}) = \sigma_1(e^{(-)}) = 0$ for each edge e of $\mathcal{R}$ contained in the closure of $\partial\partial_1 N \setminus \Omega_1$;
- $b_0 = f_g^{-1}((a_0')^+)$;
- $b_1 = f_g^{-1}(a_1^-)$;
- $b_i' = A \cap \partial_i N$ for $i \in \{0,1\}$;
- b_i'' be the augmentation of b_i' by σ_i (defined up to normal isotopy) for $i \in \{0,1\}$;
- S_0 be the surface

  $$S_0 = {F_p \sqcup \partial_0 N}\Big/{\{x \sim y : x \in F_p, y \in \partial_0 N \cap N', x^+ = f_g(y)\}},$$

  and let $\iota_0\colon F_p \sqcup \partial_0 N \to S_0$ be the quotient map;

- S_1 be the surface

$$S_1 = {F'_p \sqcup \partial_1 N} \big/ {\{x \sim y : x \in F'_p, y \in \partial_1 N \cap N', x^- = f_g(y)\}},$$

  and let $\iota_1 \colon F'_p \sqcup \partial_1 N \to S_1$ be the quotient map.

We say that Σ is *valid*, and write $\Sigma \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$, if the following conditions are satisfied:

(g.1) no component of a_0 bounds a disc in F;

(g.2) a'_0 is isotopic to a_0 in F;

(g.3) a'_0 is disjoint from F_0;

(g.4) $a_1 \cap F'_p = \triangle(a'_0 \cap F_p)$;

(g.5) a_1 is disjoint from F_1;

(h.1) A is vertical in N;

(h.2) σ_0 is compatible with b'_0;

(h.3) $\sigma_0(e) = 0$ for every edge e of $\mathcal{R}$ contained in the closure of $(\partial\partial_0 N) \setminus \Omega_0$;

(h.4) $|b_0 \cap e| = |b''_0 \cap e|$ for every edge e of $\mathcal{R}$ contained in $\partial\partial_0 N$;

(h.5) no component of b'_0 is a curve bounding a disc in $\partial_0 N$ or a boundary-parallel arc in $\partial_0 N$;

(h.6) $\partial b'_i \subseteq \Omega_i$ for $i \in \{0, 1\}$;

(h.7) $\uparrow(\partial b'_0) = \partial b'_1$;

(h.8) x and $\uparrow(x)$ are endpoints of a component of $A \cap \partial_v N$ for each point $x \in \partial b'_0$;

(i.1) the normal multicurves $\iota_0(a'_0 \cap F_p) \cup \iota_0(b_0)$ and $\iota_0(a'_0 \cap F_p) \cup \iota_0(b''_0)$ are isotopic in S_0;

(i.2) the normal multicurves $\iota_1(a_1 \cap F'_p) \cup \iota_1(b_1)$ and $\iota_1(a_1 \cap F'_p) \cup \iota_1(b''_1)$ are isotopic in S_1;

(j.1) $w(a'_0) \leqslant w(a_0) + 96t \cdot w(a_0 \cap F_g)$ and $\text{area}(\text{supp}(a'_0)) \leqslant \text{area}(\text{supp}(a_0)) + 32t$;

(j.2) $w(a_1) \leqslant w(a_0) + 2^{8426t+63} \cdot w(a_0 \cap F_g)$ and $\text{area}(\text{supp}(a_1)) \leqslant \text{area}(\text{supp}(a_0)) + 32t$;

(j.3) $w(\sigma_0) \leqslant 97t \cdot w(a_0 \cap F_g)$;

(j.4) $w(A) \leqslant 2^{8425t+48} \cdot w(a_0 \cap F_g)$. ×

We briefly justify why the objects introduced in the certificate are well-defined, assuming Σ is valid. At first glance, this might appear circular, since these objects are used to define when Σ is valid. However, it is easy to verify that our justification for each object's well-definedness relies only on properties of Σ that are established before the object itself is introduced.

Well-definedness of $\Omega_{[0,1]}$, Ω_0, Ω_1, and $\uparrow$. We note that $X \cap \mathcal{N}^+(F_p)$ is the intersection of the guts X with a union of parallelity components of M'; therefore, it is triangulated by a square-tiled triangulation as explained in Section 2.3.5. Since f_g is simplicial, and it restricts to an orientation-preserving homeomorphism $\mathrm{int}(N') \to \mathrm{int}(X)$, the triangulation $\mathcal{R}$ of N also restricts to a square-tiled triangulation of $\Omega_{[0,1]}$. Moreover, by property (a.5), we have that $\Omega_{[0,1]}$ is a vertical subset of $\partial_v N$; this justifies the existence of the upward shift map $\uparrow$.

Well-definedness of b_0 and b_1. By property (g.3), the normal curve a_0' is disjoint from F_0. As a consequence, we see that b_0 is a normal 1-manifold in $\partial_0 N \cap N'$ that is disjoint from N''; in particular, we have that $\partial b_0 \subseteq \Omega_0$. Similarly, b_1 is a normal 1-manifold in $\partial_1 N \cap N'$ that is disjoint from N''; in particular, we have that $\partial b_1 \subseteq \Omega_1$.

Well-definedness of σ_1 and b_1'. The compressed system of boundary-parallel arcs σ_1 is constructed as in Proposition 4.8, in a way that will be made more explicit in the proof of Proposition 4.11 below. In particular, by the first conclusion of Proposition 4.8, we have that σ_1 is effectively a compressed system of boundary-parallel arcs in $\partial_1 N$, and it is compatible with b_1'. This implies that the augmentation b_1'' of b_1' by σ_1 is well-defined (up to normal isotopy). Also, the second conclusion of Proposition 4.8 ensures that we can pick b_1'' so that $\uparrow(\partial b_0'') = \partial b_1''$ (recall that b_0'' is an augmentation of b_0' by σ_0).

Well-definedness of $\iota_0(a_0' \cap F_p) \cup \iota_0(b_0)$ and $\iota_0(a_0' \cap F_p) \cup \iota_0(b_0'')$. Recall that $b_0 = f_g^{-1}((a_0')^+)$; in particular, we have that $|a_0' \cap e| = |b_0 \cap e'|$ whenever e is an edge of $F_p \cap F_g$ and e' is an edge of $\partial_0 N \cap N'$ with $f_g(e') = e^+$. This implies that $\iota_0(a_0' \cap F_p)$ and $\iota_0(b_0)$ glue up in S_0 to give a normal multicurve. By property (h.4), we have that

$$|a_0' \cap e| = |b_0 \cap e'| = |b_0'' \cap e'| \tag{4.5}$$

whenever e is an edge of $F_p \cap F_g$ and e' is an edge of $\partial_0 N \cap N'$ with $f_g(e') = e^+$. Therefore, the 1-manifolds $\iota_0(a_0' \cap F_p)$ and $\iota_0(b_0'')$ glue up in S_1 to give a normal multicurve.

Well-definedness of $\iota_1(a_1 \cap F'_p) \cup \iota_1(b_1)$ and $\iota_1(a_1 \cap F'_p) \cup \iota_1(b''_1)$. Similarly to above, since $b_1 = f_g^{-1}(a_1^-)$, we have that $\iota_1(a_1 \cap F'_p)$ and $\iota_1(b_1)$ glue up in S_1 to give a normal multicurve. Note that property (g.4) implies that $|b_0 \cap e| = |b_1 \cap \uparrow(e)|$ for every edge e of $\mathcal{R}$ contained in Ω_0. Moreover, by the second conclusion of Proposition 4.8, we have that $|b''_0 \cap e| = |b''_1 \cap \uparrow(e)|$ for every edge e of $\mathcal{R}$ contained in Ω_0. From these two observations, it follows that b_1 intersects each edge of Ω_1 as many times as b''_1 does. We conclude that $\iota_1(a_1 \cap F'_p)$ and $\iota_1(b''_1)$ glue up in S_1 to give a normal multicurve.

4.3.1 Existence

We now prove that every essential normal multicurve a_0 in F admits a valid Certificate 2 for some normal multicurve $a_1 \subseteq F$. We start with an elementary result, showing that 1-manifolds can be isotoped away from a collection of discs; we invite the reader to the discussion in Section 4.1 to understand why this is necessary.

Proposition 4.9 (Avoiding forbidden discs). *Let $\mathcal{T}$ be a triangulation of a compact orientable surface F, and let D be a sub-2-complex of F such that* thick(D) *is a disjoint union of discs, each of which intersects ∂F in a single (possibly empty) arc. Let a be a normal 1-manifold in F such that ∂a is disjoint from D. Suppose that no component of a is a curve bounding a disc or an arc cobounding a disc with a subarc of an edge of $\mathcal{T}$. Then there is a normal 1-manifold a' in F such that:*

(i) a' is isotopic to a fixing ∂F;

(ii) a' is disjoint from D;

(iii) the weight of a' satisfies the bound

$$w(a') \leqslant w(a \setminus D) + 3|a \cap \partial D| \cdot \mathrm{area}(\mathrm{clos}(F \setminus D)).$$

Proof. Let $F' = \mathrm{clos}(F \setminus D)$, which is a sub-2-complex of F. We now describe an inductive procedure that, at each step, reduces the number of components of $a \cap D$, while increasing the weight of a by at most a constant amount. Start by setting $a_0 = a$. At the i-th step, for $i \geqslant 0$, consider the general position 1-manifold a_i properly embedded in F. If a_i is disjoint from D, then set $a'' = a_i$ and stop the procedure. Otherwise, there is a choice of thick(D) such that each component of $a_i \cap$ thick(D) intersects D. Pick a component b of $a_i \cap$ thick(D) that is outermost in thick(D), in the sense that it cobounds a disc E with a subarc of ∂_F thick(D) such that the interior of E is disjoint from a_i

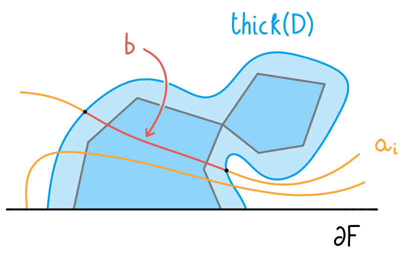

(a)

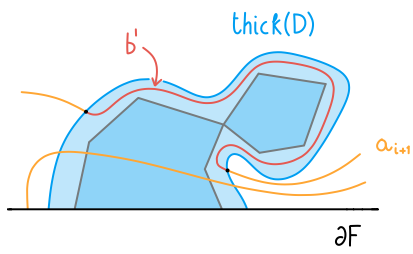

(b)

Figure 4.6. (a) The arc b is a component of $a_i \cap \text{thick}(D)$ that is outermost in $\text{thick}(D)$. (b) The arc b can be isotoped in $\text{thick}(D)$ fixing ∂b to an arc b' in $\text{thick}(D) \setminus D$.

(see Figure 4.6a). This component b is an arc properly embedded in $\text{thick}(D)$, whose endpoints lie on $\partial_F \text{thick}(D)$. Then the arc b can be isotoped in $\text{thick}(D)$, fixing ∂b, to an arc $b' \subseteq \text{thick}(D) \setminus D$, such that $\text{int}(b')$ is disjoint from a_i; this is depicted in Figure 4.6b. Note that we can – and will – choose b' so that it intersects each edge of F' at most twice. Set $a_{i+1} = (a_i \setminus b) \cup b'$. We note that a_{i+1} is isotopic to a_i fixing ∂F, and that $|a_{i+1} \cap \partial D| \leqslant |a_i \cap \partial D| - 2$. Moreover, since b' intersects each edge of F' at most twice, and there are at most $3\,\text{area}(F')$ edges in F', we immediately get the bound

$$w(a_{i+1} \setminus D) \leqslant w(a_i \setminus D) + 2 \cdot 3\,\text{area}(F'). \tag{4.6}$$

Since the quantity $|a_i \cap \partial D|$ decreases by at least 2 at each step, the procedure will eventually terminate. The final 1-manifold a'' is in general position and properly embedded in F. Moreover, it is isotopic to a fixing ∂F, and it is disjoint from D. Its weight satisfies the bound

$$\begin{aligned} w(a'') &= w(a'' \setminus D) && \text{since } a'' \text{ is disjoint from } D \\ &\leqslant w(a \setminus D) + \frac{|a \cap \partial D|}{2} \cdot 6\,\text{area}(F') && \text{by (4.6)} \\ &= w(a \setminus D) + 3|a \cap \partial D| \cdot \text{area}(F'). \end{aligned}$$

Since a'' is isotopic to a fixing ∂F, we see that no component of a'' is a curve bounding a disc in F', or an arc cobounding a disc in F' with a subarc of an edge of $\mathcal{T}$. Therefore, an application of Proposition 2.3 to the 1-manifold a'' in the surface $\text{ab}(F')$ yields the desired normal 1-manifold a'. ▢

The following proposition is the cornerstone of our monodromy certificate. Despite the statement being very natural, and probably unsurprising to experts, the proof is more involved than one might expect. In fact, it requires introducing additional technical concepts, and revising classical results of normal surface theory to introduce small generalisations; it does not contain any novel ideas or clever arguments. Therefore, we have chosen to defer the proof to Appendix A, so as not to disrupt the flow of the main text.

Proposition 4.10 (Vertical surfaces with prescribed boundary)**.** *Let $\mathcal{T}$ be a suitable pre-sutured triangulation of $M = F \times [0,1]$ with t tetrahedra, where F is a compact orientable surface. Let a be a normal 1-manifold in $\partial_0 M \cup \partial_v M$. Denote by a_0 and a_v the intersections of a with $\partial_0 M$ and $\partial_v M$ respectively. Suppose that no component of a_0, seen as a 1-manifold in $\partial_0 M$, is a boundary-parallel arc or a curve bounding a disc; moreover, suppose that a_v is vertical in $\partial_v M$. Then there is a normal surface A in M such that:*

(i) A is a vertical surface in M;

(ii) $A \cap (\partial_0 M \cup \partial_v M)$ is isotopic to a in $\partial_0 M \cup \partial_v M$, fixing $\partial\partial_h M$;

(iii) the weight of A is bounded by

$$w(A) \leqslant 2^{9t+40} \cdot w(a).$$

We are finally ready to describe the construction of the monodromy certificate. As an aid to the reader, we have organised all the objects involved in this construction in Figure 4.7; we recommend referring to it while reading the proof of Proposition 4.11 below.

Proposition 4.11 (Existence of Certificate 2)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented fibred 3-manifold M, and let F be a transversely oriented connected normal surface M, with a given $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M,F)$. Let a_0 be a normal multicurve in F, no component of which bounds a disc in F. Then there is a normal multicurve a_1 in F such that $\mathfrak{S}^*_{\text{mon}}(M,F,\Sigma_{\text{fib}},a_0,a_1)$ is non-empty.*

Proof. Denote by $\mathcal{T}'$ the triangulation of F.

Constructing c, c', c'', and a'_0. Let $F' = F_g \cup F_0$. By Proposition 2.3, there is a normal multicurve c in F satisfying the following properties:

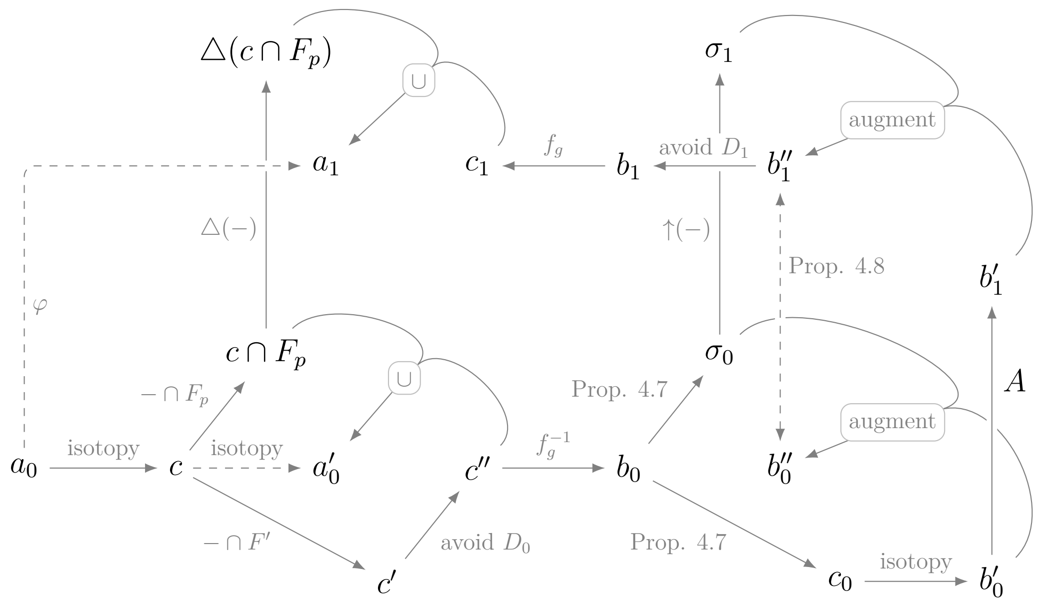

Figure 4.7. A useful diagram to follow the construction of the many objects involved in the proof of Proposition 4.11. The arrow from a_0 to a_1 labelled φ has no relevance in this construction, but it foreshadows the fact, proved in Proposition 4.12, that the curve a_1 is isotopic to the image of a_0 under the monodromy φ of F.

- c is isotopic to a_0 in F;
- $|c \cap e| \leqslant |a_0 \cap e|$ for each edge e of $\mathcal{T}'$;
- there is no general position multicurve b in F isotopic to a_0 such that $|b \cap e| \leqslant |c \cap e|$ for each edge e of $\mathcal{T}'$ and $w(b) < w(c)$.

Consider the normal 1-manifold $c' = c \cap F'$ in $\mathrm{ab}(F')$. Clearly, no component of c' is a curve bounding a disc in $\mathrm{ab}(F')$. Moreover, we claim that no component of c' is an arc cobounding a disc with a subarc of an edge in $\partial\,\mathrm{ab}(F')$. In fact, if this were the case – say that a component of c' cobounds a disc with a subarc of an edge $e \subseteq \partial\,\mathrm{ab}(F')$ – then we could use this disc to construct an isotopy of c to a general position multicurve b in F with $b \cap \mathcal{T}'^{(1)} \subseteq c \cap \mathcal{T}'^{(1)}$ and $|b \cap e| < |c \cap e|$, contradicting the minimality of c.

We can then apply Proposition 4.9 to the 1-manifold $c' \subseteq \mathrm{ab}(F')$ with forbidden region $\mathrm{thick}(D_0)$. The resulting 1-manifold $c'' \subseteq \mathrm{ab}(F')$ is isotopic to c' in $\mathrm{ab}(F')$ fixing $F_g \cap F_p$, and it is disjoint from D_0 – and, hence, from F_0. In particular, note that no component of c'' is a curve bounding a disc in $\mathrm{ab}(F')$ or an arc cobounding a disc with a subarc of an edge in $\partial\,\mathrm{ab}(F')$. Then, we take $a_0' = (c \cap F_p) \cup c''$; since a_0 is isotopic to c, we easily see that it is isotopic to a_0' as well. This guarantees that properties (g.2) and (g.3) are satisfied.

Constructing c_0 and σ_0. Note that $b_0 = f_g^{-1}((c'')^+)$. By Proposition 3.11, there is a map $\partial_0 N \to F_g \cup F_0^+$ that agrees with f_g on $\partial_0 N \cap N'$, and restricts to a homeomorphism on the interior of $\partial_0 N$. As a consequence, no component of b_0 is a curve bounding a disc in $\partial_0 N$ or an arc cobounding a disc with a subarc of an edge in $\partial\partial_0 N$.

We can then apply Proposition 4.7 to obtain a normal 1-manifold c_0 in $\partial_0 N$ and a compressed system of boundary-parallel arcs σ_0, such that:

- no component of c_0 is a curve bounding a disc or a boundary-parallel arc;
- σ_0 is compatible with c_0;
- the augmentation of c_0 by σ_0 is isotopic to b_0 fixing Ω_0.

Recall that $\partial b_0 \subseteq \Omega_0$; this immediately implies that property (h.3) holds, and moreover that $\partial c_0 \subseteq \Omega_0$.

Constructing c_v and A. We construct a normal 1-manifold c_v in $\partial_v N$ as follows. For each point $x \in \partial c_0$, we take a normal arc in $\Omega_{[0,1]}$ that intersects the 1-skeleton of $\mathcal{R}$ exactly three times, with endpoints x and $\uparrow(x)$; note that this is possible because $\Omega_{[0,1]}$ is square-tiled. We call the union of these $|\partial c_0|$ arcs c_v, noting that they can be chosen to be pairwise disjoint.

By construction, no component of c_0 is a curve bounding a disc or a boundary-parallel arc. Therefore, we can apply Proposition 4.10 to the 1-manifold $c_0 \cup c_v \subseteq \partial N$, to obtain a vertical normal surface A in N such that $A \cap \partial_0 N$ is isotopic to c_0 fixing Ω_0, and $A \cap \partial_v N$ is isotopic to c_v fixing $\partial\partial_h N$. Recalling that $b_i' = A \cap \partial_i N$ for $i \in \{0, 1\}$, it is immediate to check that properties (h.1) and (h.5) to (h.8) are satisfied.

Since $\partial c_0 = \partial b_0'$ and σ_0 is compatible with c_0, we deduce that it is also compatible with b_0' (property (h.2)). Moreover, the 1-manifolds c_0 and b_0' are isotopic in $\partial_0 N$ fixing Ω_0; by Proposition 4.6, this implies that the two augmentations b_0'' and b_0 are isotopic in $\partial_0 N$ fixing Ω_0 (in particular, note that $\partial b_0'' \subseteq \Omega_0$, and that property (h.4) holds). By means of the quotient map ι_0, this isotopy can be translated to an isotopy of S_0 that is the identity outside $\iota(\partial_0 N)$; the isotopy witnesses that property (i.1) holds.

Constructing b_1, c_1 and a_1. We want to apply Proposition 4.9 to the normal 1-manifold $b_1'' \subseteq \partial_1 N$ with forbidden region thick(D_1). By Proposition 4.8, we see that if a component of b_1'' were a curve bounding a disc, or an arc cobounding a disc with a subarc of an edge in $\partial\partial_1 N$, then there would be a component of b_0'' that is a curve bounding a disc or an arc cobounding a disc with a subarc of an edge in $\partial\partial_0 N$; this is impossible, since b_0'' is isotopic to b_0 in $\partial_0 N$ fixing Ω_0. Additionally, note that D_1 is disjoint from Ω_1 by property (d.6). Therefore, we can apply Proposition 4.9 to obtain a normal 1-manifold b_1 in $\partial_1 N$ that is isotopic to b_1'' fixing Ω_1 and is disjoint from D_1. Similarly to above, this isotopy can be translated to an isotopy of S_1 that witnesses property (i.2).

Since b_1 is disjoint from D_1, by property (d.5) we deduce that it lies in the domain of f_g; in particular, there exists a normal 1-manifold c_1 in F_g' such that

$$c_1^- = f_g(b_1).$$

Note that, for every edge e of $F_g \cap F_p$, we have that

$$|a_0' \cap e| = |b_0 \cap e'| = |b_0'' \cap e'| = |b_1'' \cap \uparrow(e')| = |b_1 \cap \uparrow(e')| = |c_1 \cap \triangle(e)|,$$

where e' is the edge of $\mathcal{R}$ such that $f_g(e') = e^+$. Therefore,

$$a_1 = \triangle(a_0' \cap F_p) \cup c_1$$

is a normal multicurve in F. It satisfies property (g.4) by definition, and property (g.5) because $f_g(b_1)$ is disjoint from F_1^-.

Quantitative bounds. Let t be the number of tetrahedra of $\mathcal{T}$. We bound the complexities of all the objects we constructed. Recall that $|c \cap e| \leqslant |a_0 \cap e|$ for each edge e of $\mathcal{T}$ (in particular, we have that $\mathrm{supp}(c) \subseteq \mathrm{supp}(a_0)$). Therefore, when we apply Proposition 4.9 to obtain c'', we have the bound

$$\begin{aligned}
w(c'') &\leqslant w(c' \setminus D_0) + 3|c' \cap \partial D_0| \cdot \mathrm{area}(\mathrm{clos}(F' \setminus D_0)) && \text{by Proposition 4.9} \\
&\leqslant w(c' \setminus D_0) + 3|c' \cap \partial D_0| \cdot \mathrm{area}(F_g) && \text{since } F' \setminus D_0 \subseteq F_g \\
&\leqslant w(c' \setminus D_0) + 96t \cdot |c' \cap \partial D_0| && \text{by property 8 of Proposition 3.6} \\
&= w(c' \setminus D_0) + 96t \cdot |c \cap \partial D_0| && \text{since } c' = c \cap F' \\
&\leqslant w(c' \setminus D_0) + 96t \cdot w(c \cap F_g) && \text{since } c \cap \partial D_0 \subseteq F_g
\end{aligned}$$

$$\leqslant w(c' \setminus D_0) + 96t \cdot w(a_0 \cap F_g) \qquad \text{by minimality of } c. \tag{4.7}$$

We deduce that

$$\begin{aligned}
w(a_0') &= w(c \setminus F') + w(c'') && \text{since } a_0' = (c \setminus F') \cup c'' \\
&\leqslant w(c \setminus F') + w(c' \setminus D_0) + 96t \cdot w(a_0 \cap F_g) && \text{by (4.7)} \\
&= w(c \setminus D_0) + 96t \cdot w(a_0 \cap F_g) && \text{since } c' = c \cap F' \\
&\leqslant w(c) + 96t \cdot w(a_0 \cap F_g) \\
&\leqslant w(a_0) + 96t \cdot w(a_0 \cap F_g) && \text{by minimality of } c
\end{aligned} \tag{4.8}$$

and

$$\begin{aligned}
\operatorname{area}(\operatorname{supp}(a_0')) &\leqslant \operatorname{area}(\operatorname{supp}(c \cap F_p)) + \operatorname{area}(\operatorname{supp}(c'')) && \text{since } a_0' = (c \cap F_p) \cup c'' \\
&\leqslant \operatorname{area}(\operatorname{supp}(c)) + \operatorname{area}(F_g) && \text{since } c'' \subseteq F_g \\
&\leqslant \operatorname{area}(\operatorname{supp}(a_0)) + 32t && \text{by property 8 of Proposition 3.6} \\
& && \text{and minimality of } c.
\end{aligned}$$

In particular, we see that property (j.1) holds.

Since $b_0 = f_g^{-1}((c'')^+)$, we have that

$$\begin{aligned}
w(b_0) &= w(c'') && \text{since } f_g \text{ is simplicial} \\
&\leqslant w(c' \setminus D_0) + 96t \cdot w(a_0 \cap F_g) && \text{by (4.7)} \\
&\leqslant w(c \cap F_g) + 96t \cdot w(a_0 \cap F_g) && \text{since } c' = c \cap F' \\
&\leqslant 97t \cdot w(a_0 \cap F_g) && \text{by minimality of } c.
\end{aligned} \tag{4.9}$$

Recall that c_0 and σ_0 are obtained from b_0 by means of Proposition 4.7. As a consequence, we have the bounds

$$w(c_0) \leqslant w(b_0) \leqslant 97t \cdot w(a_0 \cap F_g), \tag{4.10}$$

$$w(\sigma_0) \leqslant |\partial b_0| \leqslant w(b_0) \leqslant 97t \cdot w(a_0 \cap F_g). \tag{4.11}$$

In particular, note that property (j.3) is satisfied.

Moving on to c_v, it is clear from the definition that

$$w(c_0 \cup c_v) = w(c_0) + 2|\partial c_0| \leqslant 3w(c_0) \leqslant 291t \cdot w(a_0 \cap F_g). \tag{4.12}$$

The normal surface A is constructed from $c_0 \cup c_v$ by means of Proposition 4.10. Recalling that $\mathcal{R}$ has at most $936t$ tetrahedra by property (e.2), we deduce that

$$\begin{aligned} w(A) &\leqslant 2^{9\cdot 936t+40} \cdot w(c_0 \cup c_v) && \text{by Proposition 4.10} \\ &\leqslant 2^{8425t+48} \cdot w(a_0 \cap F_g) && \text{by (4.12).} \end{aligned}$$

This guarantees that property (j.4) is satisfied.

Finally, we need to give an upper bound for the weight of a_1. We start by noting that $w(b_1') \leqslant w(A)$ (since $b_1' = A \cap \partial_1 N$) and $w(\sigma_1) = w(\sigma_0)$ (by the first conclusion of Proposition 4.8). Therefore, we can bound the weight of b_1'' by

$$\begin{aligned} w(b_1'') &\leqslant w(b_1') + 2\,\mathrm{area}(\partial_1 N) \cdot w(\sigma_1) && \text{by Remark 4.3} \\ &\leqslant w(A) + 2\,\mathrm{area}(\partial_1 N) \cdot w(\sigma_0) \\ &\leqslant (2^{8425t+48} + 2 \cdot 2808t \cdot 97t) \cdot w(a_0 \cap F_g) && \text{by property 9 of Proposition 3.6 and (4.11)} \\ &\leqslant 2^{8425t+49} \cdot w(a_0 \cap F_g) && \text{by elementary algebra.} \end{aligned} \tag{4.13}$$

Recall that b_1 is obtained from b_1'' by means of Proposition 4.9. Therefore, we have the bound

$$\begin{aligned} w(b_1) &\leqslant w(b_1'' \setminus D_1) + 3|b_1'' \cap \partial D_1| \cdot \mathrm{area}(\mathrm{clos}(\partial_1 N \setminus D_1)) && \text{by Proposition 4.9} \\ &\leqslant w(b_1'') + 3w(b_1'') \cdot \mathrm{area}(\partial_1 N) \\ &\leqslant 8425t \cdot w(b_1'') && \text{by property 9 of Proposition 3.6} \\ &\leqslant 2^{8426t+62} \cdot w(a_0 \cap F_g) && \text{by (4.13).} \end{aligned} \tag{4.14}$$

Since $f_g(b_1) = c_1^-$ and f_g is simplicial, we have that $w(c_1) = w(b_1)$. Finally, since $a_1 = \triangle(a_0' \cap F_p) \cup c_1$ and $\triangle$ is a simplicial isomorphism, we conclude that

$$\begin{aligned} w(a_1) &\leqslant w(\triangle(a_0' \cap F_p)) + w(c_1) \\ &= w(a_0' \cap F_p) + w(b_1) \\ &\leqslant w(a_0') + w(b_1) \\ &\leqslant w(a_0) + (96t + 2^{8426t+62}) \cdot w(a_0 \cap F_g) && \text{by (4.8) and (4.14)} \\ &\leqslant w(a_0) + 2^{8426t+63} \cdot w(a_0 \cap F_g) && \text{by elementary algebra} \end{aligned}$$

and

$$\begin{aligned}
\operatorname{area}(\operatorname{supp}(a_1)) &\leqslant \operatorname{area}(\operatorname{supp}(\triangle(a_0' \cap F_p))) + \operatorname{area}(\operatorname{supp}(c_1)) && \\
&\leqslant \operatorname{area}(\operatorname{supp}(a_0' \cap F_p)) + \operatorname{area}(F_g') && \text{since } c_1 \subseteq F_g' \\
&\leqslant \operatorname{area}(\operatorname{supp}(c)) + 32t && \text{by Proposition 2.10} \\
&\leqslant \operatorname{area}(\operatorname{supp}(a_0)) + 32t && \text{by minimality of } c,
\end{aligned}$$

as required by property (j.2). ▢

4.3.2 Correctness

We prove that, whenever two normal multicurves a_0 and a_1 in F admit a certificate in $\mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$, then a_1 is isotopic to the image of a_0 under the monodromy of F.

Proposition 4.12 (Correctness of Certificate 2)**.** *Let M be a triangulated compact connected oriented 3-manifold, let F be a transversely oriented connected normal surface in M, and let $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$. Let a_0 and a_1 be normal multicurves in F, and let $\Sigma \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$ be a certificate. Recall that M' is homeomorphic to $F \times [0, 1]$ as a pre-sutured manifold by Proposition 3.11. Then there is a vertical surface A' properly embedded in M' such that $A' \cap \partial_h M' = a_0^+ \cup a_1^-$.*

Proof. Let $M_N = X \cup \mathcal{N}^+(F_0) \cup \mathcal{N}^+(F_1)$, with the pre-sutured manifold structure defined in Proposition 3.11. By Proposition 3.11, the simplicial map $f_g \colon N' \to X$ extends to a map $f \colon N \to M_N$ of pre-sutured manifolds, such that f restricts to a homeomorphism $\operatorname{int}(N) \to \operatorname{int}(M_N)$. Let $f^+ \colon \partial_0 N \to F_g \cup F_0$ be the map defined by $(f^+(x))^+ = f(x)$. By (4.5), it follows that

$$a_0'' = (a_0' \cap F_p) \cup f^+(b_0'')$$

is a multicurve in F; we claim that a_0'' is isotopic to a_0 in F. Define a homeomorphism $g \colon F \to S_0$ by

$$g|_{F_p} = \iota_0|_{F_p} \quad \text{and} \quad g|_{F_g \cup F_0} \circ f^+ = \iota_0|_{\partial_0 N}.$$

It is immediate from the definitions that

$$g(a_0') = \iota_0(a_0' \cap F_p) \cup \iota_0(b_0) \quad \text{and} \quad g(a_0'') = \iota_0(a_0' \cap F_p) \cup \iota_0(b_0'').$$

From property (i.1), we deduce that $g(a_0')$ and $g(a_0'')$ are isotopic in S_0, and hence that

a'_0 and a''_0 are isotopic in F. Property (g.2) then implies that a''_0 is isotopic to a_0 in F.

Similarly, let $f^-: \partial_1 N \to F'_g \cup F_1$ be the map defined by $(f^-(x))^- = f(x)$. As in the previous paragraph, we find that

$$a'_1 = (a_1 \cap F'_p) \cup f^-(b'_1)$$

is isotopic to a_1 in F.

Finally, we show how to construct a properly embedded vertical surface A'' in M' such that $A'' \cap \partial_h M' = (a''_0)^+ \cup (a'_1)^-$; since a''_0 is isotopic to a_0 and a'_1 is isotopic to a_1, this will complete the proof. The surface A'' will be the union of two surfaces A_p and A_N. The surface A_p is a vertical surface in $\mathcal{N}^+(F_p)$ such that $A_p \cap \partial_0 M' = (a'_0 \cap F_p)^+$ and $A_p \cap \partial_1 M' = (a_1 \cap F'_p)^-$, and moreover such that x^+ and $(\triangle(x))^-$ are endpoints of a component of $A_p \cap M_N$ for each point $x \in a'_0 \cap F_p \cap F_g$; such a surface exists by property (g.4) and property 2 of Proposition 3.6.

To construct A_N, we first show that we are in a position to apply Proposition 4.8. In fact, we have that $\mathcal{R}$ is a suitable pre-sutured triangulation of $N \cong G \times [0,1]$, and Ω_0 is a simplicial sub-1-manifold of $\partial\partial_0 N$. We know that the restriction of $\mathcal{R}$ to $\partial_h N$ is flapless by property (b.3), and that the restriction of $\mathcal{R}$ to $\Omega_{[0,1]}$ is square-tiled. The normal surface A is vertical in N by property (h.1), and we have that $b''_i = A_i \cap \partial_i N$ for $i \in \{0,1\}$. Properties (h.2), (h.3), and (h.6) to (h.8) ensure that the additional conditions of Proposition 4.8 are satisfied. Therefore, we can conclude that there is a vertical surface A''' properly embedded in N such that $A''' \cap \partial_i N = b''_i$ for $i \in \{0,1\}$, and x and $\uparrow(x)$ are endpoints of a component of $A''' \cap \partial_v N$ for each point $x \in \partial b''_0$. In fact, up to isotopy, we can assume that $f_g(A''' \cap \partial_v N) = A_p \cap M_N$.

In conclusion, we set $A_N = f(A''')$. The surfaces A_p and A_N glue along their common boundary in $\mathcal{N}^+(F_p) \cap M_N$ to give a properly embedded surface A'' in M'. From the construction, it follows that $A'' \cap \partial_h M' = (a''_0)^+ \cup (a'_1)^-$. Moreover, we have that A'' is vertical in M' by Proposition 3.11, since A_p is vertical in $\mathcal{N}^+(F_p)$ and A_N is vertical in M_N. ▢

4.3.3 Verification

We prove that the validity of a certificate $\Sigma \in \mathfrak{S}_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$ can be verified in polynomial time in the size of the certificate and the Euler characteristic of F.

Proposition 4.13 (Verification of Certificate 2). *There is an algorithm that takes as input*

- *a triangulation of a compact connected oriented* 3*-manifold* M*,*
- *a transversely oriented connected normal surface* F *in* M*,*
- *a valid certificate* $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$*,*
- *normal multicurves* a_0 *and* a_1 *in* F*, and*
- *a certificate* $\Sigma \in \mathfrak{S}_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$*,*

and decides whether $\Sigma \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$*. The running time of the algorithm is polynomial in* $|\Sigma|$ *and* $|\chi(F)|$*.*

Proof. As usual, we will refrain from discussing the properties that can be trivially verified by direct inspection or by a straightforward application of the algorithms listed in Section 2.4. Moreover, we remark that all the auxiliary objects defined in Certificate 2, except for S_0, ι_0, S_1, and ι_1, can be explicitly constructed in polynomial time in $|\Sigma|$.

Properties (g.). Property (g.1) can be verified as follows: we first apply Proposition 2.19 to retriangulate F with a number of triangles that is polynomial in $|\chi(F)|$ and $\text{area}(\text{supp}(a_0))$, exactly as we did in the proof of Proposition 2.20; then, we use Proposition 2.13 (applied to the new triangulation of F) to obtain the list of components of a_0; finally, we check that none of these components is a curve bounding a disc in F, by means of Proposition 2.14. Verifying property (g.2) is a straightforward application of Proposition 2.20. Properties (g.3) and (g.5) can be easily verified by means of Proposition 2.18. Property (g.4) can be checked by direct inspection.

Properties (h.). Properties (h.1) and (h.5) can be verified simultaneously thanks to Proposition 3.5. Properties (h.2) to (h.4), (h.6), and (h.7) can be checked by direct inspection. Finally, in order to verify property (h.8), we perform the following two steps. First, we use the algorithm of Proposition 2.13 to check that all components of $A \cap \partial_v N$ are arcs with one endpoint on $\partial_0 N$ and the other on $\partial_1 N$. Then, assuming that $A \cap \partial_v N$ is non-empty, we pick an arbitrary point $x \in \partial b'_0$, we find the point $x' = \uparrow(x)$, and we check that x and x' are endpoints of a component of $A \cap \partial_v N$. This can be done by using the "orbit-tracking" trick described in the proof of Proposition 2.16. To be concrete, we run the algorithm of Proposition 2.13 on $A \cap \partial_v N$ again, but with an additional coordinate in the weight function; we set this coordinate to 1 for the normal arcs containing x and x', and to 0 for all other normal arcs. The points x and x' belong to the same component of $A \cap \partial_v N$ if and only if the algorithm outputs an orbit with weight having the value 2 on the additional coordinate.

Properties (i.). We only discuss property (i.1), as property (i.2) is perfectly analogous. We start by applying Proposition 2.19 to the normal surface F and its sub-2-complex $\mathrm{supp}(a_0') \cup F_g$. The algorithm outputs a triangulated surface H, a sub-2-complex H_0 of H, and a simplicial isomorphism $h\colon\ \mathrm{supp}(a_0') \cup F_g \to H_0$, which extends to a homeomorphism $\overline{h}\colon F \to H$; the number of triangles of H is $\mathcal{O}(|\Sigma| + |\chi(F)|)$. It is then easy to construct the sub-2-complex $H' = \overline{f}(F_p)$ of H, by removing from H the triangles in $h(F_g)$ and then keeping the components that intersect $f(F_p \cap F_g)$. Consider the triangulated surface

$$S_0' = {H' \sqcup \partial_0 N}\Big/{\{h(x) \sim y : x \in F_p, y \in \partial_0 N \cap N', x^+ = f_g(y)\}},$$

where the triangulation is induced by those of H' and $\partial_0 N$, and let $\iota_0'\colon H' \sqcup \partial_0 N \to S_0'$ be the quotient map. Unlike their counterparts S_0 and ι_0, the surface S_0' and the map ι_0' can be explicitly constructed in polynomial time in $|\Sigma|$ and $|\chi(F)|$. Moreover, the homeomorphism $F_p \sqcup \partial_0 N \to H' \sqcup \partial_0 N$ obtained by gluing $\overline{h}|_{F_p}\colon F_p \to H'$ and the identity $\partial_0 N \to \partial_0 N$ induces a homeomorphism on the quotient surfaces $S_0 \to S_0'$. It is then clear that the normal multicurves $\iota_0(a_0' \cap F_p) \cup \iota_0(b_0)$ and $\iota_1(a_0' \cap F_p) \cup \iota_0(b_0'')$ are isotopic in S_1 if and only if the normal multicurves $\iota_0'(h(a_0' \cap F_p)) \cup \iota_0'(b_0)$ and $\iota_0'(h(a_0' \cap F_p)) \cup \iota_0'(b_0'')$ are isotopic in S_0'. The latter can be verified by means of [21, Theorem 1.2].

Properties (j.). These conditions are just inequalities, which can be verified directly. ▢

Chapter 5

Certifying hyperbolicity

5.1 Outline of the certificate

In this final chapter, building on the certificate for the monodromy developed in Chapter 4, we will show how to certify hyperbolicity of fibred 3-manifolds. To clarify, we say that a compact 3-manifold is *hyperbolic* if its interior admits a complete finite-volume hyperbolic metric. For fibred 3-manifolds, hyperbolicity can be directly inferred from the monodromy of a fibre, thanks to Thurston's hyperbolisation theorem, which we state below.

A mapping class of a compact orientable surface F is said to be:

- *periodic* if it has finite order in $\mathrm{Mcg}(F)$;
- *reducible* if it preserves an essential multicurve in F up to isotopy;
- *pseudo-Anosov* otherwise.

This trichotomy is known as the *Nielsen-Thurston classification* of mapping classes; see [8, Exposé 9]. A celebrated theorem of Thurston states that the geometry of a fibred 3-manifold is determined by the Nielsen-Thurston type of its monodromy.

Theorem 5.1 (Thurston's hyperbolisation for fibred 3-manifolds, [29])**.** *Let M be a compact orientable 3-manifold that fibres over the circle with fibre F. Then M is hyperbolic if and only if $\chi(F) < 0$ and the monodromy of F is pseudo-Anosov.*

The tools we have developed so far allow us to certify that a given surface F in a compact orientable 3-manifold M is a fibre of M. Moreover, given two essential curves a_0 and a_1 in F, we can certify that a_1 is the image of a_0 under the monodromy φ of F. In this chapter, we will describe how to employ these tools to certify that φ is pseudo-Anosov, and therefore that M is hyperbolic.

It is not surprising that understanding the action of φ on curves in F is enough to decide the Nielsen-Thurston type of φ. The main issue is quantifying how many iterations of φ we need to look at in order to establish whether φ is pseudo-Anosov or not. An effective bound for this number can be obtained by studying the action of φ on the *curve graph* of F (see Section 5.2 for more details). In particular, for a fixed essential curve a_0 in F, it turns out that the distance d_n in the curve graph of F between a_0 and $\varphi^n(a_0)$ grows linearly in n when φ is pseudo-Anosov, and remains bounded when φ is periodic or reducible. These statements can be made precise and quantitative (see Proposition 5.4). This implies that, for sufficiently large n, there is a threshold D such that φ is pseudo-Anosov if and only if $d_n > D$. This threshold is explicitly computable in terms of F, a_0, and φ.

In our setting, where F is a (least-weight) fibre of M, the essential normal curve $a_0 \subseteq F$ is given as part of the certificate. Thanks to the monodromy certificate of Chapter 4, we can also provide normal curves $a_1, \ldots, a_n$ in F, together with certificates showing that each a_i is the image of a_{i-1} under φ. Under these assumptions, we can quantify how large n needs to be in order to certify that φ is pseudo-Anosov: the bound turns out to be polynomial in $|\chi(F)|$ and linear in the size t of the triangulation of F. The threshold D has the same dependence on $|\chi(F)|$ and t.

All that is left, therefore, is to compute the distance d_n between a_0 and a_n in the curve graph of F. There is a wealth of algorithms that can achieve this, including a recent algorithm by Bell and Webb presented in [5]. This algorithm has the advantage of running in polynomial time in the logarithms of the weights of a_0 and a_n. However, the dependence of the running time on the Euler characteristic of F is at least exponential, which is incompatible with our goal of certifying hyperbolicity in polynomial time in $|\chi(F)|$. In order to overcome this issue, in [4] we have developed a new algorithm that runs in polynomial time in the logarithms of the weights of a_0 and a_n, but also in the Euler characteristic of F. Its output is not quite the distance d_n, but rather an estimate of d_n up to additive and multiplicative error. To account for this error, the values of n and D need to be adjusted. However, since the error term is polynomial in $|\chi(F)|$, the new values are still polynomial in $|\chi(F)|$ and linear in t.

To summarise, we provide normal curves $a_0, \ldots, a_n$ in F as a certificate, with auxiliary instances of Certificate 2 to prove that these curves are iterates of a_0 under the monodromy φ; the integer n is polynomial in $|\chi(F)|$ and linear in t – see Certificate 4 for the exact value. In order to verify the certificate, one should run the distance estimate algorithm (stated in Proposition 5.5) on the curves a_0 and a_n, and check that the output is greater than the adjusted threshold – which can also be computed explicitly, see Certificate 4 again.

One important point to make is that all the quantities we have discussed and the running times of the algorithms we employ depend on $|\chi(F)|$; the dependence is polynomial – and we have taken great care to ensure this – but it cannot be ignored. Since the Euler characteristic of a fibre can, in general, be exponential in the size of the triangulation of M, our final result will be that the hyperbolicity of a fibred 3-manifold can be certified in polynomial time in the size of the triangulation and in the Euler characteristic of a fibre. However, the natural measure of complexity of a fibred 3-manifold is simply the size of its triangulation, which is why we cannot claim to have placed the problem Hyperbolicity detection for fibred 3-manifolds in NP.

To phrase our result in terms of the complexity class NP (which is convenient for providing a statement that is both precise and concise), we have to include the Euler characteristic of the fibre in the size of the input; this explains why our statement of Theorem 5.18 is less natural than one might expect. A convenient situation in which this exploit is not necessary is when the input 3-manifold is the complement of a knot in the 3-sphere, which is described by a diagram with n crossings. In this case – assuming that the knot is fibred – the Euler characteristic of the fibre is guaranteed to be at least $-n$. The exterior of the knot can be triangulated with $\mathcal{O}(n)$ tetrahedra; therefore, our certificate for hyperbolicity can then be verified in polynomial time in n. Since n is also the natural measure of size of the input, we can then claim that the problem Hyperbolicity detection for fibred knots is in NP.

There are two caveats to this overview. The first one is almost inconsequential, but it must be addressed. The algorithm of [4] does not work when F is a torus with one boundary component or a sphere with 4 boundary components. However, for these two "low-complexity" surfaces, one can decide whether a mapping class is pseudo-Anosov with a much simpler procedure, that does not involve the curve graph at all. In Proposition 5.2, we show that knowing the intersection numbers between a_0 and a_1 and between a_0 and a_2 is enough to decide whether φ is pseudo-Anosov.

The second issue is more substantial. We have not mentioned this, but the values of n

and D – and hence the running times of the algorithms we use to certify hyperbolicity – also depend on the weight of a_0. Therefore, our ability to certify hyperbolicity in polynomial time hinges on the existence of an essential curve $a_0 \subseteq F$ with polynomial weight in the parameters t and $|\chi(F)|$. This fact is surprisingly hard to establish. Bounds on the "combinatorial injectivity radius" – that is, the smallest possible weight of an essential curve – of normal surfaces have appeared in the literature (see [6, Theorem 7.7] for instance), but not with the polynomial dependence that we need. A large portion of this chapter is devoted to proving that an essential curve a_0 with polynomial weight does indeed exist. Our proof only works in the case where the normal surface F is a least-weight fibre of a hyperbolic 3-manifold, which is precisely the case we are interested in. Unexpectedly, the argument relies on a variation of our monodromy certificate; we refer the reader to the beginning of Section 5.4 for a short outline of the proof.

The rest of this chapter is organised as follows. In Section 5.2, we describe the quantitative criteria to decide whether a mapping class is pseudo-Anosov, both in the low-complexity case (Proposition 5.2) and in the general case (Proposition 5.4); we then adapt the distance estimate algorithm of [4] to the setting of fibred 3-manifolds in Proposition 5.5. Sections 5.3 and 5.4 are devoted to showing the existence of a polynomial-weight essential curve in least-weight fibres of hyperbolic 3-manifolds. In Section 5.5, we describe in detail our certificate for hyperbolicity, both in the low-complexity case and in the general case. Finally, we conclude with Section 5.6, that contains the precise statement of our main result on certifying hyperbolicity of fibred 3-manifolds, as well as a brief discussion of the case of fibred knots.

5.2 Pseudo-Anosov detection

On surfaces of low complexity, namely the torus with one boundary component and the sphere with 4 boundary components, deciding if a mapping class is pseudo-Anosov is exceptionally easy. The following proposition provides a simple criterion based exclusively on computations of intersection numbers.

Proposition 5.2 (Criterion for the Nielsen-Thurston classification on low-complexity surfaces)**.** *Let F be either a torus with one boundary component, or a sphere with 4 boundary components. Let $\varphi \in \mathrm{Mcg}(F)$ be a mapping class, and let a be an essential curve in F. Then φ is pseudo-Anosov if and only if*

$$i(a, \varphi^2(a)) > 2i(a, \varphi(a)).$$

Proof. The key ingredient of this proof is the following algebraic identity. Let

$$\mathbf{P} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix},$$

let $\mathbf{A}$ be any 2×2 matrix, and let $\mathbf{x} \in \mathbb{Z}^2$. Then we have that

$$\mathbf{x}^t\mathbf{P}\mathbf{A}^2\mathbf{x} = \operatorname{tr}(\mathbf{A})\mathbf{x}^t\mathbf{P}\mathbf{A}\mathbf{x}, \tag{5.1}$$

where $\operatorname{tr}(\mathbf{A})$ denotes the trace of $\mathbf{A}$; this is a consequence of the identities

$$\mathbf{A}^2 = \operatorname{tr}(\mathbf{A})\mathbf{A} - \det(\mathbf{A})\mathbf{I} \quad \text{and} \quad \mathbf{x}^t\mathbf{P}\mathbf{x} = \mathbf{0},$$

where $\mathbf{I}$ is the identity matrix.

When F is a torus with one boundary component. Upon fixing a basis of $H_1(F;\mathbb{Z}) \cong \mathbb{Z}^2$, consisting of two curves intersecting once, there is an identification of $\mathrm{Mcg}(F)$ with $\mathrm{SL}(2,\mathbb{Z})$ given by the action of the mapping class group on $H_1(F;\mathbb{Z})$. Given two essential curves a_1 and a_2 in F, whose homology classes correspond to $\mathbf{x}_1$ and $\mathbf{x}_2$ in $\mathbb{Z}^2$, the intersection number between a_1 and a_2 is given by

$$i(a_1, a_2) = |\mathbf{x}_1^t\mathbf{P}\mathbf{x}_2|.$$

Finally, if $\mathbf{A} \in \mathrm{SL}(2,\mathbb{Z})$ is the matrix describing the action of φ on $H_1(F;\mathbb{Z})$, then it is well-known that φ is pseudo-Anosov if and only if $|\operatorname{tr}(\mathbf{A})| > 2$.

Now, let $\mathbf{x} \in \mathbb{Z}^2$ correspond to the homology class of a. If $i(a, \varphi(a)) = 0$ then $\varphi(a)$ is isotopic to a, and the claim in the statement is trivially true. Otherwise, it follows from (5.1) that

$$\frac{i(a, \varphi^2(a))}{i(a, \varphi(a))} = \left|\frac{\mathbf{x}^t\mathbf{P}\mathbf{A}^2\mathbf{x}}{\mathbf{x}^t\mathbf{P}\mathbf{A}\mathbf{x}}\right| = |\operatorname{tr}(\mathbf{A})|,$$

which is greater than 2 if and only if φ is pseudo-Anosov.

When F is a sphere with 4 boundary components. Let T be a torus with four boundary components, and let $\widehat{T}$ be the torus obtained by capping off each boundary component of T with a disc. [7, §2.2.5] describes a 2-sheeted covering map $\pi\colon T \to F$ induced by the *hyperelliptic involution* $\iota\colon \widehat{T} \to \widehat{T}$. For every essential curve $b \subseteq F$, this covering map has the property that $\pi^{-1}(b)$ is the union of two curves in T that are isotopic in $\widehat{T}$. Moreover, every essential curve $b \subseteq T$ can be isotoped so that $\pi(b)$ is an essential curve in F.

Let $\widetilde{\varphi}\colon T \to T$ be a lift of φ to T, and let $\widehat{\varphi}\colon \widehat{T} \to \widehat{T}$ be the unique completion (up to isotopy) of $\widetilde{\varphi}$ to $\widehat{T}$. By the one-to-one correspondence between essential curves in F and in and $\widehat{T}$ discussed above, it is easy to see that a power of $\widehat{\varphi}$ preserves an essential curve in $\widehat{T}$ up to isotopy if and only if a power of φ preserves an essential curve in F up to isotopy. We conclude that φ is pseudo-Anosov if and only if $\widehat{\varphi}$ is Anosov.

Let b be any essential curve in F, and suppose that b is in minimal position with a. Let $\widetilde{a} = \pi^{-1}(a)$ and $\widetilde{b} = \pi^{-1}(b)$; denote by $\widehat{a}$ and $\widehat{b}$ the multicurves $\widetilde{a}$ and $\widetilde{b}$ considered as multicurves in $\widehat{T}$. It is easy to see that $\widehat{a}$ and $\widehat{b}$ are in minimal position in $\widehat{T}$. In fact, unless a and b are isotopic, every complementary region of $a \cup b$ in F is a bigon minus an open disc, which lifts to a quadrilateral minus an open disc in T; it follows that every complementary region of $\widehat{a} \cup \widehat{b}$ in $\widehat{T}$ is a quadrilateral – in particular, not a bigon. Therefore, we have proved that

$$i(\widehat{a}, \widehat{b}) = 2i(a, b). \tag{5.2}$$

By the previous case, we know that $\widehat{\varphi}$ is Anosov if and only if

$$i(\widehat{a}, \widehat{\varphi}^2(\widehat{a})) > 2i(\widehat{a}, \widehat{\varphi}(\widehat{a}))$$

(the absence of a boundary component in $\widehat{T}$ and the fact that $\widehat{a}$ is a multicurve do not affect the argument carried out in the previous case). By (5.2), this is equivalent to

$$i(a, \varphi^2(a)) > 2i(a, \varphi(a)),$$

which concludes the proof. ▢

For arbitrary surfaces, criteria for deciding if a mapping class is pseudo-Anosov are necessarily more sophisticated. A combinatorial tool that has proved extremely useful to this end is the *curve graph*. Introduced by Harvey in [12], the curve graph $\mathcal{C}(F)$ of a compact connected orientable surface F with $\chi(F) < 0$ is a graph with vertices given by the isotopy classes of essential curves in F, and edges given by pairs of classes of curves that can be realised disjointly; when F is a torus with one boundary component or a sphere with 4 boundary components, the condition for two vertices to be adjacent is relaxed to require that the two isotopy classes of curves admit representatives that intersect once or twice respectively. For two essential curves a and b in F, we denote by

$\mathrm{dist}_{\mathcal{C}(F)}(a,b)$ the combinatorial length of a shortest path in $\mathcal{C}(F)$ connecting the vertices corresponding to a and b; this number is well-defined, since the curve graph can be shown to be connected. In fact, there is an explicit bound on $\mathrm{dist}_{\mathcal{C}(F)}(a,b)$ in terms of the intersection number $i(a,b)$, as given by Hempel in [14, Lemma 2.1] (see [34, Lemma 1.21] for a proof in the case of non-empty boundary).

Proposition 5.3 (Bounding distance in the curve graph with intersection number)**.** *Let F be a compact connected orientable surface with $\chi(F) < 0$. Let a and b be essential curves in F with $i(a,b) \geqslant 1$. Then*

$$\mathrm{dist}_{\mathcal{C}(F)}(a,b) \leqslant 2\log i(a,b) + 2.$$

The seminal work [27] of Masur and Minsky established a deep connection between the geometry of the curve graph and the Nielsen-Thurston classification of mapping classes, by showing that $\mathcal{C}(F)$ is a Gromov-hyperbolic space, and that pseudo-Anosov mapping classes are precisely those that act loxodromically on $\mathcal{C}(F)$. This justifies the following proposition, which provides criteria for the Nielsen-Thurston classification of mapping classes in terms of distances in the curve graph. These criteria are not new: the bound on periodic mapping classes is an easy consequence of the Nielsen realisation theorem of [19]; the bound on reducible mapping classes follows from Masur and Minsky's bounded geodesic image theorem, and its effective version [40, Theorem 4.2.1] by Webb; the bound on pseudo-Anosov mapping classes is proved by Gadre and Tsai in [10]. However, in the proof below we reference [4], where these results have been collected and proved in detail.

Proposition 5.4 (Quantitative criteria for the Nielsen-Thurston classification)**.** *Let F be a compact connected orientable surface with $\chi(F) < 0$. Let $\varphi \in \mathrm{Mcg}(F)$ be a mapping class, and let a be an essential curve in F. For $h \geqslant 1$, denote by $D(h)$ the integer*

$$D(h) = \max\{\mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^j(a)) : 1 \leqslant j \leqslant h\}.$$

Then the following hold for every $j \geqslant 0$.

- *If φ is periodic, then we have that*

$$\mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^j(a)) \leqslant D(6|\chi(F)|).$$

- *Suppose that $\chi(F) \leqslant -2$ and that F is not a sphere with 4 boundary components;*

if φ is reducible, then we have that

$$\mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^j(a)) \leqslant 2D(49\,572|\chi(F)|^3) + 2.$$

- *If φ is pseudo-Anosov, then we have that*

$$162\chi(F)^2 \cdot \mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^j(a)) \geqslant j.$$

Proof. If φ is periodic, then by [4, Proposition 4.1] its order in $\mathrm{Mcg}(F)$ is at most $6|\chi(F)|$; the desired inequality then follows immediately from the fact that the action of φ on $\mathcal{C}(F)$ is an isometry. If φ is pseudo-Anosov, then the inequality in the statement is proved in [4, Proposition 4.5]. Finally, if φ is reducible, then a relevant bound is proved in [4, Proposition 4.6]. This proposition states that there is an integer $1 \leqslant h_0 \leqslant C_{\mathrm{red}}$ such that

$$\mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^{h_0 \cdot h}(a)) \leqslant \mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^{h_0}(a)) + 2$$

for each $h \geqslant 0$. The constant C_{red} is computed explicitly in [4] in terms of a measure of complexity ξ of the surface F. More precisely, the quantity ξ is defined as $3g-3+p$, where g is the genus of F and p is the number of boundary components of F. In particular, we easily see that $\xi \leqslant 3/2 \cdot |\chi(F)|$. We can then compute the constant C_{red} as

$$C_{\mathrm{red}} = 2(100+2) \cdot 162\chi(F)^2 \cdot \xi \leqslant 49\,572|\chi(F)|^3.$$

Any integer $j \geqslant 0$ can be written as $j = h_0 \cdot h + k$ for some integers $h \geqslant 0$ and $0 \leqslant k < h_0$. We immediately get that

$$\begin{aligned}\mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^j(a)) &\leqslant \mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^{h_0 \cdot h}(a)) + \mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^k(a)) \\ &\leqslant \mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^{h_0}(a)) + 2 + \mathrm{dist}_{\mathcal{C}(F)}(a, \varphi^k(a)) \\ &\leqslant 2D(C_{\mathrm{red}}) + 2.\end{aligned}$$

□

If we want to use the criteria of Proposition 5.4 to certify hyperbolicity of fibred 3-manifolds, we need to be able to efficiently compute distances in the curve graph of the fibre. The crucial ingredient for this is [4, Corollary 3.12], which provides an algorithm to coarsely compute the distance in the curve graph between two curves, with running time that is polynomial in the logarithms of the weights of the curves and, importantly, in the Euler characteristic of the surface.

Proposition 5.5 (Coarse computation of distances in the curve graph of a normal surface). *There is an algorithm that takes as input*

- *a triangulation of a compact orientable* 3*-manifold* M *with* t *tetrahedra,*
- *a connected orientable normal surface* F *in* M *with* $\chi(F) \leqslant -2$ *that is not a sphere with* 4 *boundary components, and*
- *two essential normal curves* a *and* b *in* F*,*

and outputs an integer $d \geqslant 0$ *such that*

$$d \leqslant \mathrm{dist}_{\mathcal{C}(F)}(a,b) \leqslant 1313|\chi(F)| \cdot d + 3\,727\,496|\chi(F)|^3.$$

The running time of the algorithm is polynomial in t*,* $|\chi(F)|$, $\mathrm{area}(\mathrm{supp}(a))$*,* $\mathrm{area}(\mathrm{supp}(b))$*,* $\log w(a)$*, and* $\log w(b)$*.*

Proof. A standard application of Proposition 2.19 reduces the problem to the computation of the distance in $\mathcal{C}(G)$ between two normal curves a' and b' in G, where G is a triangulated surface homeomorphic to F; the size of the triangulation of G is bounded above by

$$22(\mathrm{area}(\mathrm{supp}(a)) + \mathrm{area}(\mathrm{supp}(b))) - 3\chi(F),$$

and moreover $w(a') = w(a)$ and $w(b') = w(b)$. The algorithm of [4, Corollary 3.12] requires the input surface to be triangulated as a *one-vertex triangulation* (if it is closed) or an *ideal triangulation* (if it has boundary). A one-vertex triangulation is simply a triangulation with a single vertex; an ideal triangulation of G is a triangulation of a surface G' such that removing open discs around the vertices of G' yields a surface homeomorphic to G. Therefore, in order to apply [4, Corollary 3.12], we need an additional pre-processing step to convert the triangulation of G into a one-vertex or ideal triangulation, which we now describe.

We start by capping off each boundary component of G with a disc, obtained by coning the boundary component over an additional vertex, which we mark as ideal; see Figure 5.1. This operation increases the number of triangles in the triangulation of G by at most a constant multiplicative factor. We now describe a procedure to get rid of non-ideal vertices of G one by one. Fix a non-ideal vertex v of G. If v is not adjacent to any vertex other than itself, then we have obtained a one-vertex triangulation of the closed surface G; we can therefore terminate the procedure. Otherwise, we distinguish a few cases.

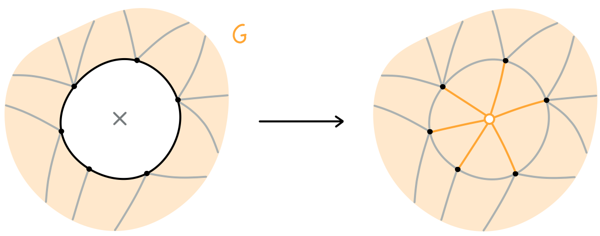

Figure 5.1. We cap off each boundary component of G with a disc, obtained by coning over an ideal vertex.

- If v has degree 2, then it is adjacent to exactly two triangles, say T_1 and T_2. Let e_i be the edge of T_i that is not adjacent to v for $i \in \{1, 2\}$. Then the edges e_1 and e_2 are different, and we can collapse $T_1 \cup T_2$ onto e_1, as shown in Figure 5.2a. This operation removes the vertex v, and reduces the number of triangles in the triangulation.

- If v has degree 1, then it is adjacent to a single triangle T_1. Let e be the edge of T_1 that is not adjacent to v, and let T_2 be the triangle adjacent to T_1 along e. Since the triangles T_1 and T_2 are different, we can then apply a *flip* (also known as 2-2 *Pachner move*) along e, replacing it with an edge transverse to it as shown in Figure 5.2b. After this move, the vertex v has degree 2, and we can remove it as described in the previous case.

- If v has degree at least 3, then we are presented with two cases. If every edge adjacent to v ends in a vertex other than v, then performing a flip along any of these edges decreases the degree of v, as shown in Figure 5.2c; note that this is always allowed, since each of these edges is necessarily adjacent to two distinct triangles. Otherwise, there must be a triangle T such that exactly two vertices of T are equal to v. Denote by e the edge of T that starts and ends at v. Then a flip along e will decrease the degree of v, as shown in Figure 5.2d; note that this is allowed, since e must be adjacent to a triangle that is not T.

After fixing the vertex v, we can repeat the steps above until v is removed; since the degree of v decreases at each step, this will happen after a number of steps that is at most linear in the size of the triangulation of G. We also remark that each step does not increase the number of triangles in the triangulation. After v has been removed, we switch to a different non-ideal vertex and iterate again. In the end, we will obtain a

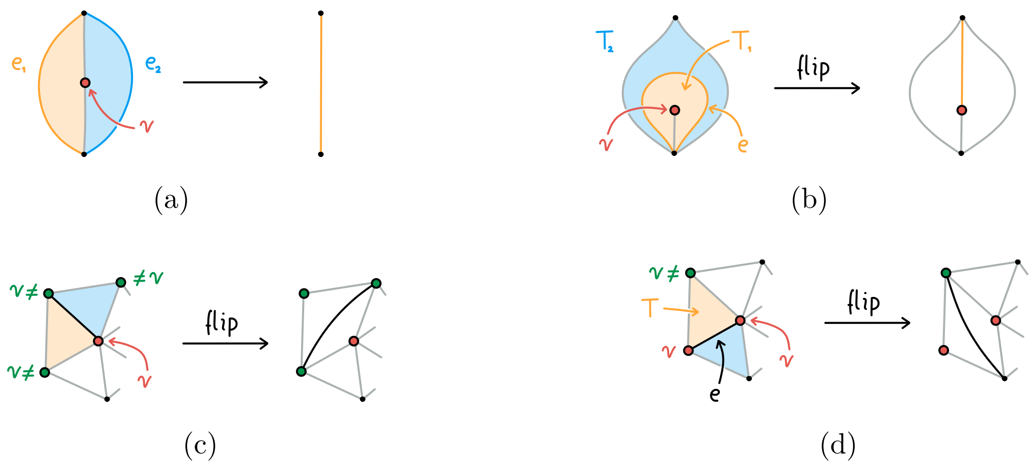

Figure 5.2. (a) If v has degree 2, then we can remove it by collapsing $T_1 \cup T_2$ onto e_1. (b) If v has degree 1, then we can perform a flip along e and reduce to the degree-2 case. (c) If v has degree at least 3 and it is not self-adjacent, then we can perform a flip along any edge adjacent to v to reduce the degree of v. (d) If v has degree at least 3 and there is a triangle T with vertices v, v, and not v, then we can perform a flip along the $v - v$ edge e to reduce the degree of v.

one-vertex or ideal triangulation $\mathcal{T}'$ of G. It is easy to see that the total number of steps is at most quadratic in the size of the original triangulation of G, and that each step can be performed in polynomial time.

During this process, we also need to keep track of the normal vectors of the curves a' and b' with respect to the changing triangulation. This is straightforward to do when performing a flip. However, when performing a collapse as in Figure 5.2a, the curves a' and b' may not be normal with respect to the new triangulation. In particular, the curves a' and b' could intersect e_1 twice with opposite orientations, without crossing any other edge of the new triangulation in between. For this reason, after performing a collapse, we run the algorithm of [21, Theorem 1.3] to the curves a' and b'. This algorithm takes as input a possibly non-normal essential curve in a triangulated surface, and outputs a normal curve that is isotopic to the input curve. The running time of the algorithm is polynomial in the size of the triangulation and the logarithm of the weight of the input curve. Moreover, the weight of the output curve is at most the weight of the input curve; we refer the reader to [21] for a precise statement, including how to represent a non-normal curve. In conclusion, we see that we can augment the procedure described above to also output the normal vectors of a' and b' with respect to the triangulation $\mathcal{T}'$.

Since each flip increases the weights by a factor of at most 2, and each collapse does not increase the weights, the logarithms of the weights of a' and b' with respect to $\mathcal{T}'$ will increase by a factor that is at most polynomial in the size of the original triangulation of G.

We are now in a position to apply [4, Corollary 3.12] to estimate the distance in $\mathcal{C}(G)$ between a' and b'. The algorithm of [4, Corollary 3.12] takes as input the triangulation $\mathcal{T}'$ of G and the normal vectors of a' and b' with respect to $\mathcal{T}'$, and outputs an integer $d' \geqslant 0$ such that

$$d' - L_+ \leqslant \mathrm{dist}_{\mathcal{C}(G)}(a', b') \leqslant L_\times \cdot d' + L_+$$

for constants L_+ and $L_\times$ that are computed explicitly in [4]. Recall that the integer ξ used in [4] is bounded above by $3/2 \cdot |\chi(F)|$; under the current assumptions on F, we also have that $\xi \geqslant 2$. Setting $d = \max\{d' - L_+, 0\}$ and tracking the constants, we obtain that

$$d \leqslant \mathrm{dist}_{\mathcal{C}(G)}(a', b') \leqslant 875\xi \cdot d + \frac{4\,417\,773}{4} \cdot \xi^3 \leqslant 1313|\chi(F)| \cdot d + 3\,727\,496|\chi(F)|^3,$$

as desired. ▢

5.3 Intersections of curves with inessential subsurfaces

Proposition 5.6 below is a key ingredient in the proof of existence of small-weight essential curves (see Proposition 5.11). The statement, admittedly technical, can be paraphrased as follows. Let F be a triangulated surface (in our case, the fibre of a fibred 3-manifold), and let G be a subsurface of F that is "inessential", in the sense that every curve contained in G is inessential in F. Then essential least-weight curves in F "dislike" G, in the sense that only a small fraction of their length (or, more accurately, of their weight) can be contained in G. This fraction gets larger as G gets more complicated, but – very importantly – it is inversely proportional to the smallest possible weight of an essential curve in F. We refer the reader to the beginning of Section 5.4 for a discussion of the role played by this proposition in the proof on Proposition 5.11.

Proposition 5.6 (Least-weight curves dislike inessential regions)**.** *Le $\mathcal{T}$ be a triangulation of a compact connected orientable surface F. Let G be a subsurface of F with m components, such that ∂G is a general position multicurve in F. Suppose that F contains an essential curve, but every curve contained in G is inessential in F. Denote by ω the* *smallest possible weight of a general position essential curve in F, and suppose that*

$$\omega > 11m \cdot w(\partial G).$$

Let a be an essential least-weight general position curve in F that intersects ∂G transversely and such that $a \cap \partial G$ is disjoint from the 1-skeleton of $\mathcal{T}$. Then

$$\frac{w(a \cap G)}{w(a)} \leqslant \frac{48m^2 \cdot w(\partial G)}{\omega}.$$

The proof of Proposition 5.6, in turn, relies on the following lemma, whose evocative name is inspired by Figure 5.4

Lemma 5.7 (Flower lemma)**.** *In the setting of Proposition 5.6, let C be a subsurface of F that is either a disc with $\partial C \subseteq \operatorname{int}(F)$ or a collared annulus. Let b be a subarc of a such that $\partial b \subseteq \partial C$ and a small neighbourhood of ∂b in b is contained in C. Suppose that $b \cap \partial C$ is disjoint from the 1-skeleton of $\mathcal{T}$. Moreover, suppose that $w(b) + w(\partial_F C) < \omega$. Then there is a (possibly empty) subarc b' of b whose interior is disjoint from C and such that*

$$w(b') \geqslant w(b) - w(\partial_F C).$$

Proof. In the context of this proof, we call a compact subsurface P of F a *petal* if the following hold:

(i) P is a disc or a collared annulus;

(ii) $P \cap C$ is a non-empty subarc of $\partial_F C$;

(iii) $\operatorname{clos}(\partial_F P \setminus C)$ is a non-empty subarc of b.

In this setting, we say that the subarc $\operatorname{clos}(\partial_F P \setminus C)$ of b *bounds* P. We say that P is *maximal* if it is not strictly contained in any other petal. If P is a collared annulus, then we say that P is a *punctured petal.* Figure 5.3a shows examples of petals.

Every component of $\operatorname{clos}(b \setminus C)$ **bounds a petal.** In fact, let c be such a component, and let r be a subarc of $\partial_F C$ connecting the endpoints of c. Note that the curve $c \cup r$ has weight

$$w(c \cup r) = w(c) + w(r) \leqslant w(b) + w(\partial_F C) < \omega,$$

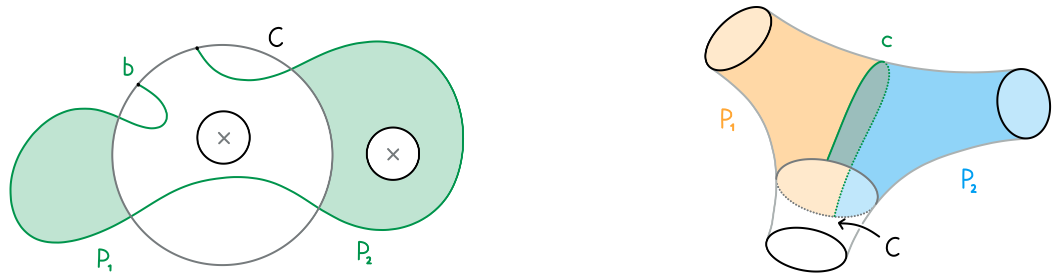

(a) (b)

Figure 5.3. (a) Two petals P_1 and P_2; the former is not punctured, the latter is. (b) If a component c of $\text{clos}(b \setminus C)$ bounds two petals P_1 and P_2 (one on each side), then F is a sphere with at most 3 boundary components.

therefore it bounds a disc or a collared annulus P. If $\text{int}(P)$ is disjoint from C, then P is a petal bounded by c. Otherwise, we see that C is contained in P, and the closure of $P \setminus C$ is a petal bounded by c.

If C is an annulus, then no petal is punctured. In fact, if C is an annulus and P is a punctured petal, then $C \cup P$ is a subsurface of F that is homeomorphic to a pair of pants. The curve $\partial_F(C \cup P)$ has weight

$$w(\partial_F(C \cup P)) \leqslant w(\partial_F C) + w(b) < \omega,$$

therefore it bounds a disc or a collared annulus on the side opposite $C \cup P$. This, however, contradicts the assumption that F is not an annulus nor a pair of pants.

Every component of $\text{clos}(b \setminus C)$ bounds exactly one petal. Suppose, for a contradiction, that a component c of $\text{clos}(b \setminus C)$ bounds a petal on each side. Then F would be the union of three surfaces – namely, the two petals and C – that only intersect along their boundaries, as shown in Figure 5.3b. Each of these surfaces is either a disc or a collared annulus, therefore F would be a sphere with at most 3 boundary components, which contradicts our assumption.

For every two petals, either they are disjoint, or one is contained in the other. In fact, let P_1 and P_2 be two petals, bounded by c_1 and c_2 respectively. If c_1 is contained in P_2, then P_1 is also contained in P_2 by uniqueness of P_1. Similarly, if c_2 is contained in P_1, then P_2 is also contained in P_1. Otherwise, it is easy to see that P_1 and P_2 are disjoint.

There is at most one maximal punctured petal. In fact, let P_1 and P_2 be two disjoint (and hence distinct) maximal punctured petals. Then $C \cup P_1 \cup P_2$ is a subsurface of F that is homeomorphic to a pair of pants. The curve $\partial_F(C \cup P_1 \cup P_2)$ has weight

$$w(\partial_F(C \cup P_1 \cup P_2)) \leqslant w(\partial_F C) + w(b) < \omega,$$

therefore it bounds a disc or a collared annulus on the side opposite $C \cup P_1 \cup P_2$. This, however, contradicts the assumption that F is not an annulus nor a pair of pants.

Constructing the flower. Let S (the "flower") be the union of C with all the (maximal) petals; note that S contains b. It is possible that one or both endpoints of b are contained in the interior of S. To fix this, let b_0 be the largest subarc of a containing b that is contained in S; since a is essential in F and S is a disc or a collared annulus, we see that a is not contained in S, and therefore b_0 is well-defined. In particular, note that $\partial b_0 \subseteq \partial S$.

Suppose first that there are no punctured petals. Then b_0 is isotopic in S, fixing the endpoints, to an arc b_1 properly embedded in C, as shown in Figure 5.4a. We can isotope b_1, fixing the endpoints, to an arc $b_2 \subseteq \partial_F C$. By Corollary 2.5, we deduce that

$$w(b) \leqslant w(b_0) \leqslant w(b_2) \leqslant w(\partial_F C).$$

In particular, we can take b' to be the empty arc.

Suppose now that there is a maximal punctured petal, bounded by a subarc c of b. Note that, in this case, the subsurface S of F is a collared annulus. Therefore, depending on the isotopy class of b_0 in S relative its endpoints, we have that b_0 can be isotoped in S, fixing the endpoints, to an arc b_1 that is either contained in C, or is the union of c with two disjoint arcs contained in C. In the first case, we conclude as above, taking b' to be the empty arc. In the second case, depicted in Figure 5.4b, we can isotope b_1, fixing the endpoints, to the union b_2 of c with two disjoint subarcs of $\partial_F C$. By Corollary 2.5, we deduce that

$$w(b) \leqslant w(b_0) \leqslant w(b_2) \leqslant w(c) + w(\partial_F C).$$

In particular, we can take $b' = c$. ▢

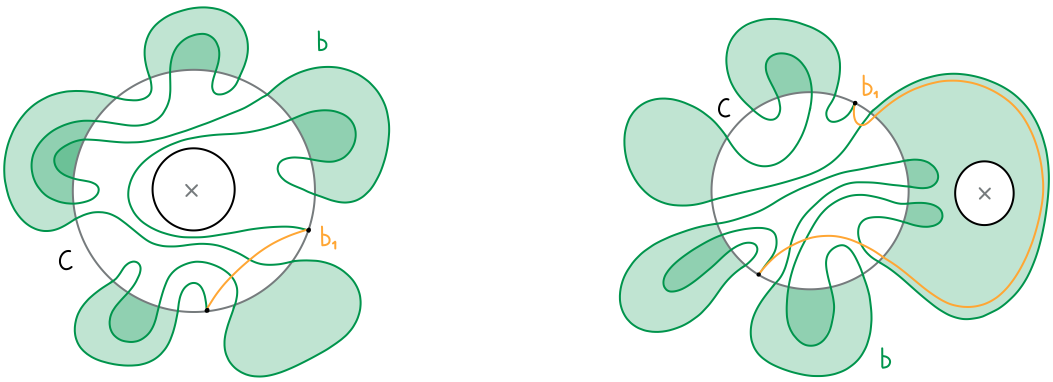

(a) (b)

Figure 5.4. (a) The arc b_0 is isotopic to an arc b_1 in C. (b) The arc b_0 is isotopic to a union b_1 of c (here pushed slightly inside S) with two arcs in C.

Proof of Proposition 5.6. Let G' be the union of G with all the components of $\operatorname{clos}(F \setminus C)$ that are discs or collared annuli, where C ranges over all components of G. Clearly, the number m' of components of G' is bounded above by m. It is not hard to see that each component C of G' is a disc or a collared annulus with $\partial_F C \subseteq \partial G$.

Since a is essential, it is not contained in G'. If a is disjoint from G', then it is disjoint from G, and we are done. Otherwise, every component of $a \cap G'$ – or, in fact, every arc in G' with endpoints in $\partial_F G'$ – is isotopic in G', fixing the endpoints, to a subarc of $\partial_F G' \subseteq \partial G$, as depicted in Figure 5.5a. Since a is least-weight, Corollary 2.5 implies that every component b of $a \cap G'$ satisfies

$$w(b) \leqslant \kappa, \tag{5.3}$$

where for convenience we have set $\kappa = w(\partial G)$.

Fix now an orientation on a; for every point $p \in a$, this induces two vectors tangent to a at p: the *forward vector* and the *backward vector*. If p and q are two distinct points of a, denote by $[p, q]$ the unique subarc of a between p and q such that the forward vector at p points inside $[p, q]$. We also let $[p, p] = \{p\}$ for every $p \in a$.

We now inductively define a sequence of points in $a \cap \partial G'$. Start by picking a point $q_0 \in a \cap \partial G'$ such that the forward vector at q_0 points inside G'. Suppose we have defined $q_0, \ldots, q_{i-1}, p_1, \ldots, p_{i-1}$ for some $i \geqslant 1$. Then, we let p_i be the first point of $a \cap \partial G'$ after q_{i-1}, according to the orientation of a. We also let q_i be the first point of $a \cap \partial G'$ after p_i such that

(i) the forward vector at q_i points inside G', and

(ii) $w([p_i, q_i]) > 2\kappa$.

If $i > 1$ and $[p_i, q_1]$ contains p_1, then we set $n = i - 1$ and stop the process. Note that this construction is well-defined: for every $i \geqslant 1$, we have that

$$\begin{aligned} w([p_i, q_{i-1}]) &= w(a) - w([q_{i-1}, p_i]) \\ &\geqslant w(a) - \kappa && \text{by (5.3)} \\ &\geqslant \omega - \kappa \\ &> 2\kappa && \text{since } \omega > 3\kappa. \end{aligned}$$

We also remark that, by construction, we have that either $q_n = q_0$ or $w([p_{n+1}, q_0]) \leqslant 2\kappa$; either way, (5.3) implies that

$$w([q_n, q_0]) \leqslant 3\kappa. \tag{5.4}$$

For every $1 \leqslant i \leqslant n$, let $a_i = [p_i, q_i]$. These are pairwise disjoint subarcs of a; every a_i satisfies $w(a_i) > 2\kappa$ and

$$w(a_i \setminus \mathrm{int}(G')) \geqslant w(a_i) - 3\kappa. \tag{5.5}$$

To prove (5.5), let p' be the first point of $a \cap \partial G'$ before q_i. If $p' = p_i$, then the inequality is trivial. Otherwise, let q' be the first point of $a \cap \partial G'$ before p'. We have that

$$\begin{aligned} w(a_i) &= w([p_i, q']) + w([q', p']) + w([p', q_i]) \\ &\leqslant 2\kappa + \kappa + w([p', q_i]) && \text{by (5.3) and definition of } q_i \\ &\leqslant 2\kappa + \kappa + w(a_i \setminus \mathrm{int}(G')), \end{aligned}$$

as desired. See Figure 5.5b for a depiction of the subarcs a_i.

Set

$$\theta = \frac{\omega - (2m + 3)\kappa}{2m},$$

and consider the set of indices

$$\mathcal{I}_{\text{short}} = \{1 \leqslant i \leqslant n : w(a_i) < \theta\},$$

and the complement $\mathcal{I}_{\text{long}} = \{1, \ldots, n\} \setminus \mathcal{I}_{\text{short}}$.

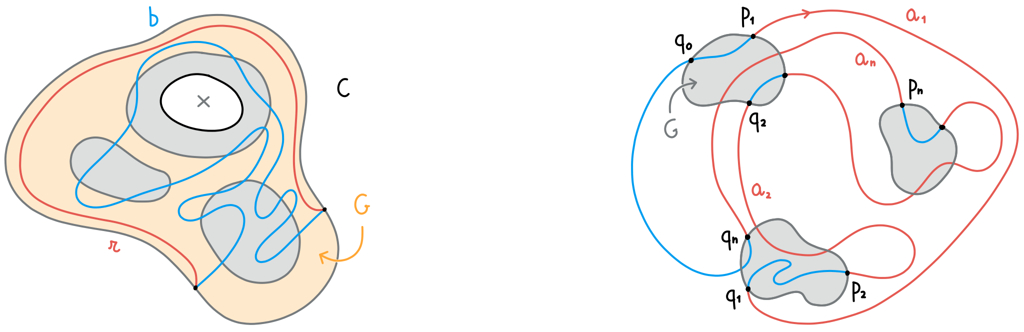

(a) (b)

Figure 5.5. (a) Every arc $b \subseteq C$ is isotopic in C to an arc $r \subseteq G$ with $w(r) \leqslant w(\partial G)$. (b) The subarcs a_i of a, together with the points p_i and q_i employed in the construction.

If $n \leqslant 2m$ then $\mathcal{I}_{\text{long}}$ is non-empty. If, in this setting, the set $\mathcal{I}_{\text{long}}$ were empty, then we would have that

$$\begin{aligned} w(a) &= \sum_{i=1}^{n}\Big(w([q_{i-1},p_i]) + w(a_i)\Big) + w([q_n,q_0]) & \\ &< \sum_{i=1}^{n}(\kappa+\theta) + w([q_n,q_0]) && \text{by (5.3) and definition of } \mathcal{I}_{\text{short}} \\ &\leqslant 2m\cdot(\kappa+\theta) + 3\kappa && \text{by (5.4)} \\ &= \omega && \text{by definition of } \theta, \end{aligned}$$

which contradicts the definition of ω.

There are at most $2m$ consecutive integers in $\mathcal{I}_{\text{short}}$. For ease of notation, suppose for a contradiction that $\{1,\dots,2m+1\} \subseteq \mathcal{I}_{\text{short}}$ (the argument for the general case is identical). Since G' has $m' \leqslant m$ components, there must be three indices $1 \leqslant i < j < k \leqslant 2m+1$ such that p_i, p_j, and p_k lie in the same component C of G'.

Consider the arc $b_1 = [q_{i-1}, p_j]$. We can write b_1 as the union of $a_i, \dots, a_{j-1}$ (each of which has weight strictly smaller than θ) and $j-i+1$ arcs contained in G' (each of which has weight at most κ by (5.3)). As a consequence, we find that

$$\begin{aligned} w(b_1) &< (j-i)\cdot\theta + (j-i+1)\cdot\kappa \\ &\leqslant 2m\cdot(\theta+\kappa) \\ &\leqslant \omega - \kappa && \text{by definition of } \theta. \end{aligned}$$

By Lemma 5.7, it follows that $w(c_1) \geqslant w(b_1) - \kappa$ for some component c_1 of $\operatorname{clos}(b_1 \setminus C)$. Note that

$$w(c_1) \geqslant w(b_1) - \kappa \geqslant w(a_i) - \kappa > \kappa.$$

We can apply the same argument to the arc $b_2 = [q_{j-1}, p_k]$ to conclude that $w(c_2) \geqslant w(b_2) - \kappa > \kappa$ for some component c_2 of $\operatorname{clos}(b_2 \setminus C)$. Finally, by considering $b = [q_{i-1}, p_k]$, we find a component c of $\operatorname{clos}(b \setminus C)$ such that $w(c) \geqslant w(b) - \kappa$. Note that c must be disjoint from at least one of c_1 and c_2; without loss of generality, suppose that c is disjoint from c_1. Moreover, we have that $c_1 \subseteq b_1 \subseteq b$. We then get a contradiction:

$$w(b) \geqslant w(c) + w(c_1) > w(b) - \kappa + \kappa = w(b).$$

Final estimates. It follows from the previous two steps that

$$\begin{aligned} |\mathcal{I}_{\text{long}}| &\geqslant \max\left\{1, \frac{n-2m}{2m+1}\right\} \\ &\geqslant \frac{n+1}{4m+2} && \text{by elementary algebra.} \end{aligned}$$

We conclude with the following estimates:

$$\begin{aligned} w(a \setminus \operatorname{int}(G)) &\geqslant w(a \setminus \operatorname{int}(G')) \\ &\geqslant \sum_{i \in \mathcal{I}_{\text{long}}} w(a_i \setminus \operatorname{int}(G')) \\ &\geqslant \sum_{i \in \mathcal{I}_{\text{long}}} (w(a_i) - 3\kappa) && \text{by (5.5)} \\ &\geqslant |\mathcal{I}_{\text{long}}| \cdot (\theta - 3\kappa) \\ &\geqslant \frac{n+1}{4m+2} \cdot \frac{\omega - (8m+3)\cdot\kappa}{2m}; \end{aligned}$$

$$\begin{aligned} w(a \cap G) &\leqslant w(a \cap G') \\ &\leqslant \sum_{i=1}^{n} (w([q_{i-1}, p_i]) + w(a_i \cap G')) + w([q_n, q_0]) \\ &\leqslant \sum_{i=1}^{n} (\kappa + w(a_i \cap G')) + w([q_n, q_0]) && \text{by (5.3)} \\ &\leqslant \sum_{i=1}^{n} (\kappa + 3\kappa) + w([q_n, q_0]) && \text{by (5.5)} \\ &\leqslant (4n+3)\cdot\kappa && \text{by (5.4)} \end{aligned}$$

$$\leqslant 4(n+1) \cdot \kappa;$$

$$\begin{aligned}
\frac{w(a)}{w(a \cap G)} &= \frac{w(a \setminus \operatorname{int}(G))}{w(a \cap G)} + 1 \\
&\geqslant \frac{\omega - (8m+3) \cdot \kappa}{16m \cdot (2m+1) \cdot \kappa} + 1 \\
&\geqslant \frac{\omega}{48m^2 \cdot \kappa} && \text{by elementary algebra.}
\end{aligned}$$

Admittedly, this argument does not work when $\kappa = w(\partial G) = 0$. However, in this case, each component of G is contained in a single triangle of $\mathcal{T}$. Therefore, the inequality in the statement holds trivially, since $w(a \cap G) = 0$. ▢

5.4 Short curves on least-weight normal fibres

Our goal is now to prove Proposition 5.11, stating that least-weight fibres of hyperbolic 3-manifold contain polynomial-weight essential curves. The proof goes as follows. Let M be a triangulated fibred hyperbolic 3-manifold, and let F be a normal fibre of M admitting a certificate in $\mathfrak{S}^*_{\text{fib}}(M, F)$ – for instance, the fibre F could be the least-weight fibre of M. Using the notation of Certificate 1, if F_g or F'_g contain a curve that is essential in F, then we are done; this follows from the fact that F_g and F'_g have small area (see Proposition 3.6), which implies that if they contain an essential curve, then they contain a short one (this is the content of Lemma 5.8 below). Therefore, we can assume that F_g and F'_g are both "inessential" in F, meaning that every curve contained in F_g or F'_g is inessential in F.

Let ω be the smallest possible weight of a general position essential curve in F; our goal is to show that ω is bounded above by a polynomial function of t (the size of the triangulation of M) and $|\chi(F)|$. Let a_0 be a least-weight general position essential curve in F. By Proposition 5.6, we know that the weight of $a_0 \cap F_g$ is only a small fraction of the weight of a_0; more precisely, we have that

$$w(a_0 \cap F_g) \leqslant \frac{\beta}{\omega} \cdot w(a_0),$$

where β depends polynomially on t. Before continuing, let us remark that we are glossing over the fact that F_g is, in general, a sub-2-complex of F; to be precise, we would need to replace F_g with an actual subsurface G of F in the argument; this is done in detail in

Lemma 5.9 below. Going back to the proof, recall from Certificate 2 that the image a_1 of a_0 under the monodromy φ of F – or, at least, the one coming from the monodromy certificate – has weight

$$w(a_1) \leqslant w(a_0) + C \cdot w(a_0 \cap F_g),$$

where C is an exponential function of t. In other words, if the weight of $a_0 \cap F_g$ is small, then the weight of a_1 is only marginally larger than the weight of a_0. To be more quantitative, we find that

$$w(a_1) \leqslant \left(1 + \frac{C \cdot \beta}{\omega}\right) \cdot w(a_0).$$

It is not restrictive to assume that a_1 is least-weight. We can then proceed inductively, defining a sequence of least-weight curves $a_0, a_1, \ldots$ in F such that $a_i = \varphi^i(a_0)$ and

$$w(a_i) \leqslant \left(1 + \frac{C \cdot \beta}{\omega}\right)^i \cdot w(a_0)$$

for each $i \geqslant 1$. Using Propositions 2.2 and 5.3, we can show that the distance between a_0 and a_i in the curve graph of F grows at a linear rate of at most $2C \cdot \beta/\omega$. However, since φ is pseudo-Anosov, Proposition 5.4 implies that the same distance must grow at a linear rate of at least $(162\chi(F)^2)^{-1}$. As a consequence, we find that $\omega \leqslant 324C \cdot \beta \cdot \chi(F)^2$.

If C were a polynomial function of t, then we would be done. However, from how Certificate 2 is stated, the quantity C depends exponentially on t. Nonetheless, we show in Lemma 5.10 that, under the assumption that F'_g is inessential, the constant C can be replaced with a polynomial function C' of t. Then the argument above goes through, showing that ω is at most a polynomial function of t and $|\chi(F)|$. We now carry out this argument in detail, starting with the auxiliary results and culminating with the proof of Proposition 5.11.

Lemma 5.8 (Small essential subsurfaces contain short essential curves)**.** *Let $\mathcal{T}$ be a triangulation of a compact orientable surface F, and let G be a subsurface of F. Suppose that ∂G is a general position multicurve in F. Denote by n the number of edges of $\mathcal{T}$ that intersect G. Suppose that G contains a curve that is essential in F. Then G contains a general position curve a that is essential in F and such that*

$$w(a) \leqslant w(\partial G) + 2n.$$

Proof. In the scope of this proof, we will call components of $e \cap G$ *partial edges*, where e

(a) (b)

Figure 5.6. If a is non-separating and intersects a partial edge e of $\mathcal{T}$ at least twice, then we can perform a cut-and-paste operation on a. (a) If two consecutive intersections have the same orientation, cutting and pasting results in a new non-separating curve a' of smaller weight. (b) If two consecutive intersections have opposite orientations, cutting and pasting results in two new curves a' and a'' of smaller weight, at least one of which is non-separating.

is any edge of $\mathcal{T}$. Up to replacing G with one of its components, we can assume that G is connected. Let us first suppose that G admits a non-separating curve c – that is, a curve such that $G \setminus c$ is connected. Let a be a general position non-separating curve in G of minimal weight amongst all such curves. We claim that a intersects each partial edge at most once. In fact, for the sake of contradiction, suppose that a intersects some partial edge e twice, and consider two consecutive intersections on e.

- If the two intersections have the same orientation, then the cut-and-paste operation displayed in Figure 5.6a would produce a general position non-separating curve of smaller weight than a.

- If the two intersections have opposite orientations, then the cut-and-paste operation displayed in Figure 5.6b would produce two general position curves of smaller weight than a, at least one of which is non-separating.

In either case, we obtain a contradiction to the minimality of a. Since there are at most $n + w(\partial G)/2$ partial edges, we deduce that the weight of a is at most $n + w(\partial G)/2$.

Suppose now that there are no non-separating curves in G – in other words, that G is planar. If any component of ∂G is essential in F, then we can take a to be that component. Otherwise, we see that each boundary component of G must bound a disc or a collared annulus in F; moreover, since G contains a curve that is essential in F, these discs and collared annuli cannot contain G. There must be at least two components c_1

(a) (b)

Figure 5.7. (a) If b intersects a partial edge e that is not disjoint from c_1, then we can construct an arc b' of smaller weight by following e from c_1 to the intersection of e with b, and then following b towards c_2. In the picture, the arrow on b represents the direction going from c_1 to c_2. (b) The curve a follows c_1, then b, then c_2, and then b in the opposite direction.

and c_2 of ∂G that bound a collared annulus in F. Let b a general position arc properly embedded in G that connects c_1 and c_2, of minimal weight amongst all such arcs. By a cut-and-paste argument similar to the previous case, it is easy to see that b intersects each partial edge at most once. Moreover, we claim that b is disjoint from partial edges that intersect c_1 or c_2. In fact, if b intersects a partial edge e that intersects c_1, then we can perform a cut-and-paste operation on b to obtain an arc of smaller weight, contradicting the minimality of b; this is shown in Figure 5.7a We can then take a to be a curve that follows c_1, then b, then c_2, and then b in the opposite direction, as depicted in Figure 5.7b. An easy counting argument shows that there are at most

$$n + \frac{1}{2}(w(\partial G) - w(c_1 \cup c_2))$$

partial edges disjoint from c_1 and c_2. Therefore, we can bound the weight of a as follows:

$$\begin{aligned} w(a) &= w(c_1) + w(c_2) + 2w(b) \\ &\leqslant w(c_1) + w(c_2) + 2\left(n + \frac{1}{2}(w(\partial G) - w(c_1 \cup c_2))\right) \\ &= 2n + w(\partial G). \end{aligned}$$ ▢

Lemma 5.9 (Bounds on the thickening of the horizontal boundary of the guts)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected orientable 3-manifold M with t tetrahedra. Let*

F be a transversely oriented normal surface in M, and let $M' = M \setminus\setminus F$. Denote by X the guts of M', and let G be a sub-2-complex of F such that $G^+ \subseteq \partial_0 X$ or $G^- \subseteq \partial_1 X$. Then $\mathrm{thick}(G)$ *intersects at most $432t^2$ triangles of F. Moreover, there is an arbitrarily small open neighbourhood N of ∂F in F such that $G' = \mathrm{thick}(G) \setminus N$ is a subsurface of F, with $\partial G'$ a general position multicurve in F satisfying*

$$w(\partial G') \leqslant 2592t^2.$$

Proof. Suppose, without loss of generality, that $G^+ \subseteq \partial_0 X$. Denote by $\mathcal{R}$ the set of triangles of F that intersect $\mathrm{thick}(G)$. We claim that

$$\mathcal{R} = \left\{ R_1 \in \mathcal{R} : \begin{array}{c} T_1 \text{ and } T_2 \text{ are tetrahedra of } \mathcal{T}, \\ R_1 \text{ and } R_2 \text{ are triangles of } F,\ R_1 \subseteq T_1,\ R_2 \subseteq T_2, \\ R_2^+ \subseteq \partial_0 X,\ e_1 \text{ is an edge of } \mathrm{ab}(T_1),\ e_2 \text{ is an edge of } \mathrm{ab}(T_2), \\ e_1 \text{ and } e_2 \text{ are identified in } M,\ p_1 \in R_1 \cap e_1,\ p_2 \in R_2 \cap e_2, \\ p_1 \text{ and } p_2 \text{ are identified in } M \end{array} \right\}.$$

In other words, each triangle $R_1 \in \mathcal{R}$ can be obtained with the following procedure. Choose two tetrahedra T_1 and T_2 of $\mathcal{T}$, two edges e_1 and e_2 of $\mathrm{ab}(T_1)$ and $\mathrm{ab}(T_2)$ respectively, and a triangle R_2 of F contained in T_2 and intersecting e_2. Suppose that e_1 and e_2 are identified in M, and that $R_2^+ \subseteq \partial_0 X$. Let p_2 be the point in which e_2 intersects R_2, and let p_1 be the unique point of e_1 that is identified with p_2 in M. Finally, we take R_1 to be a triangle of F contained in T_1 and intersecting e_1 at p_1.

We now justify this claim. Let R_1 be a triangle of F such that $R_1^+ \subseteq \partial_0 X$. Then R_1 can be obtained by the procedure above, setting $T_1 = T_2$ to be the tetrahedron of $\mathcal{T}$ that contains R_1, with $R_2 = R_1$ and $e_1 = e_2$ to be an edge of T_1 intersecting R_1. If, instead, a triangle R_1 intersects $\mathrm{thick}(G)$, but does not satisfy $R_1^+ \subseteq \partial_0 X$, then there is a triangle R_2 of F that is contained in G and is adjacent to R_1 in F; note that R_2 cannot be contained in the same normal disc as R_1. In this case, we take T_1 and T_2 to be the tetrahedra of $\mathcal{T}$ that contain R_1 and R_2 respectively, and we take e_1 and e_2 to be edges of $\mathrm{ab}(T_1)$ and $\mathrm{ab}(T_2)$ respectively, witnessing the adjacency of R_1 and R_2; see Figure 5.8.

We can now bound the size of $\mathcal{R}$ with a counting argument. There are t^2 choices of tetrahedra T_1 and T_2. The edges e_1 and e_2 can each be chosen in 6 ways. The triangle R_2 can lie on at most 6 different normal discs of F (4 triangles and 2 quadrilaterals) for the condition $R_2^+ \subseteq \partial_0 X$ to hold; in particular, there are at most 6 possible choices for the point p_2. The point p_1 is then uniquely determined, and this leaves 2 choices for the

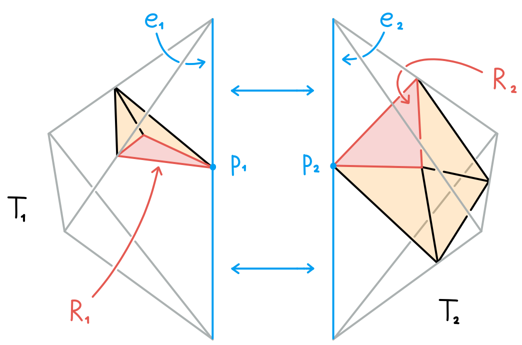

Figure 5.8. The tetrahedra T_1 and T_2, the triangles R_1 and R_2, the edges e_1 and e_2, and the points p_1 and p_2 described in the construction. The normal discs containing R_1 and R_2 are also highlighted.

triangle R_1. In total, we find that

$$|\mathcal{R}| \leqslant t^2 \cdot 6^2 \cdot 6 \cdot 2 = 432t^2.$$

Take N to be an arbitrarily small regular neighbourhood of ∂F in F. We can assume that $c = \partial_F N$ is a general position multicurve in F, intersecting each edge of F at most twice (except for the edges of F that are contained in ∂F, which are necessarily disjoint from c). Since $\partial\,\mathrm{thick}(G)$ also intersects each interior edge of F at most twice, we conclude that $\partial G'$ intersects each interior edge of F at most 4 times. Moreover, every edge intersecting $\partial G'$ is adjacent to triangles of $\mathcal{R}$ on both sides There can be at most $3|\mathcal{R}|/2$ such edges, hence the desired bound

$$w(\partial G') \leqslant 4 \cdot \frac{3}{2}|\mathcal{R}| \leqslant 2592t^2. \qquad \square$$

Lemma 5.10 (Polynomial bounds on Certificate 2 for fibres with inessential guts)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M, let F be a transversely oriented connected normal surface in M, and let a_0 and a_1 be two normal curves in F. Let $\Sigma_{\mathrm{fib}} \in \mathfrak{S}^*_{\mathrm{fib}}(M,F)$ and $\Sigma \in \mathfrak{S}^*_{\mathrm{mon}}(M,F,\Sigma_{\mathrm{fib}},a_0,a_1)$ be valid certificates. Suppose that every curve contained in $\mathrm{thick}(F'_g)$ is inessential in F. Then, if we denote by t the number of tetrahedra of $\mathcal{T}$, the curve a_1 is isotopic in F to a general position curve a'_1 with*

$$w(a'_1) \leqslant w(a_0) + 125\,712t^3 \cdot w(a_0 \cap F_g).$$

Proof. By Lemma 5.9, there exists a subsurface G' of F, obtained by removing a small neighbourhood of ∂F from thick(F'_g), such that $\partial G'$ is a general position multicurve in F with

$$w(\partial G') \leqslant 2592t^2. \tag{5.6}$$

By Proposition 4.12, we know that the isotopy class of a_1 in F is uniquely determined by M, F, and a_0. In particular, we can assume that Σ and a_1 are obtained by the construction in the proof of Proposition 4.11. The only properties that we will need from this construction is that the normal curve a_1 is disjoint from F_1, and it satisfies

$$\begin{aligned}
w(a_1 \setminus G') &\leqslant w(a_1 \cap F'_p) && \text{since } a_1 \subseteq G' \cup F'_p \\
&= w(\triangle(a'_0 \cap F_p)) && \text{by definition of } a_1 \\
&= w(a'_0 \cap F_p) && \text{since } \triangle \text{ is a simplicial isomorphism} \\
&\leqslant w(c) && \text{by definition of } a'_0 \\
&\leqslant w(a_0) && \text{by minimality of } c
\end{aligned} \tag{5.7}$$

and

$$\begin{aligned}
|a_1 \cap \partial G'| &= |a_1 \cap \partial F'_g| \\
&= |c_1 \cap \partial F'_g| && \text{by definition of } a_1 \\
&= |b_1 \cap \partial \partial_1 N| && \text{since } b_1 = f_g^{-1}(c_1^+) \\
&= |b''_1 \cap \partial \partial_1 N| && \text{since } b_1 \text{ and } b''_1 \text{ are isotopic fixing } \partial\partial_1 N \\
&= |b''_0 \cap \partial \partial_0 N| && \text{by Proposition 4.8} \\
&\leqslant w(b_0) && \text{since } b_0 \text{ and } b''_0 \text{ are isotopic fixing } \partial\partial_0 N \\
&\leqslant 97t \cdot w(a_0 \cap F_g) && \text{by (4.9).}
\end{aligned} \tag{5.8}$$

The intermediate steps above involve a variety of auxiliary objects introduced in the proof of Proposition 4.11, but we will only make use of the two final inequalities, where these objects do not appear.

We now describe a procedure to "simplify" a_1 with respect to G'. Suppose that a component s of $a_1 \cap G'$ satisfies $w(s) > w(\partial G')$. Note that a_1 is not fully contained in G', since a_1 is essential in F and G' does not contain any curves that are essential in F. Therefore, we see that s is a subarc of a_1. Let b be a component of $\partial G'$ intersecting s. Let s' be the smallest subarc of a_1 that contains s and such that $\partial s' \subseteq b$ (possibly $s' = s$). Since b is inessential in F, both intersections of s' with b occur from the same side of b – namely, from the side of G'. Let G be the component of G' containing b; we

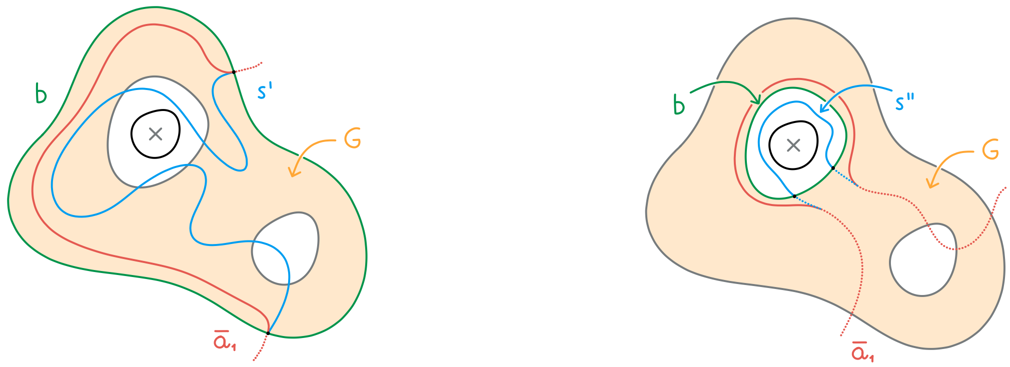

Figure 5.9. (a) If b bounds a disc or a collared annulus in F that contains G, then the subarc s' of a_1 can be isotoped to reduce its weight. (b) Otherwise, a subarc s'' of a_1 lying outside G can be isotoped inside G, reducing the number of intersections of a_1 with $\partial G'$.

analyse two cases.

- If b bounds a disc or a collared annulus in F that contains G, then s' can be isotoped fixing its endpoints to a subarc of b, and then pushed slightly inside G. By Lemma 2.4, since $w(b) \leqslant w(\partial G') < w(s')$, we deduce that a_1 can be isotoped to a general position curve $\overline{a}_1$ satisfying

$$w(\overline{a}_1 \setminus G') \leqslant w(a_1 \setminus G'), \quad |\overline{a}_1 \cap \partial G'| \leqslant |a_1 \cap \partial G'|, \quad w(\overline{a}_1) < w(a_1).$$

- Otherwise, the curve b bounds a disc or a collared annulus in F that does not contain G – and, in fact, only intersects G in b. In this case, let s'' be a subarc of $\operatorname{clos}(a_1 \setminus G)$ that shares an endpoint with s' and such that $s'' \cap b = \partial s''$; in particular, we have that $s'' \cap G = \partial s''$. We see that s'' can be isotoped fixing its endpoints to a subarc of b, and then pushed slightly inside G. By Lemma 2.4, we deduce that a_1 can be isotoped to a general position curve $\overline{a}_1$ satisfying

$$w(a_1 \setminus G') \leqslant w(a_1 \setminus G'), \quad |\overline{a}_1 \cap \partial G'| \leqslant |a_1 \cap \partial G'| - 2, \quad w(\overline{a}_1) \leqslant w(a_1) + w(\partial G').$$

The two constructions described above are depicted in Figure 5.9. We also note that, in both cases, the curve $\overline{a}_1$ can be chosen to be transverse to $\partial G'$.

We can now replace a_1 with $\overline{a}_1$, and iterate the procedure until every component s of

$a_1 \cap G'$ satisfies $w(s) \leqslant w(\partial G')$. Note that this procedure has to terminate, since the quantity

$$(w(\partial G') + 1) \cdot |a_1 \cap \partial G'| + w(a_1)$$

decreases at each step. Denote by a_1' the final general position curve obtained by this procedure. In particular, this curve satisfies:

- a_1' is isotopic to a_1 in F;
- $w(a_1' \setminus G') \leqslant w(a_1 \setminus G')$;
- $|a_1' \cap \partial G'| \leqslant |a_1 \cap \partial G'|$;
- each component s of $a_1' \cap G'$ has weight $w(s) \leqslant w(\partial G')$.

We then obtain the bound

$$\begin{aligned} w(a_1') &= w(a_1' \setminus G') + w(a_1' \cap G') \\ &\leqslant w(a_1' \setminus G') + w(\partial G') \cdot |a_1' \cap \partial G'|/2 \\ &\leqslant w(a_1 \setminus G') + w(\partial G') \cdot |a_1 \cap \partial G'|/2 \\ &\leqslant w(a_0) + w(\partial G') \cdot 97t \cdot w(a_0 \cap F_g)/2 && \text{by (5.7) and (5.8)} \\ &\leqslant w(a_0) + 125\,712t^3 \cdot w(a_0 \cap F_g) && \text{by (5.6),} \end{aligned}$$

as desired. ▢

Proposition 5.11 (Least-weight normal fibres of hyperbolic 3-manifolds contain short essential curves)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected orientable fibred hyperbolic 3-manifold M with t tetrahedra. Let F be a least-weight normal fibre of M. Then there is an essential normal curve a in F with*

$$w(a) \leqslant 2^{53} t^7 \cdot \chi(F)^2.$$

Proof. Fix an arbitrary transverse orientation of F. By Proposition 3.9, there exists a valid certificate $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$. Lemma 5.9 implies that there is a subsurface G of F, obtained by removing a small neighbourhood of ∂F from thick(F_g), such that G intersects at most $432t^2$ triangles of F, and moreover ∂G is a general position multicurve in F with

$$w(\partial G) \leqslant 2592t^2.$$

If thick(F_g) contains a curve that is essential in F, then we can apply Lemma 5.8 to conclude that G contains a general position curve that is essential in F and has weight at most

$$w(\partial G) + 2 \cdot 3 \cdot 432t^2 \leqslant 5184t^2 \leqslant 2^{53}t^7 \cdot \chi(F)^2.$$

A similar conclusion holds if thick(F'_g) contains a curve that is essential in F. Therefore, we can assume that every curve contained in thick(F_g) or in thick(F'_g) is inessential in F.

Let $\varphi \in \mathrm{Mcg}(F)$ be the monodromy of F. Fix a least-weight essential normal curve a_0 in F. For $i \geqslant 0$, we inductively choose a normal curve a_{i+1} in F such that a_{i+1} is isotopic in F to $\varphi(a_i)$, it is least-weight, and it satisfies

$$w(a_{i+1}) \leqslant w(a_i) + 125\,712t^3 \cdot w(a_i \cap F_g). \tag{5.9}$$

The existence of such a curve is guaranteed by Proposition 2.3 and Lemma 5.10.

We are now in a position to apply Proposition 5.6 to the subsurface G of F. Note that, trivially, G is a union of at most area(F_g) components; recall that area$(F_g) \leqslant 32t$ by Proposition 3.6. Moreover, since each a_i is a normal curve, we can assume that $w(a_i \cap G) = w(a_i \cap F_g)$. Denote by ω the smallest possible weight of a general position essential curve in F, and suppose that

$$\omega > 912\,384t^3 \tag{5.10}$$

(in particular, $\omega > 11|G| \cdot w(\partial G)$). For each $i \geqslant 0$, Proposition 5.6 then implies that

$$\begin{aligned} w(a_i \cap F_g) &\leqslant \frac{48|G|^2 \cdot w(\partial G)}{\omega} \cdot w(a_i) \\ &\leqslant \frac{127\,401\,984t^4}{\omega} \cdot w(a_i), \end{aligned}$$

and hence that

$$\begin{aligned} w(a_{i+1}) &\leqslant w(a_i) + 125\,712t^3 \cdot w(a_i \cap F_g) && \text{by (5.9)} \\ &\leqslant w(a_i) + \frac{16\,015\,958\,212\,608t^7}{\omega} \cdot w(a_i) \\ &= \left(1 + \frac{\alpha}{\omega}\right) \cdot w(a_i), \end{aligned} \tag{5.11}$$

where $\alpha = 16\,015\,958\,212\,608t^7$.

We can then estimate

$$\begin{aligned}
\operatorname{dist}_{\mathcal{C}(F)}(a_0, a_i) &\leqslant 2\log i(a_0, a_i) + 2 && \text{by Proposition 5.3} \\
&\leqslant 2\log(w(a_0)\cdot w(a_i)) + 2 && \text{by Proposition 2.2} \\
&\leqslant 2\log\left(\left(1+\frac{\alpha}{\omega}\right)^i \cdot w(a_0)^2\right) + 2 && \text{by (5.11)} \\
&= 2i \cdot \log\left(1+\frac{\alpha}{\omega}\right) + 4\log w(a_0) + 2. && (5.12)
\end{aligned}$$

Since M is hyperbolic, the monodromy φ is pseudo-Anosov. In particular, we have that

$$\begin{aligned}
\frac{1}{162\chi(F)^2} &\leqslant \liminf_{i\to\infty} \frac{\operatorname{dist}_{\mathcal{C}(F)}(a_0, a_i)}{i} && \text{by Proposition 5.4} \\
&\leqslant 2\log\left(1+\frac{\alpha}{\omega}\right) && \text{by (5.12)} \\
&\leqslant \frac{2\alpha}{\omega} && \text{since } \log(1+x)\leqslant x \text{ for } x \geqslant 0.
\end{aligned}$$

By rearranging, we conclude that

$$\begin{aligned}
\omega &\leqslant 324\alpha \cdot \chi(F)^2 \\
&= 5\,189\,170\,460\,884\,992 t^7 \cdot \chi(F)^2 \\
&\leqslant 2^{53} t^7 \cdot \chi(F)^2 && \text{by elementary algebra.}
\end{aligned}$$

Note that, if (5.10) is not satisfied, then $\omega \leqslant 2^{53}t^7 \cdot \chi(F)^2$ holds anyway. Therefore, we conclude that there is an essential normal curve a in F with

$$w(a) \leqslant \omega \leqslant 2^{53}t^7 \cdot \chi(F)^2. \qquad \square$$

5.5 A certificate for hyperbolicity

We now describe in detail our certificates for hyperbolicity of fibred 3-manifolds. We have two different certificates: one (Certificate 3) for fibres with low complexity – namely, the torus with one boundary component and the sphere with 4 boundary components – and one (Certificate 4) for all other fibres. The low-complexity certificate consists of one starting curve a_0 and two iterates of it under the monodromy; thanks to Proposition 5.2, this is enough to certify that the monodromy is pseudo-Anosov. The general certificate, instead, requires n iterates, for some relatively large – but still polynomial – value of

n. For each certificate, we follow the usual pattern: after describing the certificate, we prove its existence, correctness, and algorithmic verification.

Certificate 3 (Certificate for hyperbolicity with low-complexity fibres)**.** Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M with t tetrahedra, let F be a transversely oriented connected normal surface in M, and let $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ be a valid certificate. We say that a certificate Σ lies in $\mathfrak{S}_{\text{hyp.low}}(M, F, \Sigma_{\text{fib}})$ if it consists of:

- normal curves a_0, a_1, and a_2 in F;
- certificates $\Sigma_0 \in \mathfrak{S}_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$ and $\Sigma_1 \in \mathfrak{S}_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_1, a_2)$.

We define the *size* of Σ to be the number

$$|\Sigma| = t + \log(w(F) + 1) + |\Sigma_0| + |\Sigma_1| + \sum_{i=0}^{2} \left(\text{area}(\text{supp}(a_i)) + \log(w(a_i) + 1)\right).$$

We say that Σ is *valid*, and write $\Sigma \in \mathfrak{S}^*_{\text{hyp.low}}(M, F, \Sigma_{\text{fib}})$, if the following conditions are satisfied:

(k.1) F is either a torus with one boundary component or a sphere with 4 boundary components;

(k.2) a_0 is essential in F;

(k.3) $i(a_0, a_2) > 2i(a_0, a_1)$;

(l.1) $w(a_i) \leqslant t^7 \cdot 2^{(8426t+64)\cdot i+55}$ for $0 \leqslant i \leqslant 2$;

(l.2) $\text{area}(\text{supp}(a_i)) \leqslant 32t \cdot i + 2^{55}t^7$ for $0 \leqslant i \leqslant 2$;

(m.1) $\Sigma_1 \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_0, a_1)$;

(m.2) $\Sigma_2 \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_1, a_2)$. ×

Proposition 5.12 (Existence of Certificate 3)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M, let F be a transversely oriented connected normal surface in M, and let $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ be a valid certificate. Suppose that M is hyperbolic, and that F is either a torus with one boundary component or a sphere with 4 boundary components. Then $\mathfrak{S}^*_{\text{hyp.low}}(M, F, \Sigma_{\text{fib}})$ is non-empty.*

Proof. By Proposition 3.11, the 3-manifold M is fibred, and F is a fibre of M. Let $\varphi \in \text{Mcg}(F)$ be the monodromy of F, which is pseudo-Anosov by Theorem 5.1. Denote by t the number of tetrahedra of $\mathcal{T}$; by Proposition 5.11, there is an essential normal curve a_0 in F with

$$\mathrm{supp}(a_0) \leqslant w(a_0) \leqslant 2^{53} t^7 \chi(F)^2 \leqslant 2^{55} t^7.$$

Applying Proposition 4.11 twice, we find normal curves a_1 and a_2 and valid certificates $\Sigma_1 \in \mathfrak{S}^*_{\mathrm{mon}}(M, F, \Sigma_{\mathrm{fib}}, a_0, a_1)$ and $\Sigma_2 \in \mathfrak{S}^*_{\mathrm{mon}}(M, F, \Sigma_{\mathrm{fib}}, a_1, a_2)$. In particular, properties (l.1) and (l.2) are satisfied. Finally, property (k.3) is a consequence of Proposition 5.2 and the fact that φ is pseudo-Anosov. ▢

Proposition 5.13 (Correctness of Certificate 3)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M, and let F be a transversely oriented connected normal surface in M. Let $\Sigma_{\mathrm{fib}} \in \mathfrak{S}^*_{\mathrm{fib}}(M, F)$ and $\Sigma \in \mathfrak{S}^*_{\mathrm{hyp.low}}(M, F, \Sigma_{\mathrm{fib}})$ be valid certificates. Then M is hyperbolic.*

Proof. We know by Proposition 3.11 that M is fibred, and that F is a fibre of M. Let $\varphi \in \mathrm{Mcg}(F)$ be the monodromy of F. Note that, by Proposition 4.12 and properties (m.1) and (m.2), we have that a_i is isotopic in F to $\varphi^i(a_0)$ for $0 \leqslant i \leqslant 2$. Therefore, we can apply Proposition 5.2 to deduce that φ is pseudo-Anosov by property (k.3). We conclude that M is hyperbolic by Theorem 5.1. ▢

Proposition 5.14 (Verification of Certificate 3)**.** *There is an algorithm that takes as input*

- *a triangulation of a compact connected oriented 3-manifold M,*
- *a transversely oriented connected normal surface F in M,*
- *a valid certificate $\Sigma_{\mathrm{fib}} \in \mathfrak{S}^*_{\mathrm{fib}}(M, F)$, and*
- *a certificate $\Sigma \in \mathfrak{S}_{\mathrm{hyp.low}}(M, F, \Sigma_{\mathrm{fib}})$,*

and decides whether $\Sigma \in \mathfrak{S}^*_{\mathrm{hyp.low}}(M, F, \Sigma_{\mathrm{fib}})$. The running time of the algorithm is polynomial in $|\Sigma|$.

Proof. Property (k.1) can be verified using Proposition 2.15. Property (k.2) can be checked with the algorithm of Proposition 2.14. Property (k.3) can be verified using Proposition 2.21. Properties (l.1) and (l.2) are just inequalities, that can be verified directly. Finally, properties (m.1) and (m.2) can be verified using the algorithm of Proposition 4.13. ▢

Certificate 4 (Certificate for hyperbolicity)**.** Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M with t tetrahedra, let F be a transversely oriented connected normal surface in M, and let $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ be a valid certificate. Let $n = 2^{49}\chi(F)^6 \cdot t$. We say that a certificate Σ lies in $\mathfrak{S}_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$ if it consists of:

- normal curves $a_0, \ldots, a_n$ in F;
- certificates $\Sigma_i \in \mathfrak{S}_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_{i-1}, a_i)$ for $1 \leqslant i \leqslant n$.

We define the *size* of Σ to be the number

$$|\Sigma| = t + \log(w(F) + 1) + \sum_{i=0}^{n} (\text{area}(\text{supp}(a_i)) + \log(w(a_i) + 1)) + \sum_{i=1}^{n} |\Sigma_i|.$$

Let d be the output of the algorithm of Proposition 5.5 applied to the 3-manifold M, the normal surface F, and the normal curves a_0 and a_n. We say that Σ is *valid*, and write $\Sigma \in \mathfrak{S}^*_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$, if the following conditions are satisfied:

(n.1) $\chi(F) \leqslant -2$ and F is not a sphere with 4 boundary components;

(n.2) a_0 is essential in F;

(n.3) $d > 1\,683\,266\,941|\chi(F)|^3 \cdot t$;

(o.1) $w(a_i) \leqslant \chi(F)^2 \cdot t^7 \cdot 2^{(8426t+63)\cdot i+53}$ for $0 \leqslant i \leqslant n$;

(o.2) $\text{area}(\text{supp}(a_i)) \leqslant 32t \cdot i + 2^{53}\chi(F)^2 \cdot t^7$ for $0 \leqslant i \leqslant n$;

(p.1) $\Sigma_i \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_{i-1}, a_i)$ for $1 \leqslant i \leqslant n$. ×

Proposition 5.15 (Existence of Certificate 4)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M, let F be a transversely oriented connected normal surface in M, and let $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ be a valid certificate. Suppose that M is hyperbolic, that $\chi(F) \leqslant -2$, and that F is not a sphere with 4 boundary components. Then $\mathfrak{S}^*_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$ is non-empty.*

Proof. By Proposition 3.11, the 3-manifold M is fibred, and F is a fibre of M. Let $\varphi \in \text{Mcg}(F)$ be the monodromy of F; by Theorem 5.1, we know that φ is pseudo-Anosov. Denote by t the number of tetrahedra of $\mathcal{T}$; by Proposition 5.11, there is an essential normal curve a_0 in F with

$$\text{supp}(a_0) \leqslant w(a_0) \leqslant 2^{53}t^7 \cdot \chi(F)^2.$$

Consistently with Certificate 4, let $n = 2^{49}\chi(F)^6 \cdot t$. By Proposition 4.11, we can inductively find normal curves $a_1, \ldots, a_n$ in F and, for each $1 \leqslant i \leqslant n$, a valid certificate $\Sigma_i \in \mathfrak{S}^*_{\text{mon}}(M, F, \Sigma_{\text{fib}}, a_{i-1}, a_i)$. In particular, for each $1 \leqslant i \leqslant n$, the essential curve a_i is isotopic in F to $\varphi(a_{i-1})$ (see Proposition 4.12), and satisfies properties (o.1) and (o.2).

All that remains is to verify property (n.3). Since the monodromy φ is pseudo-Anosov, we have that

$$\begin{aligned} n &\leqslant 162\chi(F)^2 \cdot \text{dist}_{\mathcal{C}(F)}(a_0, a_n) && \text{by Proposition 5.4} \\ &\leqslant 162\chi(F)^2 \cdot (1313|\chi(F)| \cdot d + 3\,727\,496|\chi(F)|^3) && \text{by Proposition 5.5.} \end{aligned}$$

It follows that

$$\begin{aligned} d &\geqslant \frac{1}{1313|\chi(F)|} \cdot \left(\frac{n}{162\chi(F)^2} - 3\,727\,496|\chi(F)|^3 \right) \\ &\geqslant \frac{1}{1313|\chi(F)|} \cdot \left(\frac{n}{162\chi(F)^2} - 1\,863\,748\chi(F)^4 \cdot t \right) \\ &> 1\,683\,266\,941|\chi(F)|^3 \cdot t && \text{by elementary algebra.} \end{aligned}$$

In particular, we see that property (n.3) is satisfied. ▢

Proposition 5.16 (Correctness of Certificate 4)**.** *Let $\mathcal{T}$ be a triangulation of a compact connected oriented 3-manifold M, and let F be a transversely oriented connected normal surface in M. Let $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ and $\Sigma \in \mathfrak{S}^*_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$ be valid certificates. Then M is hyperbolic.*

Proof. We know by Proposition 3.11 that M is fibred, and that F is a fibre of M. By property (n.1), the surface F satisfies $\chi(F) \leqslant -2$; in particular, Theorem 5.1 shows that M is hyperbolic if and only if the monodromy φ of F is pseudo-Anosov. Therefore, it suffices to show that φ is not periodic nor reducible. Note that, by Proposition 4.12, we have that a_i is isotopic in F to $\varphi^i(a_0)$ for $1 \leqslant i \leqslant n$.

Let $1 \leqslant j \leqslant 49\,572|\chi(F)|^3$ be an integer. Then we have the following inequalities:

$$\begin{aligned} &\text{dist}_{\mathcal{C}(F)}(a_0, \varphi^j(a_0)) \\ &\leqslant 2\,\text{dist}_{\mathcal{C}(F)}(a_0, \varphi^j(a_0)) + 2 \\ &\leqslant 4 \log i(a_0, a_j) + 6 && \text{by Proposition 5.3} \\ &\leqslant 4 \log(w(a_0) \cdot w(a_j)) + 6 && \text{by Proposition 2.2} \\ &\leqslant 4 \log(\chi(F)^4 \cdot t^{14} \cdot 2^{(8426t+63)\cdot j+106}) + 6 && \text{by property (o.1)} \end{aligned}$$

$$
\begin{aligned}
&\leqslant 16 \log|\chi(F)| + 52 \log t + (33\,704t + 252) \cdot j + 430 \\
&\leqslant 16|\chi(F)| + 52t + (33\,704t + 252) \cdot 49\,572|\chi(F)|^3 + 430 \\
&\leqslant 1\,683\,266\,941|\chi(F)|^3 \cdot t && \text{by elementary algebra} \\
&< d && \text{by property (n.3)} \\
&\leqslant \mathrm{dist}_{\mathcal{C}(F)}(a_0, \varphi^n(a_0)) && \text{by Proposition 5.5.}
\end{aligned}
$$

Since this holds for every choice of j, Proposition 5.4 implies that φ is not periodic nor reducible. ▢

Proposition 5.17 (Verification of Certificate 4)**.** *There is an algorithm that takes as input*

- *a triangulation of a compact connected oriented* 3*-manifold* M*,*
- *a transversely oriented connected normal surface* F *in* M*,*
- *a valid certificate* $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$*, and*
- *a certificate* $\Sigma \in \mathfrak{S}_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$*,*

and decides whether $\Sigma \in \mathfrak{S}^*_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$*. The running time of the algorithm is polynomial in* $|\Sigma|$*.*

Proof. Property (n.1) can be verified using Proposition 2.15. Property (n.2) can be checked with the algorithm of Proposition 2.14. Properties (n.3), (o.1), and (o.2) are just equalities and inequalities, that can be verified directly (after computing the integer d with Proposition 5.5). Finally, property (p.1) can be verified using the algorithm of Proposition 4.13. ▢

5.6 Algorithmic hyperbolicity detection for fibred 3-manifolds

We now have all the tools we need to state and prove our main result about certifying hyperbolicity of fibred 3-manifolds. Informally speaking, the statement is that, if a 3-manifold is fibred and hyperbolic, then this can be certified in polynomial time in the size of the triangulation and the Euler characteristic of a fibre. However, the number of binary digits required to describe a triangulated 3-manifold is (almost) proportional to the size of the triangulation, and it is independent of the complexity of a fibre. Therefore, we cannot quite say that deciding hyperbolicity of fibred 3-manifolds is in NP. Since the easiest way to state a certification result is to place a problem in the complexity class NP, we need to modify the statement of the decision problem, to include the complexity of a fibre as part of the input. This is the reason why we add the seemingly artificial bound B on the Euler characteristic of a fibre in the statement of Hyperbolicity detection for fibred 3-manifolds with bound on the fibre below. Since the number of binary digits required to describe an integer B is logarithmic in B, and the complexity of our certificate is polynomial in the Euler characteristic of a fibre, we require that said Euler characteristic is bounded by the logarithm of B.

Problem (Hyperbolicity detection for fibred 3-manifolds with bound on the fibre)**.**

Input: a triangulation of a compact connected oriented 3-manifold M and an integer $B \geqslant 1$.

Output: whether M is hyperbolic and fibres over the circle with fibre F such that $|\chi(F)| \leqslant \log B$.

The size of the input is measured by the number of tetrahedra in the triangulation of M and the logarithm of B. ×

Theorem 5.18 (Hyperbolicity detection for fibred 3-manifolds with bound on the fibre is in NP.)**.** *The problem Hyperbolicity detection for fibred 3-manifolds with bound on the fibre is in* NP*.*

Proof. The verification algorithm takes as input a triangulation $\mathcal{T}$ of a compact connected oriented 3-manifold M, an integer $B \geqslant 1$, and a certificate consisting of three parts. The first part is a transversely oriented normal surface F in M; the algorithm verifies that F is connected and orientable, using Propositions 2.15 and 2.16; moreover, the algorithm checks that $0 > \chi(F) > -\log(B)$. The second part is a certificate $\Sigma_{\text{fib}} \in \mathfrak{S}_{\text{fib}}(M, F)$; the algorithm verifies that $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ using Proposition 3.12. The third part is a certificate Σ_{hyp}, that lies in $\mathfrak{S}_{\text{hyp.low}}(M, F, \Sigma_{\text{fib}})$ if F is a torus with one boundary component or a sphere with 4 boundary components, and in $\mathfrak{S}_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$ otherwise; the algorithm verifies that $\Sigma_{\text{hyp}} \in \mathfrak{S}^*_{\text{hyp.low}}(M, F, \Sigma_{\text{fib}})$ or $\Sigma_{\text{hyp}} \in \mathfrak{S}^*_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$ using Proposition 5.14 or Proposition 5.17 respectively. All these verifications can be performed in polynomial time in $|\mathcal{T}|$, $\log B$, $|\chi(F)|$, $|\Sigma_{\text{fib}}|$, and $|\Sigma_{\text{hyp}}|$. If all the checks are successful, then Proposition 3.11 implies that M fibres over the circle with fibre F. Moreover, Proposition 5.13 or Proposition 5.16 then imply that M is hyperbolic.

Conversely, if M is hyperbolic and fibres over the circle with fibre F_0 such that $|\chi(F_0)| \leqslant \log B$, then Propositions 3.11 and 3.13 guarantee that there exists a transversely oriented connected normal surface F in M and a certificate $\Sigma_{\text{fib}} \in \mathfrak{S}^*_{\text{fib}}(M, F)$ such that $|\chi(F)| \leqslant \log B$ and $|\Sigma_{\text{fib}}|$ is bounded above by a polynomial in $|\mathcal{T}|$. If F_0 is a torus with one boundary component or a sphere with 4 boundary components, then Proposition 5.12 guarantees the existence of a certificate $\Sigma_{\text{hyp}} \in \mathfrak{S}^*_{\text{hyp.low}}(M, F, \Sigma_{\text{fib}})$ such that $|\Sigma_{\text{hyp}}|$ is bounded above by a polynomial in $|\mathcal{T}|$. Otherwise, Proposition 5.15 guarantees the existence of a certificate $\Sigma_{\text{hyp}} \in \mathfrak{S}^*_{\text{hyp}}(M, F, \Sigma_{\text{fib}})$ such that $|\Sigma_{\text{hyp}}|$ is bounded above by a polynomial in $|\mathcal{T}|$ and $\log B$. ▢

If we are in a setting where the Euler characteristic of a fibre is guaranteed to be bounded above by the size of the input, then the artificial bound B we introduced in the statement of Hyperbolicity detection for fibred 3-manifolds with bound on the fibre is not needed. This happens, for example, in the case of knot diagrams. A diagram of a knot K in S^3 can be represented as a 4-valent graph in the plane; the amount of information required to encode it is proportional to the number n of crossings in the diagram. Denote by M_K the 3-manifold obtained by removing an open regular neighbourhood of K from S^3. If K is *fibred* – that is, if M_K is fibred – and *hyperbolic* – that is, if M_K is hyperbolic – then we can certify this in polynomial time in the size of a triangulation of M_K and the Euler characteristic of a fibre, thanks to Theorem 5.18. It is not hard to construct a triangulation of M_K with linear size in n (see [13, Lemma 7.2]). Moreover, it is easy to see that the Euler characteristic of a fibre of M_K is at most n. Therefore, the certificate of fibredness and hyperbolicity of K can be verified in polynomial time in n; in other words, as we prove in Corollary 5.19 below, the problem of deciding hyperbolicity of fibred knots is in NP.

Problem (Hyperbolicity detection for fibred knots)**.**

Input: a diagram of a knot K in the 3-sphere.

Output: whether K is fibred and hyperbolic.

The size of the input is measured by the number of crossings in the diagram. ×

Corollary 5.19 (Hyperbolicity detection for fibred knots is in NP)**.** *The problem* Hyperbolicity detection for fibred knots *is in* NP.

Proof. Let D be a diagram of a knot $K \subseteq S^3$ with n crossings. The algorithm of [13, Lemma 7.2] computes a triangulation $\mathcal{T}$ of the 3-manifold M_K obtained by removing an open regular neighbourhood of K from S^3. This algorithm runs in polynomial time in n, and the number of tetrahedra of $\mathcal{T}$ is bounded above by a linear function of n.

We claim that K is fibred and hyperbolic if and only if the answer to Hyperbolicity detection for fibred 3-manifolds with bound on the fibre with input $M = M_K$ and $B = 2^{n+1}$ is "yes". In fact, by Seifert's algorithm, the knot K admits a Seifert surface F_0 with $|\chi(F_0)| \leqslant n+1$. Since every incompressible Seifert surface for K, intersected with M_K, is a fibre of M_K, this implies that M_K is fibred if and only if it admits a fibre F with $|\chi(F)| \leqslant n+1$.

Theorem 5.18 then implies that a positive answer to Hyperbolicity detection for fibred knots can be certified in polynomial time in $|\mathcal{T}|$ and $\log B$. Both quantities are bounded above by a linear function of n, from which we conclude that Hyperbolicity detection for fibred knots is in NP. ▢

Appendix A

Proof of Proposition 4.10

This appendix is entirely devoted to the proof of Proposition 4.10, whose statement we recall here for the reader's convenience.

Proposition 4.10 (Vertical surfaces with prescribed boundary)**.** *Let $\mathcal{T}$ be a suitable pre-sutured triangulation of $M = F \times [0,1]$ with t tetrahedra, where F is a compact orientable surface. Let a be a normal 1-manifold in $\partial_0 M \cup \partial_v M$. Denote by a_0 and a_v the intersections of a with $\partial_0 M$ and $\partial_v M$ respectively. Suppose that no component of a_0, seen as a 1-manifold in $\partial_0 M$, is a boundary-parallel arc or a curve bounding a disc; moreover, suppose that a_v is vertical in $\partial_v M$. Then there is a normal surface A in M such that:*

(i) A is a vertical surface in M;

(ii) $A \cap (\partial_0 M \cup \partial_v M)$ is isotopic to a in $\partial_0 M \cup \partial_v M$, fixing $\partial\partial_h M$;

(iii) the weight of A is bounded by

$$w(A) \leqslant 2^{9t+40} \cdot w(a).$$

We have decided to defer the proof of Proposition 4.10 to this appendix for several reasons. First, simply put, the proof is long enough that it would have disrupted the flow of the main text, without adding much to the exposition. Second, this is a statement that experts in the field would very readily believe (at least in spirit, ignoring the exact numbers); increasing the length of the main text by a substantial amount to prove a statement that is not particularly surprising would have been a disservice to the reader. Third, despite the alleged "obviousness" of the statement, our proof of Proposition 4.10 requires the introduction of two very technical concepts, that are necessary to overcome subtle issues that would otherwise arise in the argument. These concepts – namely, *$\mathcal{E}$-prioritised weight* and *twist number* – are not used anywhere else in the paper, and we believe that they are not particularly useful outside the context of this proof. Therefore, we have decided to introduce them in this appendix, where they can be presented in a self-contained manner without "polluting" the main text. Finally, in this appendix we will require that the reader is familiar with several concepts from Matveev's [28], that are not necessary to understand the rest of the paper.

The idea of the proof is quite straightforward. One of the cornerstones of normal surface theory is that most "interesting" surfaces in a triangulated 3-manifold can be represented by a fundamental normal surface (or by a bounded sum of fundamental normal surfaces, as was the case in our proof of Proposition 3.13). Since fundamental normal surfaces have bounded weight (see Proposition 2.6), it should follow that any "interesting" surface can be described with bounded amount of data. Of course, we cannot expect that every vertical surface in $M = F \times [0, 1]$ has a fundamental representative, since there are only finitely many fundamental normal surfaces in a triangulated 3-manifold, but infinitely many possible isotopy classes for the 1-manifold $a \subseteq \partial_0 M \cup \partial_v M$. However, this expectation can be satisfied if we adopt a more suitable notion of "fundamental".

Recall from Section 2.2.5 that normal surfaces are in one-to-one correspondence with non-negative integral solutions of a set of linear equations, namely the matching equations and the consistency equations. Fundamental normal surfaces correspond to solutions that cannot be written as a non-trivial sum of other solutions. Let us consider an additional set of equations, that we call *boundary equations*. For a given 1-manifold a properly embedded in $\partial_0 M \cup \partial_v M$, we say that a normal surface $A \subseteq M$ satisfies the boundary equations if $A \cap (\partial_0 M \cup \partial_v M)$ is an integer multiple of a (although this is slightly imprecise, as explained below). The set of matching equations, consistency equations, and boundary equations defines a system of linear equations, and we can consider fundamental solutions of this system. The bound on the weight of fundamental normal surfaces in Proposition 2.6 is a statement that relies purely on linear algebra, and in fact applies to any system of linear equations. Therefore, fundamental solutions of the system of matching equations, consistency equations, and boundary equations – call this the "augmented" linear system – will also have a bound on their weight, that depends –

as it should – both on the number t of tetrahedra and on the weight of a.

All that is left to do is showing the existence of a vertical surface that is a fundamental solution of the augmented linear system. The strategy to prove a claim of this kind has been streamlined by Matveev in [28], culminating in Theorem 4.1.36 therein (see Theorem A.2 below for the statement). The idea is then to take a least-weight normal vertical surface A in M such that $A \cap (\partial_0 M \cup \partial_v M)$ is isotopic to a preserving $\partial\partial_h M$, and prove that A must be a fundamental solution of the augmented linear system. If, for the sake of contradiction, the surface A were not fundamental, then it could be written as a normal sum $A = G_1 + G_2$, where G_1 and G_2 are non-trivial solutions of the augmented system. Since A is incompressible and boundary-incompressible, then [28, Theorem 4.1.36] implies that G_1 and G_2 are also incompressible and boundary-incompressible; moreover, one of them – say G_1 – will have the same intersection with $\partial_0 M \cup \partial_v M$ as A. This means that G_1 is vertical and, in particular, admissibly isotopic to A; since A is least-weight, we obtain a contradiction.

There are two issues with this argument. The first one is that we have no control over the weight of $A \cap (\partial_0 M \cup \partial_v M)$. To be precise, the boundary equations are set up to ensure that the intersection of a normal surface with $\partial_0 M \cup \partial_v M$ is an integer multiple of $A \cap (\partial_0 M \cup \partial_v M)$, not of a; this is necessary for the argument to go through. In particular, this means that the bound we get on the weight of fundamental solutions depends on the weight of $A \cap (\partial_0 M \cup \partial_v M)$. However, if A is chosen to be least-weight, there is no reason a priori to expect that the weight of $A \cap (\partial_0 M \cup \partial_v M)$ should be bounded in terms of the weight of a. In fact, it is entirely conceivable that the optimal way to minimise the weight of A in M requires the weight of $A \cap (\partial_0 M \cup \partial_v M)$ to be arbitrarily large compared to that of a. On the other hand, the hypothesis that A is least-weight is necessary to apply [28, Theorem 4.1.36]. To overcome this issue, we show in Section A.1 that Matveev's Theorem 4.1.36 also applies to normal surfaces that minimise, in their admissible isotopy class, first the weight of their intersection with $\partial_0 M \cup \partial_v M$, and then their weight in M; we call these *least-*$w_{\mathcal{E},*}$ *surfaces*. It is immediate to check that, if A is chosen to be least-$w_{\mathcal{E},*}$, then

$$w(A \cap (\partial_0 M \cup \partial_v M)) \leqslant w(a).$$

This will allow us to bound the weight of fundamental solutions in terms of t and $w(a)$.

The second problem is that in the definition of least-weight (or, for that matter, least-$w_{\mathcal{E},*}$) surfaces, we allow admissible isotopies, that is, isotopies that preserve – but don't necessarily fix pointwise – the boundary pattern $\partial\partial_h M$. This is unavoidable if we want [28, Theorem 4.1.36] (or the generalisation we prove in Theorem A.9) to hold. As a consequence, the surface A we obtain from the argument outlined above will not, in general, satisfy condition (ii) of Proposition 4.10; the only guarantee is that $A \cap (\partial_0 M \cup \partial_v M)$ is isotopic to a preserving $\partial\partial_h M$. We address this issue in Proposition A.16, by showing that, under these conditions, there exists a surface that is admissibly isotopic to A, satisfies condition (ii) of Proposition 4.10, and has weight bounded above by a polynomial function of t, $w(a)$, and $w(A)$.

A.1 Least-$w_{\mathcal{E},*}$ normal surfaces and normal sums

Definition A.1 ($\mathcal{E}$-prioritised weight of a surface)**.** Let F be a general position surface properly embedded in a triangulated 3-manifold with boundary pattern, and let $\mathcal{E}$ be a collection of edges of the triangulation. The *$\mathcal{E}$-weight* of F is the number $w_{\mathcal{E}}(F) = |\mathcal{E} \cap F|$ of intersections between F and the edges of $\mathcal{E}$. The *$\mathcal{E}$-prioritised weight* of F is the ordered pair

$$w_{\mathcal{E},*}(F) = (w_{\mathcal{E}}(F), w(F)).$$

We say that F is *least-$w_{\mathcal{E},*}$* if it minimises the $\mathcal{E}$-prioritised weight amongst all general position properly embedded surfaces that are admissibly isotopic to F, where $\mathcal{E}$-prioritised weights are compared lexicographically. ×

Our goal, achieved in Theorem A.9, is to prove a generalisation of [28, Theorem 4.1.36]. The original statement is as follows.

Theorem A.2 ([28, Theorem 4.1.36])**.** *Let a least-weight ("minimal" in [28]) connected normal surface F in an irreducible boundary-irreducible triangulated 3-manifold with boundary pattern (M, Γ) be presented in the form $F = G_1 + G_2$. If F is incompressible and boundary-incompressible, then so are G_1 and G_2. Moreover, neither G_1 nor G_2 is a sphere, a projective plane, or a disc intersecting Γ at most once.*

The version we require differs from [28, Theorem 4.1.36] in two ways:

1. we drop the assumption that F is connected;

2. we replace the assumption that F is least-weight with the assumption that F is least-$w_{\mathcal{E},*}$ for a given collection of edges $\mathcal{E}$.

The proof of this modified statement does not involve any new ideas or techniques that are not already present in Matveev's [28]. In fact, it suffices to review the intermediate results used in the proof of [28, Theorem 4.1.36] and verify that the arguments can be adapted to our new setting. We will try to strike a reasonable balance between brevity and completeness: we will not replicate the full proofs of the modified intermediate results, but we will highlight all the necessary changes. We recommend that this section be read alongside Sections 3.3 and 4.1 of [28]. In particular, we expect the reader to be familiar with the notions of *returns*, *regular* and *irregular switches*, *patches* (and the related terminology such as *double curves*, *trace curves*, *opposite*, *adjacent*, *clean disc patches*, *companions*), *reduced form*, *compressing and boundary-compressing discs for* 2*-dimensional subpolyhedra*, *good* and *bad angles* for such discs, and *weight* of a compressing or boundary-compressing disc of a surface in a normal sum.

Remark A.3 (Essential surfaces are isotopic to least-$w_{\mathcal{E},*}$ normal surfaces)*.* It follows from Proposition 2.7 that every incompressible boundary-incompressible general position properly embedded surface F in a triangulated compact irreducible boundary-irreducible 3-manifold with boundary pattern is admissibly isotopic to a least-$w_{\mathcal{E},*}$ normal surface F', provided that F does not have any component that is a sphere, a clean disc, or an inessential semi-clean disc. Moreover, if F contains a return, then F' can be chosen so that $w(F') < w(F)$ and $w_{\mathcal{E},*}(F') < w_{\mathcal{E},*}(F)$; this is explained in [28, Remark 3.2.23]. ×

Remark A.4 (Least-$w_{\mathcal{E},*}$ normal surfaces have no sphere, clean disc, or inessential semi-clean disc components)*.* Let (M, Γ) be an irreducible boundary-irreducible triangulated 3-manifold with boundary pattern, and let $\mathcal{E}$ be a collection of edges of the triangulation. Let F be a least-$w_{\mathcal{E},*}$ normal surface. Then F does not have any sphere, clean disc, or inessential semi-clean disc components. In fact, if S were a sphere component, then S would bound a 3-ball by irreducibility of M, and could therefore be isotoped inside a tetrahedron of the triangulation, thus reducing the weight of F. Similarly, by boundary-irreducibility of (M, Γ), a clean disc component could be admissibly isotoped to lie inside a tetrahedron, thus reducing the weight of F. Finally, if D is an inessential semi-clean disc component of F, then it can be admissibly isotoped to a disc D' linked (see [28, Definition 3.3.17]) to an edge e contained in Γ; if at least one of the intersections of D with Γ is not contained in an edge of $\mathcal{E}$, then we can pick $e \notin \mathcal{E}$, so that $w_{\mathcal{E}}(D') = 0$. Either way, we have that $w(D') = 2 < w(D)$ and $w_{\mathcal{E}}(D') \leqslant w_{\mathcal{E}}(D)$, contradicting the minimality of $w_{\mathcal{E},*}(F)$. ×

Lemma A.5 (Matveev's Lemma 4.1.2 revisited)**.** *Let M be a triangulated* 3*-manifold, and let $\mathcal{E}$ be a collection of edges of the triangulation. Let a normal surface F in M be* *presented in the form $F = G_1 + G_2$. Suppose that $G_1 \cup G_2$ has a self-opposite disc patch E. Then M contains a projective plane whose $\mathcal{E}$-prioritised weight is less than that of F.*

Proof. The proof of [28, Lemma 4.1.2] works by noting that the connected component P of – say – G_1 containing E is a projective plane with $w(P) \leqslant w(G_1) < w(F)$. However, we also have that $w_{\mathcal{E}}(P) \leqslant w_{\mathcal{E}}(G_1) \leqslant w_{\mathcal{E}}(F)$, so that $w_{\mathcal{E},*}(P) < w_{\mathcal{E},*}(F)$. ▢

Lemma A.6 (Matveev's Lemma 4.1.5 revisited)**.** *Let (M, Γ) be an irreducible boundary-irreducible triangulated 3-manifold with boundary pattern, and let $\mathcal{E}$ be a collection of edges of the triangulation. Let a least-$w_{\mathcal{E},*}$ normal surface F in (M, Γ) be presented in the form $F = G_1 + G_2$. Let E be a non-self-opposite clean disc patch of $G_1 \cup G_2$. Then E does not admit an opposite clean companion disc.*

Proof. Suppose for a contradiction that E admits an opposite clean companion disc E', and denote by s the double curve of $G_1 \cup G_2$ that corresponds to the twin trace curves in ∂E and $\partial E'$.

Suppose first that E' does not contain E. Then $S = E \cup E'$ is a sphere or a properly embedded clean disc; since (M, Γ) is irreducible and boundary-irreducible, the surface S bounds a 3-ball (if S it is a sphere) or cobounds a 3-ball with a clean disc in ∂M (if S is a disc). In either case, we see that F has a component that is entirely contained in a clean 3-ball, and is therefore a sphere or a clean disc; by Remark A.4, this contradicts the assumption that F is least-$w_{\mathcal{E},*}$.

Therefore, we have that E' contains E. The proof of [28, Lemma 4.1.5] then shows that, by replacing the regular switch along s with the irregular one, we get the disjoint union of two surfaces F_1 and T, where F_1 is admissibly isotopic to F. Note that $w_{\mathcal{E},*}(F_1 \cup T) = w_{\mathcal{E},*}(F)$, and in particular $w_{\mathcal{E},*}(F_1) \leqslant w_{\mathcal{E},*}(F)$. Moreover, one of F_1 and T must contain a return. If F_1 contains a return, then by Remark A.3 we can admissibly isotope F_1 to reduce its boundary-prioritised weight, contradicting the minimality of $w_{\mathcal{E},*}(F)$. If T contains a return, then $w(T) > 0$, and therefore $w_{\mathcal{E},*}(F_1) < w_{\mathcal{E},*}(F)$, which again leads to a contradiction. ▢

Lemma A.7 (Matveev's Lemma 4.1.8 revisited)**.** *Let (M, Γ) be an irreducible boundary-irreducible triangulated 3-manifold with boundary pattern, and let $\mathcal{E}$ be a collection of edges of the triangulation. Let a least-$w_{\mathcal{E},*}$ normal surface F in (M, Γ) be presented in a reduced form $F = G_1 + G_2$. Suppose that F is incompressible and boundary-incompressible. Then all the patches of $G_1 \cup G_2$ are incompressible and boundary-incompressible, and none of them is a clean disc patch.*

Proof. The proof is essentially the same as that of [28, Lemma 4.1.8]. The only changes we need to implement are:

- use Lemma A.5 instead of [28, Lemma 4.1.2];
- use Lemma A.6 instead of [28, Lemma 4.1.5];
- observe that conclusion 4 of [28, Lemma 4.1.4] cannot hold if F is least-$w_{\mathcal{E},*}$. ▢

Lemma A.8 (Matveev's Lemma 4.1.33 revisited)**.** *Let (M, Γ) be an irreducible boundary-irreducible triangulated 3-manifold with boundary pattern, and let $\mathcal{E}$ be a collection of edges of the triangulation. Let a least-$w_{\mathcal{E},*}$ normal surface F in (M, Γ) be presented in the form $F = G_1 + G_2$ such that there are no clean disc patches. Then $G_1 \cup G_2$ does not admit a compressing disc with precisely one bad angle.*

Proof. The proof of [28, Lemma 4.1.33] works by constructing a surface $F_1 = (F \setminus \partial T) \cup A$ that is admissibly isotopic to F. The surface A does not intersect any edge of the triangulation, while $F \cap \partial T$ does. Therefore, Matveev concludes that $w(F_1) < w(F)$, but by the very same argument we also have that $w_{\mathcal{E},*}(F_1) < w_{\mathcal{E},*}(F)$. ▢

Theorem A.9 (Matveev's Theorem 4.1.36 revisited)**.** *Let (M, Γ) be an irreducible boundary-irreducible triangulated 3-manifold with boundary pattern, and let $\mathcal{E}$ be a collection of edges of the triangulation. Let a least-$w_{\mathcal{E},*}$ normal surface F in (M, Γ) be presented in the form $F = G_1 + G_2$. Suppose that F is incompressible and boundary-incompressible, and that no component of F is a sphere, a projective plane, or a disc intersecting Γ at most once. Then the same is true of G_1 and G_2.*

Proof. It suffices to prove the theorem in the case where $F = G_1 + G_2$ is in reduced form. Then Lemma A.7 tells us that all the patches of $G_1 \cup G_2$ are incompressible and boundary-incompressible, and none of them is a clean disc patch. If, for the sake of contradiction, any component of G_1 or G_2 is a sphere, a projective plane, or a disc intersecting Γ at most once, then it must be decomposed into at least two patches, but every decomposition of such a component into patches must contain a clean disc patch.

We now prove that G_1 is incompressible and boundary-incompressible; by symmetry, the same will hold for G_2. Suppose for a contradiction that G_1 admits a non-trivial compressing or boundary-compressing disc Δ; amongst all such discs, choose one that minimises the weight $c(\Delta) + c_\partial(\Delta)$. Note that this weight cannot be zero, because otherwise Δ would be a non-trivial compressing or boundary-compressing disc for a patch of $G_1 \cup G_2$. If $\Delta \cap G_2$ contains a circle or an arc with both endpoints in ∂M, then an innermost circle or outermost arc argument yields a disc $D_0 \subseteq \Delta$ satisfying condition (1) of [28, Lemma 4.1.35], contradicting the minimality of Δ.

Therefore, $\Delta \cap G_2$ cuts Δ into regions that are homeomorphic to discs; we call such regions *polygons*. In particular, call a polygon *exterior* if it intersects ∂M, and *interior* otherwise. Denote by n the number of components of $\Delta \cap G_2$ that have one endpoint in ∂M, and by m the number of components that don't intersect ∂M. It is easy to see that there are exactly m interior polygons and $n+1$ exterior polygons.

Every interior (respectively exterior) polygon is a compressing (respectively boundary-compressing) disc for $G_1 \cup G_2$. Every point of $\partial\Delta \cap G_2$ that does not lie on ∂M is a common vertex of two angles belonging to two different polygons; one of these angles is good, and the other is bad. Therefore, the total number of bad angles in the polygons is $2m + n$. An easy counting argument shows that there must be either one exterior polygon with no bad angles, or one interior polygon with at most one bad angle. However, Lemma A.8 implies that we cannot have an interior polygon (that is, a compressing disc for $G_1 \cup G_2$) with precisely one bad angle. We conclude that there is a polygon (interior or exterior) with no bad angles. By [28, Lemma 4.1.34], we can find a disc $D_0 \subseteq M$ satisfying conditions (2) or (3) of [28, Lemma 4.1.35], contradicting the minimality of Δ. This concludes the proof, showing G_1 cannot have a non-trivial compressing or boundary-compressing disc. ▢

A.2 Admissible isotopies and isotopies fixing the boundary

The goal of this section is to convert admissible isotopies of interval bundles into isotopies that fix the boundary pattern, while maintaining some control over weights of vertical surfaces. More precisely, the setting is the following. Let $M = F \times [0,1]$ for some compact orientable surface F, and let A be a vertical surface in M. Suppose that A is isotopic preserving $\Gamma = \partial\partial_h M$ to a normal surface A'. For our purposes, the surface A will be a vertical surface such that $A \cap (\partial_0 M \cup \partial_v M)$ is the prescribed normal 1-manifold a, and A' will be a fundamental normal surface (with respect to the "augmented" system defined at the beginning of Appendix A). The aim is to find a normal surface A'' that is isotopic to A fixing Γ, with weight bounded in terms of $w(a)$, $w(A')$, and the number t of tetrahedra in the triangulation of M; in particular, we want the weight of A'' to be independent of the weight of A, over which we have no control.

The proof consists of three main steps.

1. We introduce the notion of *twist number* to quantify the amount of twisting along Γ that is induced by an admissible isotopy taking A to A'; see Definitions A.10 and A.12 and Proposition A.11.

2. Then, we show that the twist number of such an isotopy can be bounded in terms of the weights of the intersection of A and A' with $\partial_0 M \cup \partial_v M$, independently of $w(A)$; see Proposition A.13.

3. Finally, we prove that the twisting along Γ can be "undone" by an admissible isotopy of A' that increases its weight by a controlled amount, depending on t, $|A' \cap \Gamma|$, and the twist number; see Proposition A.14.

Since the twist number only depends on $w(a)$ and $w(A')$, combining these three steps will yield a normal surface A'' with the desired bound on its weight; this is achieved in Corollary A.15.

This outline is somewhat misleading: for technical reasons – arising from the fact that Γ contains curves that are isotopic in ∂M – obtaining the desired surface A'' in the setting of interval bundles require two successive applications of Corollary A.15; this is done in Proposition A.16, which contains exactly the statement that we need.

We use the following notation for isotopies. An isotopy φ of a space X is a continuous one-parameter family of homeomorphisms $\varphi_t \colon X \to X$, parametrised by $t \in [0,1]$; we write $\varphi = (\varphi_t)_{t\in[0,1]}$ to denote this parametrisation.

Definition A.10 (Twist number in S^1)**.** Let $\varphi = (\varphi_t)_{t\in[0,1]}$ be an isotopy of S^1. Consider the universal cover $p \colon \mathbb{R} \to S^1$ given by $p(\theta) = e^{i\theta}$. Let $\widetilde{\varphi}$ be a lift of φ to $\mathbb{R}$. Then the *twist number* of φ is the number

$$\operatorname{twist}(\varphi) = \left\lceil \frac{1}{2\pi} \max\left\{ |\widetilde{\varphi}_1(\theta) - \widetilde{\varphi}_0(\theta)| : \theta \in \mathbb{R} \right\} \right\rceil . \qquad \times$$

It is clear that $\operatorname{twist}(\varphi)$ does not depend on the choice of the lift $\widetilde{\varphi}$.

Proposition A.11 (Elementary properties of twist number)**.** *Let $\varphi = (\varphi_t)_{t\in[0,1]}$ and $\psi = (\psi_t)_{t\in[0,1]}$ be isotopies of S^1, and let f and g be homeomorphisms $S^1 \to S^1$. Then the following hold:*

1. $\operatorname{twist}(\varphi \circ \psi) \leqslant \operatorname{twist}(\varphi) + \operatorname{twist}(\psi)$*;*

2. $\operatorname{twist}(f \circ \varphi \circ f^{-1}) = \operatorname{twist}(\varphi)$*;*

3. *if f and g are orientation-preserving, then they are isotopic through an isotopy ρ such that* $\operatorname{twist}(\rho) \leqslant 1$ *and, for every $x \in S^1$, either ρ fixes x or the map $t \mapsto \rho_t(x)$ is injective.*

Proof. We prove each statement in turn.

1. Let $\widetilde{\varphi}$ and $\widetilde{\psi}$ be lifts of φ and ψ to $\mathbb{R}$. By definition of $\operatorname{twist}(\psi)$, we have that

$$\widetilde{\psi}_1(\theta) \leqslant \widetilde{\psi}_0(\theta) + 2\pi \cdot \operatorname{twist}(\psi) \qquad \text{for every } \theta \in \mathbb{R}. \tag{A.1}$$

Suppose, for the sake of simplicity, that φ is orientation-preserving, so that $\widetilde{\varphi}_t$ is increasing for every $t \in [0,1]$. Then, for $\theta \in \mathbb{R}$, we have that

$$\begin{aligned} \widetilde{\varphi}_1(\widetilde{\psi}_1(\theta)) &\leqslant \widetilde{\varphi}_1(\widetilde{\psi}_0(\theta) + 2\pi \cdot \operatorname{twist}(\psi)) && \text{by (A.1), since } \widetilde{\varphi}_1 \text{ is increasing} \\ &= \widetilde{\varphi}_1(\widetilde{\psi}_0(\theta)) + 2\pi \cdot \operatorname{twist}(\psi) && \text{since } \widetilde{\varphi}_1(x+2\pi) = \widetilde{\varphi}_1(x) + 2\pi \\ &\leqslant \widetilde{\varphi}_0(\widetilde{\psi}_0(\theta)) + 2\pi \cdot \operatorname{twist}(\varphi) && \text{by definition of } \operatorname{twist}(\varphi), \\ &\quad + 2\pi \cdot \operatorname{twist}(\psi) \end{aligned}$$

and similarly

$$\widetilde{\varphi}_1(\widetilde{\psi}_1(\theta)) \geqslant \widetilde{\varphi}_0(\widetilde{\psi}_0(\theta)) - 2\pi(\operatorname{twist}(\varphi) + \operatorname{twist}(\psi)).$$

When φ is orientation-reversing, the argument is exactly the same (except for the fact that $\widetilde{\varphi}_t$ is decreasing and $\widetilde{\varphi}_1(x+2\pi) = \widetilde{\varphi}_1(x) - 2\pi$).

2. Let $\rho = (\rho_t)_{t \in [0,1]}$ be the isotopy of S^1 such that $\rho_t = f$ for every $t \in [0,1]$. It is clear from the definition that $\operatorname{twist}(\rho) = \operatorname{twist}(\rho^{-1}) = 0$. We can then apply statement 1 to get

$$\operatorname{twist}(f \circ \varphi \circ f^{-1}) = \operatorname{twist}(\rho \circ \varphi \circ \rho^{-1}) \leqslant \operatorname{twist}(\rho) + \operatorname{twist}(\varphi) + \operatorname{twist}(\rho^{-1}) = \operatorname{twist}(\varphi).$$

By symmetry, we also have that $\operatorname{twist}(f \circ \varphi \circ f^{-1}) \geqslant \operatorname{twist}(\varphi)$, so the conclusion follows.

3. Let $\widetilde{f}$ and $\widetilde{g}$ be lifts of f and g to $\mathbb{R}$ such that $0 \leqslant m < 2\pi$, where

$$m = \max\{\widetilde{f}(\theta) - \widetilde{g}(\theta) : \theta \in \mathbb{R}\}.$$

Let $\theta_0 \in [0, 2\pi]$ be a point attaining the maximum. For $\theta_0 \leqslant \theta < \theta_0 + 2\pi$ we have that

$$\begin{aligned}
\widetilde{f}(\theta) - \widetilde{g}(\theta) &\geqslant \widetilde{f}(\theta_0) - \widetilde{g}(\theta) && \text{since } \widetilde{f} \text{ is increasing} \\
&= \widetilde{g}(\theta_0) - \widetilde{g}(\theta_0) + m && \text{by definition of } \theta_0 \\
&> \widetilde{g}(\theta_0) - \widetilde{g}(\theta_0 + 2\pi) + m && \text{since } \widetilde{g} \text{ is increasing} \\
&= m - 2\pi && \text{since } \widetilde{g}(x + 2\pi) = \widetilde{g}(x) + 2\pi \\
&\geqslant -2\pi && \text{since } m \geqslant 0.
\end{aligned}$$

We conclude that

$$|\widetilde{f}(\theta) - \widetilde{g}(\theta)| < 2\pi \qquad \text{for every } \theta \in \mathbb{R}. \tag{A.2}$$

We now define an isotopy $\widetilde{\rho} = (\widetilde{\rho}_t)_{t\in[0,1]}$ of $\mathbb{R}$ as follows:

$$\widetilde{\rho}_t(\theta) = t \cdot \widetilde{g}(\theta) + (1 - t) \cdot \widetilde{f}(\theta).$$

One checks that this isotopy descends to an isotopy ρ of S^1 with $\rho_0 = f$ and $\rho_1 = g$. Inequality (A.2) immediately gives that $\operatorname{twist}(\rho) \leqslant 1$. Suppose now that $\widetilde{\rho}_t(\theta) = \widetilde{\rho}_{t'}(\theta) + 2k\pi$ for some $t, t' \in [0, 1]$, $\theta \in \mathbb{R}$, and $k \in \mathbb{Z}$. Some easy algebra shows that, in this case, we must have

$$(t' - t) \cdot (\widetilde{f}(\theta) - \widetilde{g}(\theta)) = 2k\pi.$$

Since $|t' - t| \leqslant 1$ and $|\widetilde{f}(\theta) - \widetilde{g}(\theta)| < 2\pi$, we either have $\widetilde{f}(\theta) = \widetilde{g}(\theta)$ (and, hence, ρ fixes $e^{i\theta}$) or $t = t'$. ▢

Definition A.12 (Twist number in a manifold)**.** Let X be a surface or a 3-manifold, and let Γ be a closed 1-manifold embedded in M. Let $\varphi = (\varphi_t)_{t\in[0,1]}$ be an isotopy of M preserving Γ. If Γ is connected, we define the *twist number* of φ relative to Γ as the number

$$\operatorname{twist}_\Gamma(\varphi) = \operatorname{twist}(f^{-1} \circ \varphi|_\Gamma \circ f),$$

where f is a homeomorphism $S^1 \to \Gamma$. If Γ is disconnected, we let

$$\operatorname{twist}_\Gamma(\varphi) = \sum_b \operatorname{twist}_b(\varphi),$$

where the sum ranges over the components b of Γ. ×

We remark that statement 2 of Proposition A.11 implies that the definition above does not depend on the choice of the homeomorphism f.

Proposition A.13 (Bounding twist number with intersection number)**.** *Let F be a compact connected orientable surface, and let a and c be arcs properly embedded in F that are not boundary-parallel. Let Ω be a boundary component of F that intersects both a and c. Let $\varphi = (\varphi_t)_{t\in[0,1]}$ be an isotopy of F such that φ_0 is the identity. If F is an annulus, suppose that φ fixes a component of ∂F. Then we have the bound*

$$\mathrm{twist}_\Omega(\varphi) \leqslant 2(|a\cap c| + |\varphi_1(a)\cap c| + 8).$$

Proof. Firstly, let us perform a small isotopy on c, if necessary, to ensure that a and c do not intersect on ∂F, and neither do $\varphi_1(a)$ and c; this can be done without increasing the number of intersections of a and c nor of $\varphi_1(a)$ and c. We can also assume that a and c intersect transversely, and so do $\varphi_1(a)$ and c. Let $\overline{F}$ be the double of F, that is the surface obtained by gluing two copies of F along their boundary via the identity map. We think of F as being embedded in $\overline{F}$ as one of these two copies, and we denote by $\iota\colon \overline{F}\to\overline{F}$ the involution that exchanges $\mathrm{int}(F)$ with $\overline{F}\setminus F$.

Homeomorphism around a boundary component. Pick a boundary component b of F, and let N be a regular neighbourhood of b in F. If we pick N small enough, we can identify it with $b\times[0,1]$ in such a way that $b\times\{0\}$ corresponds to b, and $a\cap N$ and $c\cap N$ are unions of arcs of the form $\{x\}\times[0,1]$. By composing with ι, this induces an identification of the annulus $A = N\cup\iota(N)$ with $b\times[-1,1]\subseteq\overline{F}$ such that $(a\cup\iota(a))\cap A$ and $(c\cup\iota(c))\cap A$ are unions of arcs of the form $\{x\}\times[-1,1]$. Apply statement 3 of Proposition A.11 with $f=\varphi_1|_b$ and $g=\mathrm{id}$ to get an isotopy ρ of b from $\varphi_1|_b$ to the identity.

We now define a homeomorphism $g_b\colon\overline{F}\to\overline{F}$ as follows. For $(x,y)\in b\times[-1,1]=A$, set

$$g_b(x,y) = \begin{cases} (\rho_{-y}(x), y) & \text{if } -1\leqslant y\leqslant 0,\\ (\varphi_{1-y}(x), y) & \text{if } 0\leqslant y\leqslant 1; \end{cases}$$

outside A, we set g_b to be the identity. See Figure A.1 for a sketch of the homeomorphism g_b. We remark a few properties of g_b.

1. The homeomorphism g_b of $\overline{F}$ preserves b and coincides with φ_1 on it.

2. For every component a' of $\iota(a)\cap(b\times[-1,0])$ and every component c' of $\iota(c)\cap(b\times[-1,0])$, the two arcs $g_b(a')$ and c' intersect in at most one point. This follows

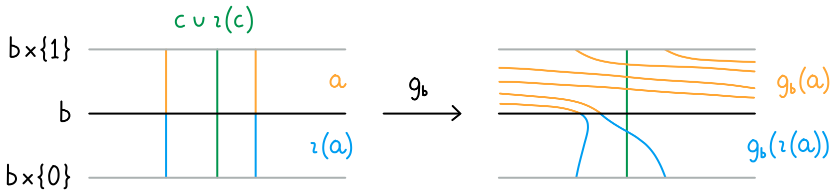

Figure A.1. The homeomorphism g_b around the boundary component b.

from the fact that ρ is obtained by applying statement 3 of Proposition A.11; in particular, this shows that $g_b(a')$ is either of the form $\{x\} \times [-1, 0]$, in which case it is disjoint from c', or $\{(\rho_t(x), -t) : t \in [0, 1]\}$, in which case it intersects c' in at most one point by injectivity of $t \mapsto \rho_t(x)$.

3. The restriction $g_b|_F$ is isotopic to φ_1 through an isotopy that is constant on ∂F. To show this explicitly, consider the isotopy ψ of F defined as follows: for $t \in [0, 1]$, set ψ_t to be the identity on $F \setminus (b \times [0, t])$, while for $(x, y) \in b \times [0, t]$ set

$$\psi_t(x, y) = (\varphi_{t-y}(x), y).$$

It is then easy to check that $(\varphi_t \circ \psi_t^{-1} \circ \psi_1)_{t\in[0,1]}$ is an isotopy of F from $\psi_1 = g_b|_F$ to φ_1 that is constantly equal to $\varphi_1|_{\partial F}$ on ∂F.

Note that, if φ fixes b, then we can (and will) take g_b to be the identity.

Writing g_b as a power of a Dehn twist. Since g_b restricts to a homeomorphism of the annulus A that fixes ∂A, it is isotopic (as a homeomorphism of $\overline{F}$) to a power $T_b^{k_b}$ of a Dehn twist about the core b of A. In order to estimate the exponent k_b, consider a lift $\widetilde{g} \colon \mathbb{R} \times [-1, 1] \to \mathbb{R} \times [-1, 1]$ of $g_b|_A$ to the universal cover $\mathbb{R} \times [1, 1]$ of $A = b \times [-1, 1]$. Since g_b preserves sets of the form $b \times \{t\}$ for $t \in [-1, 1]$, we see that $\widetilde{g}$ preserves the sets $\mathbb{R} \times \{t\}$; for each $t \in [-1, 1]$, denote by $\widetilde{g}_t \colon \mathbb{R} \to \mathbb{R}$ the homeomorphism such that

$$\widetilde{g}(x, t) = (\widetilde{g}_t(x), t).$$

Consider lifts $\widetilde{\varphi}$ and $\widetilde{\rho}$ of φ and ρ respectively to $\mathbb{R}$; in particular, pick them so that $\widetilde{\varphi}_1 = \widetilde{\rho_0} = \widetilde{g}_0$. It is easy to see that the exponent k_b such that g_b is isotopic to $T_b^{k_b}$ satisfies

$$2\pi|k_b| = |\widetilde{g}_1(\theta) - \widetilde{g}_{-1}(\theta)| \qquad \text{for every } \theta \in \mathbb{R}.$$

If we choose θ to maximise $|\widetilde{\varphi}_1(\theta) - \widetilde{\varphi}_0(\theta)|$, we have that

$$\begin{aligned}
|k_b| &= |\widetilde{g}_1(\theta) - \widetilde{g}_{-1}(\theta)|/2\pi \\
&= |\widetilde{\varphi}_1(\theta) - \widetilde{\rho}_0(\theta)|/2\pi && \text{by our choice of } \widetilde{\varphi} \text{ and } \widetilde{\rho} \\
&\geqslant |\widetilde{\varphi}_1(\theta) - \widetilde{\varphi}_0(\theta)|/2\pi - |\widetilde{\rho}_1(\theta) - \widetilde{\rho}_0(\theta)|/2\pi \\
&\geqslant \operatorname{twist}_b(\varphi) - \operatorname{twist}(\rho) && \text{by definition of twist number} \\
& && \text{and our choice of } \theta \\
&\geqslant \operatorname{twist}_b(\varphi) - 1 && \text{since } \operatorname{twist}(\rho) \leqslant 1.
\end{aligned}$$

We remark that, when φ fixes b, the exponent k is zero.

Patching homeomorphisms around all boundary components. For each boundary component b of F, we have constructed a homeomorphism g_b of $\overline{F}$ that is supported in a regular neighbourhood of b in $\overline{F}$; in fact, it is not restrictive to assume that these supports are pairwise disjoint. We can then define $g = g_{b_1} \circ \ldots g_{b_m}$, where $b_1, \ldots, b_m$ are the components of ∂F. This is a homeomorphism of $\overline{F}$ that is supported in a regular neighbourhood A of ∂F in $\overline{F}$. Recall that, in the case where F is an annulus, we are assuming that φ fixes a component of ∂F; therefore, so does g.

From the properties of the individual homeomorphisms g_b, we can infer properties of g.

4. The homeomorphism g preserves ∂F, and coincides with φ_1 on it; this is an immediate consequence of property 1 above.

5. The homeomorphism g is isotopic to

$$f = \prod_{b \subseteq \partial F} T_b^{k_b}$$

in $\overline{F}$, where the product ranges over the components b of ∂F. Recall that k_b is an integer such that $|k_b| \geqslant \operatorname{twist}_b(\varphi) - 1$, and that k_b is zero if φ fixes b.

6. The homeomorphism $g|_F$ of F is isotopic to φ_1 through an isotopy that is constant on ∂F. This is an immediate consequence of property 3 above.

7. If we think of $g(\iota(a))$ and $\iota(c)$ as arcs properly embedded in $\iota(F)$, then

$$|g(\iota(a)) \cap \iota(c)| \leqslant |a \cap c| + 4.$$

In fact, by property 2 above, the arcs $g(\iota(a))$ and $\iota(c)$ intersect in at most 4 points in A; moreover, since g is the identity on $\iota(F)\setminus A$, the arcs $g(\iota(a))$ and $\iota(c)$ intersect at most $|a \cap c|$ times outside A.

Bounding twist number. Let $\overline{a}$ and $\overline{c}$ respectively denote the curves $a \cup \iota(a)$ and $c \cup \iota(c)$ in $\overline{F}$. Their intersection number is bounded by

$$i(\overline{a}, \overline{c}) \leqslant 2|a \cap c|. \tag{A.3}$$

Let $\overline{a}' = g(\overline{a})$. By properties 6 and 7 above, the curve $\overline{a}'$ can be isotoped in $\overline{F}$ to a curve $\overline{a}''$ that coincides with $\varphi_1(a)$ on F, and intersects $\iota(c)$ in at most $|a \cap c| + 4$ points. In particular, we have that

$$i(\overline{a}', \overline{c}) \leqslant |\overline{a}'' \cap \overline{c}| \leqslant |a \cap c| + |\varphi_1(a) \cap c| + 4. \tag{A.4}$$

By property 5 above, the curve $\overline{a}'$ is also isotopic to $f(\overline{a})$.

It is reasonable to expect that, if the exponents k_b are large, then the intersection number $i(f(\overline{a}), \overline{c})$ should also be large. This statement is made precise, for instance, in [15, Lemma 4.2]. In our setting, this lemma reads

$$i(\overline{a}, \overline{c}) + i(f(\overline{a}), \overline{c}) \geqslant \sum_{\substack{b \subseteq \partial F \\ k_b \neq 0}} (|k_b| - 2) \cdot i(\overline{a}, b) \cdot i(\overline{c}, b), \tag{A.5}$$

where the sum ranges over the components b of ∂F with $k_b \neq 0$. Two remarks are in order. Firstly, [15, Lemma 4.2] requires that the curves we Dehn twist about – namely, components of ∂F – are not pairwise isotopic. This is always the case, except when F is an annulus; however, in this situation, recall that we are assuming that φ fixes a component of ∂F; this implies that there is at most one non-zero exponent k_b, and therefore the lemma still applies. The second remark is that, even though [15, Lemma 4.2] as stated requires the exponents to be "natural numbers", the proof therein explicitly asserts that this assumption is not required.

Note that, by the assumption that a and c are not boundary-parallel, we have that $i(\overline{a}, b) = |a \cap b|$ and $i(\overline{c}, b) = |c \cap b|$ for every component b of ∂F. Finally, we get the bound

$$\begin{aligned} \operatorname{twist}_\Omega(\varphi) &\leqslant |k_\Omega| + 1 && \text{since } |k_\Omega| \geqslant \operatorname{twist}_\Omega(\varphi) - 1 \\ &\leqslant (|k_\Omega| + 1) \cdot |a \cap \Omega| \cdot |c \cap \Omega| && \text{since } a \text{ and } c \text{ intersect } \Omega \end{aligned}$$

$$
\begin{aligned}
&\leqslant \sum_{b\subseteq\partial F} (|k_b|+1)\cdot|a\cap b|\cdot|c\cap b| \\
&\leqslant \sum_{b\subseteq\partial F} (|k_b|-2)\cdot i(\overline{a},b)\cdot i(\overline{c},b)+12 && \text{since } i(\overline{a},b)\cdot i(\overline{c},b)\leqslant 4\\
&\leqslant i(\overline{a},\overline{c})+i(f(\overline{a}),\overline{c})+12 && \text{by (A.5)}\\
&\leqslant 3|a\cap c|+|\varphi_1(a)\cap c|+16 && \text{by (A.3) and (A.4)},
\end{aligned}
$$

where the sums range over the components b of ∂F. We conclude by noting that, when $|\varphi_1(a)\cap c|<|a\cap c|$, we obtain the desired inequality by replacing a with $\varphi_1(a)$ and φ with $(\varphi_{1-t})_{t\in[0,1]}$. ▢

Proposition A.14 (Bounding weight with twist number)**.** *Let $\mathcal{T}$ be a triangulation of a compact orientable 3-manifold with t tetrahedra, and let Γ be a closed simplicial sub-1-manifold of ∂M. Let A be a general position surface properly embedded in M, and let $\varphi=(\varphi_t)_{t\in[0,1]}$ be an isotopy of Γ such that φ_0 is the identity. Suppose that $\varphi_1(A\cap\Gamma)$ does not contain any vertices of $\mathcal{T}$. Then there is an isotopy $\psi=(\psi_t)_{t\in[0,1]}$ of M preserving Γ such that:*

(i) *ψ_0 is the identity;*

(ii) *ψ coincides with φ on Γ;*

(iii) *there is an arbitrarily small neighbourhood of Γ in M outside which ψ_t is the identity for every $t\in[0,1]$;*

(iv) *$\psi_1(A)$ is a general position surface whose weight is bounded by*

$$
w(\psi_1(A))\leqslant w(A)+12t\cdot\sum_{b\subseteq\Gamma}|A\cap b|\cdot\mathrm{twist}_b(\varphi),
$$

where the sum ranges over the components b of Γ.

Proof. We work with one component of Γ at a time.

Constructing a regular neighbourhood. Let b be a component of Γ, and let N be an arbitrarily small regular neighbourhood of b in M. We can identify N with $S^1\times U$, where

$$
\begin{aligned}
S^1 &= \{e^{i\theta}:\theta\in\mathbb{R}\} \text{ and}\\
U &= \{(y_1,y_2)\in\mathbb{R}^2 : y_2\geqslant 0, y_1^2+y_2^2\leqslant 1\},
\end{aligned}
$$

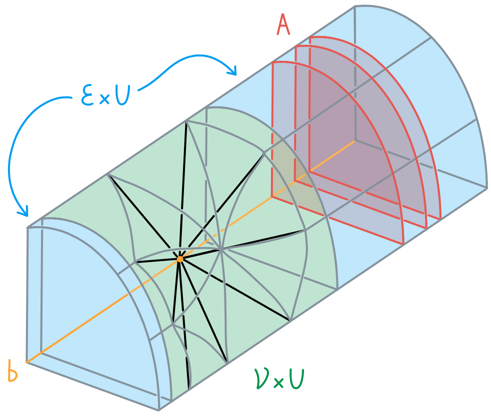

Figure A.2. A portion of the regular neighbourhood N of b. The regular neighbourhood N is decomposed into stars of vertices (labelled $\mathcal{V}\times U$) and neighbourhoods of truncated edges (labelled $\mathcal{E}\times U$).

so that b corresponds to $S^1 \times \{(0,0)\}$.

We want N to agree with the combinatorial structure of the triangulation $\mathcal{T}$ and the general position surface A therein. Informally, we would like N to be partitioned into *stars of vertices* (regions that contain a vertex of $\mathcal{T}$, intersect $\mathcal{T}^{(1)}$ in "straight lines", and do not intersect A) and *neighbourhoods of truncated edges* (regions that only intersect $\mathcal{T}^{(1)}$ in the interior of an edge on ∂F and intersect A in parallel copies of U). This decomposition is shown in Figure A.2.

Formally, for a point $e^{i\theta} \in S^1$, we call a *straight line* in N based at $e^{i\theta}$ a subset of N of the form

$$\left\{\left(e^{i\theta+i\lambda\cdot t}, t\cdot y\right) : t \in [0,1]\right\}$$

for some $\lambda \in \mathbb{R}$ and $y \in U$ with $\|y\| = 1$. If we pick N to be small enough, we can find an identification $N = S^1 \times U$ such that:

- there is a decomposition $S^1 = \mathcal{V} \cup \mathcal{E}$, where $\mathcal{V}$ and $\mathcal{E}$ are non-empty finite unions of closed connected subsets of S^1 such that $\mathcal{V} \cap \mathcal{E} = \partial\mathcal{V} = \partial\mathcal{E}$;
- every component of $\mathcal{V} \times U$ contains exactly one vertex of $\mathcal{T}$;
- for each edge e of $\mathcal{T}$ not contained in b, each component of $e \cap (\mathcal{V} \times U)$ is a straight line based at a vertex of $\mathcal{T}$;
- $\mathcal{V} \times U$ does not intersect A;

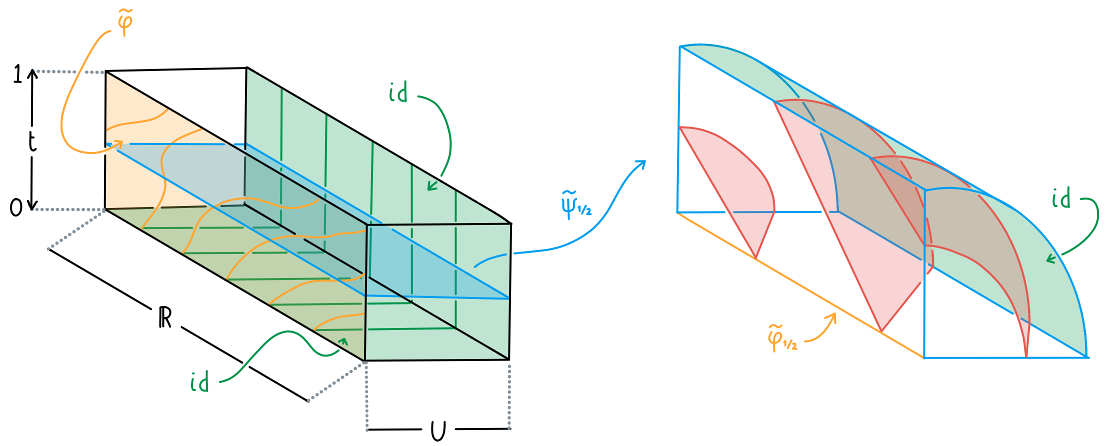

Figure A.3. On the left, a schematic representation of the isotopy $\widetilde{\psi}$ of $\mathbb{R} \times U$. On the right, emphasis on the homeomorphism $\widetilde{\psi}_{1.2}$ of $\mathbb{R} \times U$; we have drawn some surfaces in $\mathbb{R} \times U$ that are images under $\widetilde{\psi}_{1/2}$ of discs of the form $\{\theta\} \times U$, to help visualise the effect of $\widetilde{\psi}_{1/2}$.

- $(\mathcal{E} \times U) \cap \mathcal{T}^{(1)} = (\mathcal{E} \times U) \cap b$;
- every component of $(\mathcal{E} \times U) \cap A$ is of the form $\{x\} \times U$ for some $x \in \mathcal{E}$.

The reader should think of $\mathcal{V} \times U$ as stars of vertices, and $\mathcal{E} \times U$ as neighbourhoods of truncated edges; again, see Figure A.2.

An isotopy of N. Consider the universal cover $\mathbb{R} \to b$, where we identify b and S^1, and the induced universal cover $\mathbb{R} \times U \to N$. Let $\widetilde{\varphi}$ be the lift of φ to $\mathbb{R}$ such that $\widetilde{\varphi}_0$ is the identity. We define an isotopy $\widetilde{\psi} = (\widetilde{\psi}_t)_{t\in[0,1]}$ of $\mathbb{R} \times U$ as follows:

$$\widetilde{\psi}_t(\theta, y) = \Big(\|y\| \cdot \theta + (1 - \|y\|) \cdot \widetilde{\varphi}_t(\theta), y\Big).$$

A schematic representation of $\widetilde{\psi}$ is shown in Figure A.3. One easily checks that $\widetilde{\psi}$ descends to an isotopy ψ of N that fixes $\operatorname{clos}(\partial N \setminus \partial M)$, coincides with φ on b, and such that ψ_0 is the identity.

We now analyse the effect of ψ_1 on a component $D = \{e^{i\theta}\} \times U$ of $A \cap N$. Let $\widetilde{D} = \{\theta\} \times U$ be a lift of D to $\mathbb{R} \times U$. Then

$$\widetilde{\psi}_1(\widetilde{D}) = \Big\{\Big(\|y\| \cdot \theta + (1 - \|y\|) \cdot \widetilde{\varphi}_1(\theta), y\Big) : y \in U\Big\}, \tag{A.6}$$

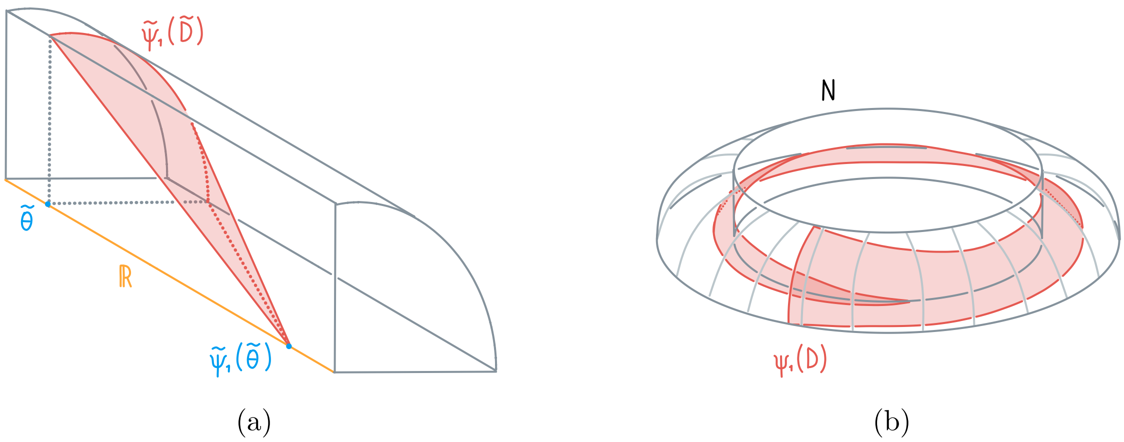

Figure A.4. (a) The image under $\widetilde{\psi}_1$ of a lift $\widetilde{D} = \{\theta\} \times U$ of a component $D = \{e^{i\theta}\} \times U$ of $A \cap N$. (b) The disc $\psi_1(D)$ in N.

as depicted in Figure A.4a.

Let s be a component of the intersection between an edge of $\mathcal{T}$ not contained in b and $\mathcal{V} \times U$, which by construction of N is a straight line in N based at a vertex of $\mathcal{T}$. Let V be the component of $\mathcal{V}$ containing the vertex $s \cap b$. Since $e^{i\theta} \notin V$, it is easy to see that $\widetilde{\psi}_1(\widetilde{D})$ intersects at most

$$\left\lceil \frac{|\widetilde{\varphi}_1(\theta) - \theta|}{2\pi} \right\rceil \leqslant \operatorname{twist}_b(\varphi)$$

lifts of V.

Now, every lift $\widetilde{s}$ of s to $\mathbb{R} \times U$ is contained in $\widetilde{V} \times U$ for some lift $\widetilde{V}$ of V, and is of the form

$$\widetilde{s} = \{(\eta + \lambda \cdot t, t \cdot y) : t \in [0,1]\} \tag{A.7}$$

for some $\lambda \in \mathbb{R}$, $y \in U$ with $\|y\| = 1$, and $\eta \in \mathbb{R}$ such that $e^{i\eta}$ is the vertex $s \cap b$ of $\mathcal{T}$. Some easy algebra shows that $\widetilde{\psi}_1(\widetilde{D})$ intersects $\widetilde{s}$ in at most one point. The only subtle point is showing that $\widetilde{\varphi}_1(\theta) \neq \eta$. This is true because $e^{i\eta}$ is a vertex of $\mathcal{T}$, while $\exp(i\widetilde{\varphi}_1(\theta))$ is a point of $\varphi_1(A \cap \Gamma)$, which by assumption does not contain any vertices of $\mathcal{T}$. Moreover, when $\widetilde{\psi}_1(\widetilde{D})$ and $\widetilde{s}$ do intersect, they do so transversely. This can again be easily checked from (A.6) and (A.7).

We conclude that $\psi_1(D)$ intersects s at most $\operatorname{twist}_b(\varphi)$ times; see Figure A.4b for a representation of the disc $\psi_1(D)$ in N. There are at most $6t$ edges in $\mathcal{T}$, and every edge not contained in b intersects $\mathcal{V} \times U$ in at most two straight lines, contributing at most $2\operatorname{twist}_b(\varphi)$ intersections. Moreover, the disc $\psi_1(D)$ also intersects b once.

There are $|A \cap \Gamma|$ components of $A \cap N$, so we get a total of

$$|\psi_1(A \cap N) \cap \mathcal{T}^{(1)}| \leqslant |A \cap b| \cdot (12t \cdot \mathrm{twist}_b(\varphi) + 1) \tag{A.8}$$

intersections.

Patching isotopies around all components of Γ. We now apply the same procedure to each component of Γ, obtaining an isotopy ψ_b of an arbitrarily small neighbourhood N_b of b in M for each component b of Γ. Since each ψ_b fixes $\mathrm{clos}(\partial N_b \setminus \partial M)$, we can define an isotopy ψ of M that coincides with ψ_b on each N_b, and is the identity elsewhere. Let N be the union of all the N_b. We see that $\psi_1(A)$ is in general position with respect to $\mathcal{T}$, since $\psi_1(A) \cap N$ and $\psi_1(A) \setminus N$ are. Finally, we get the desired bound:

$$\begin{aligned}
w(\psi_1(A)) &= |\psi_1(A \cap N) \cap \mathcal{T}^{(1)}| + |\psi_1(A \setminus N) \cap \mathcal{T}^{(1)}| \\
&= |\psi_1(A \cap N) \cap \mathcal{T}^{(1)}| + w(A) - |A \cap \Gamma| && \text{since } \psi_1 \text{ is the identity outside } N \\
&= \sum_{b \subseteq \Gamma} |\psi_1(A \cap N_b) \cap \mathcal{T}^{(1)}| + w(A) - |A \cap \Gamma| \\
&\leqslant \sum_{b \subseteq \Gamma} |A \cap b| \cdot (12t \cdot \mathrm{twist}_b(\varphi) + 1) && \text{by (A.8)} \\
&\quad + w(A) - |A \cap \Gamma| \\
&= \sum_{b \subseteq \Gamma} |A \cap b| \cdot 12t \cdot \mathrm{twist}_b(\varphi) + w(A),
\end{aligned}$$

where the sums range over the components b of Γ. ▢

Corollary A.15 (Bounding weight up to isotopy with weight on the boundary)**.** *Let $\mathcal{T}$ be a triangulation of a compact orientable 3-manifold M with t tetrahedra, and let F be a simplicial subsurface of ∂M. Let A be a general position surface properly embedded in M, and denote by a the intersection $A \cap F$, which is a general position 1-manifold in F. Suppose that no component of a is a boundary-parallel arc in F. Let $\varphi = (\varphi_t)_{t \in [0,1]}$ be an isotopy of M preserving F such that φ_0 is the identity. Suppose that $A' = \varphi_1(A)$ is in general position with respect to $\mathcal{T}$. Additionally, suppose that φ fixes a boundary component of every annulus component of F. Then there is an isotopy $\psi = (\psi_t)_{t \in [0,1]}$ of M preserving F such that:*

(i) ψ_0 is the identity;

(ii) ψ fixes ∂F;

(iii) there is an arbitrarily small neighbourhood of ∂F in M outside which ψ and φ coincide;

(iv) $\psi_1(A)$ is a general position surface whose weight is bounded by

$$w(\psi_1(A)) \leqslant 193t^2 \cdot (w(A') + w(a)).$$

Proof. Let ψ' be the isotopy of M given by Proposition A.14 applied to the surface A', the isotopy $(\varphi_t^{-1})_{t\in[0,1]}$, and $\Gamma = \partial F$. Set $\psi_t = \psi'_t \circ \varphi_t$ for $t \in [0,1]$. By construction, we see that ψ_0 is the identity, ψ fixes ∂F, and ψ coincides with φ outside an arbitrarily small neighbourhood of ∂F. Moreover, $\psi_1(A) = \psi'_1(A')$ is a general position surface whose weight is bounded by

$$w(\psi_1(A)) \leqslant w(A') + 12t \cdot \sum_{b\subseteq\Gamma} |a \cap b| \cdot \mathrm{twist}_b(\varphi). \tag{A.9}$$

Consider a component b of Γ that intersects a. Let a_b be a component of a that intersects b and, amongst all such components, minimises $w(a_b) + w(\varphi_1(a_b))$. In particular, this choice of a_b guarantees that

$$|a \cap b| \cdot (w(a_b) + w(\varphi_1(a_b))) \leqslant 2w(a) + 2w(\varphi_1(a)). \tag{A.10}$$

Let G be the component of F containing b. Note that G cannot be a disc, since by assumption a_b is not boundary-parallel in G. If G has at least two boundary components, let c be a simplicial arc properly embedded in G that has precisely one endpoint on b. Otherwise, take a simplicial path c' in G that starts and ends on b, is non-trivial in the relative homology group $H_1(G, b; \mathbb{Z}/2\mathbb{Z})$, and has minimal length amongst all such paths. It is easy to see that c' is transverse to b, it has no triple self-intersections, and it runs along each edge of G at most once. By resolving the self-intersections of c', we can construct an arc c properly embedded in G that is not boundary-parallel. Either way, we see that the arc c satisfies

$$|a_b \cap c| \leqslant w(a_b) \quad \text{and} \quad |\varphi_1(a_b) \cap c| \leqslant w(\varphi_1(a_b)),$$

since c is essentially a union of edges of $\mathcal{T}$. We then apply Proposition A.13 to get the bound

$$\mathrm{twist}_b(\varphi) \leqslant 2(w(a_b) + w(\varphi_1(a_b)) + 8). \tag{A.11}$$

Finally, we can combine (A.9) to (A.11) to obtain

$$
\begin{aligned}
w(\psi_1(A)) &\leqslant w(A') + 12t \cdot \sum_{b \subseteq \Gamma} |a \cap b| \cdot \mathrm{twist}_b(\varphi) \\
&\leqslant w(A') + 12t \cdot \sum_{b \subseteq \Gamma} |a \cap b| \cdot 2(w(a_b) + w(\varphi_1(a_b)) + 8) \\
&\leqslant w(A') + 192t \cdot |a \cap \Gamma| + 48t \cdot \sum_{b \subseteq \Gamma} (w(a) + w(\varphi_1(a))) \\
&\leqslant w(A') + 192t \cdot |a \cap \Gamma| + 96t^2 \cdot (w(a) + w(\varphi_1(a))),
\end{aligned}
$$

where the sums range over the components b of Γ that intersect a; the last inequality follows from the (obvious) fact that Γ has at most $2t$ components. We obtain the inequality in the statement by recalling that $w(\varphi_1(a)) \leqslant w(A')$ and that

$$|a \cap \Gamma| \leqslant \min\{w(a), w(A')\},$$

together with some elementary algebra. ▢

Proposition A.16 (Weights of isotopic vertical surfaces in interval bundles)**.** *Let $\mathcal{T}$ be a suitable triangulation of $M = F \times [0,1]$ with t tetrahedra, where F is a compact orientable surface, and let $\Gamma = \partial\partial_h M$. Let A be a general position incompressible boundary-incompressible vertical surface properly embedded in (M, Γ), and denote by a the intersection $A \cap (\partial_0 M \cup \partial_v M)$, which is a general position 1-manifold in $\partial_0 M \cup \partial_v M$. Suppose that A is isotopic preserving Γ to a general position surface A'. Then A is isotopic fixing Γ to a general position surface A'' whose weight is bounded by*

$$w(A'') \leqslant 74\,884 t^5 \cdot (w(A') + w(a)).$$

Proof. It is not restrictive to assume that F (and hence M) is connected. If F is a sphere or a disc, then the conclusion holds vacuously, since A is empty. Otherwise, let $\Gamma_i = \partial\partial_i M$ for $i \in \{0,1\}$. Let $\varphi = (\varphi_t)_{t\in[0,1]}$ be an isotopy of M preserving Γ such that φ_0 is the identity and $\varphi_1(A) = A'$.

We first deal with the case where F is not an annulus. By applying Corollary A.15 to the simplicial subsurface $\partial_0 M \subseteq \partial M$, the general position surface $A \subseteq M$, and the isotopy φ, we see that A is isotopic fixing Γ_0 to a general position surface A_1 whose weight is bounded by

$$w(A_1) \leqslant 193t^2 \cdot (w(A') + w(a)) \tag{A.12}$$

We now apply Corollary A.15 to the simplicial subsurface $\partial_v M \subseteq \partial M$, the general position surface $A \subseteq M$, and the isotopy from A to A_1 fixing Γ_0; note that $\partial_v M$ is a union of annuli, but this isotopy fixes one of the boundary components of each annulus. Therefore, we see that A is isotopic fixing Γ to a general position surface A'' whose weight is bounded by

$$\begin{aligned} w(A'') &\leqslant 193t^2 \cdot (w(A_1) + w(a)) \\ &\leqslant 193t^2 \cdot \Big(193t^2 \cdot (w(A') + w(a)) + w(a)\Big) && \text{by (A.12)} \\ &\leqslant 37\,442t^4 \cdot (w(A') + w(a)) && \text{by elementary algebra.} \end{aligned}$$

When F is an annulus, we need to modify the isotopy φ before we can apply Corollary A.15. Write $M = S^1 \times [0,1] \times [0,1]$, and let $\psi = (\psi_t)_{t\in[0,1]}$ be the isotopy of S^1 such that

$$\varphi_t(x,0,0) = (\psi_t(x),0,0) \qquad \text{for every } t \in [0,1],\ x \in S^1.$$

By statement 3 of Proposition A.11, there is an isotopy $\rho = (\rho_t)_{t\in[0,1]}$ of S^1 with $\operatorname{twist}(\rho) \leqslant 1$ such that ρ_0 is the identity and $\rho_1 = \psi_1^{-1}$.

Extend now ψ and ρ to isotopies of M in the natural way (that is, by fixing the two $[0,1]$ factors). By applying Proposition A.14 to the boundary pattern $\Gamma_{00} = S^1 \times \{0\} \times \{0\}$, the general position surface A', and the isotopy ρ, we see that A' is isotopic to a general position surface A'_1 whose weight is bounded by

$$w(A'_1) \leqslant w(A') + 12t \cdot |a \cap \Gamma_{00}|, \tag{A.13}$$

through an isotopy that preserves Γ and agrees with ρ on Γ_{00}. In particular, A'_1 and $\rho_1(A')$ are isotopic fixing Γ_{00}.

Note that $\rho_1(A')$ is isotopic to A fixing Γ_{00}, through the isotopy $\psi^{-1} \circ \varphi$. We deduce that A is isotopic to A'_1 through an isotopy φ' fixing Γ_{00}. We are now in a position to apply the argument we described above for the non-annular case, replacing A' with A'_1 and φ with φ'. Because of this additional step, the final bound on the weight of A'' in the annular case is

$$\begin{aligned} w(A'') &\leqslant 37\,442t^4 \cdot (w(A'_1) + w(a)) \\ &\leqslant 37\,442t^4 \cdot (w(A') + w(a) + 12t \cdot |a \cap \Gamma_{00}|) && \text{by (A.13)} \\ &\leqslant 37\,442t^4 \cdot \left(\frac{3}{2}t + 1\right)(w(A') + w(a)) && \text{since } 4|a \cap \Gamma_{00}| = |a \cap \Gamma| = |A' \cap \Gamma| \end{aligned}$$

$$\leqslant 74\,884 t^5 \cdot (w(A') + w(a)) \qquad \text{by elementary algebra.} \qquad \square$$

A.3 Vertical surfaces with prescribed boundary

In Section 2.2.5, we defined a fundamental normal surface to be a normal surface that cannot be written as a normal sum of two non-empty normal surfaces. From the point of view of normal vectors, this means that the normal vector of a fundamental surface is not the sum of two non-zero and non-negative solutions of the matching and consistency equations. This notion can be generalised to arbitrary systems of equations and inequalities.

Definition A.17 (Fundamental and vertex solution)**.** Let $\mathcal{S}$ be a system of equations and inequalities in $\mathbb{R}^n$. A solution of $\mathcal{S}$ is *integral* if all its entries are integers. An integral solution of $\mathcal{S}$ is *fundamental* if it cannot be written as a sum of two integral solutions of $\mathcal{S}$, both of which are non-zero. An integral solution $\mathbf{x}$ of $\mathcal{S}$ is a *vertex solution* if, whenever $\alpha \cdot \mathbf{x} = \mathbf{y} + \mathbf{z}$, with $\mathbf{y}$ and $\mathbf{z}$ both solutions of $\mathcal{S}$ and α is a positive real number, then both $\mathbf{y}$ and $\mathbf{z}$ are scalar multiples of $\mathbf{x}$. ×

Proposition A.18 (Bounding fundamental solutions with ℓ^2 norms of the columns)**.** *Let* $\mathbf{A}$ *be an* $m \times n$ *matrix with integer entries. Consider the system*

$$\begin{cases} \mathbf{A} \cdot \mathbf{x} = \mathbf{0}, & \\ x_i = 0 & \text{for } i \in \mathcal{I}, \\ x_i \geqslant 0 & \text{for } 1 \leqslant i \leqslant n, \end{cases} \tag{A.14}$$

where $\mathbf{x} \in \mathbb{R}^n$ *and* $\mathcal{I}$ *is a subset of* $\{1, \ldots, n\}$*. For each* $1 \leqslant i \leqslant n$*, let* $k_i \geqslant 1$ *be an upper bound to the* ℓ^2 *norm of the* i*-th column of* $\mathbf{A}$*. Then every fundamental solution of the system (A.14) has* ℓ^∞ *norm bounded above by*

$$n \cdot (k_1^2 + 1)^{1/2} \cdots (k_n^2 + 1)^{1/2}.$$

Proof. The argument consists entirely of elementary linear algebra, and is outlined, for instance, in the proof of [22, Theorem 8.1]. There, it is shown that the entries of any vertex solution of (A.14) are bounded above by the absolute value of the determinant of some $(n-1) \times (n-1)$ minor of an $n \times n$ matrix $\widehat{\mathbf{A}}$; each row of $\widehat{\mathbf{A}}$ is either

- a row of $\mathbf{A}$,

- a row where one entry is 1 and the others are 0, or
- a row where all the entries are ones;

the matrix $\widehat{\mathbf{A}}$ has rank n, so there are no repeated rows. It is then immediate to see that the absolute value of the determinant of an $(n-1) \times (n-1)$ minor of $\widehat{\mathbf{A}}$ is bounded above by

$$(k_1^2 + 1)^{1/2} \cdots (k_n^2 + 1)^{1/2}.$$

In the proof of [22, Theorem 8.1], it is also shown that the entries of any fundamental solution of (A.14) are bounded above by n times the maximum entry of a vertex solution. This concludes the proof. ▢

We are finally ready to fulfil the purpose of this appendix, that is, provide a proof of Proposition 4.10; we recall the statement here for convenience.

Proposition 4.10 (Vertical surfaces with prescribed boundary). *Let $\mathcal{T}$ be a suitable pre-sutured triangulation of $M = F \times [0,1]$ with t tetrahedra, where F is a compact orientable surface. Let a be a normal 1-manifold in $\partial_0 M \cup \partial_v M$. Denote by a_0 and a_v the intersections of a with $\partial_0 M$ and $\partial_v M$ respectively. Suppose that no component of a_0, seen as a 1-manifold in $\partial_0 M$, is a boundary-parallel arc or a curve bounding a disc; moreover, suppose that a_v is vertical in $\partial_v M$. Then there is a normal surface A in M such that:*

(i) A is a vertical surface in M;

(ii) $A \cap (\partial_0 M \cup \partial_v M)$ is isotopic to a in $\partial_0 M \cup \partial_v M$, fixing $\partial\partial_h M$;

(iii) the weight of A is bounded by

$$w(A) \leqslant 2^{9t+40} \cdot w(a).$$

Proof. It is not restrictive to assume that M is connected. Note that if a is empty then the conclusion follows vacuously. Therefore, we can assume that a is non-empty.

Finding a normal surface. Let $\Gamma = \partial\partial_h M$. Note that F cannot be a sphere or a disc, since a is non-empty. Therefore, the 3-manifold with boundary pattern (M, Γ) is irreducible and boundary-irreducible. There exists a surface A_1 properly embedded in M such that A_1 is vertical in M and $A_1 \cap (\partial_0 M \cup \partial_v M) = a$. To see this, fix an identification of M with $F \times [0,1]$ such that $\partial_v M = \partial F \times [0,1]$ and a_v is a union of interval fibres. Then we can simply take $A_1 = a_0 \times [0,1] \subseteq F \times [0,1]$. Up to admissible isotopy, we can also assume that A_1 is in general position with respect to $\mathcal{T}$.

Let $\mathcal{E}$ be the set of edges of $\mathcal{T}$ that are contained in $\partial_0 M \cup \partial_v M$. Note that, because of the assumptions on a_0, the surface A_1 is incompressible and boundary-incompressible in (M, Γ) by Proposition 3.4. Moreover, no component of A_1 is a sphere or a disc intersecting Γ at most twice. By Remark A.3, there exists a least-$w_{\mathcal{E},*}$ normal surface A_2 in M that is admissibly isotopic to A_1. Let $a_2 = A_2 \cap (\partial_0 M \cup \partial_v M)$; this is a normal 1-manifold in $\partial_0 M \cup \partial_v M$. In particular, we have that

$$w(a_2) = w_{\mathcal{E}}(A_2) \leqslant w_{\mathcal{E}}(A_1) = w(a).$$

We remark that A_2 does not necessarily satisfy condition (ii); in fact, the 1-manifold a_2 is isotopic to a in $\partial_0 M \cup \partial_v M$ preserving Γ, but not necessarily fixing Γ pointwise. We will address this issue at the end of the proof, by means of Proposition A.16.

Setting up linear equations. We now set up a system of linear equations in the variables $x_1, \ldots, x_{7t}$, and λ. The variables $x_1, \ldots, x_{7t}$ correspond to the $7t$ types of normal discs defined by the triangulation $\mathcal{T}$, while λ is an auxiliary variable to make the system homogeneous. We consider three types of equations.

- *Matching equations*: these are the standard matching equations for normal surfaces, as detailed in Section 2.2.5; each equation is of the form

$$x_i + x_j = x_k + x_l.$$

- *Consistency equations*: these are equations ensuring that no tetrahedron contains more than one type of normal quadrilateral. In fact, we impose even stronger constrains, requiring that

$$x_i = 0$$

  for every $1 \leqslant i \leqslant 7t$ such that $(\mathbf{v}_{A_2})_i = 0$.

- *Boundary equations*: these are equations ensuring that a normal surface satisfying them intersects $\partial_0 M \cup \partial_v M$ in a normal 1-manifold that is a multiple of a_2. For every triangle T of $\mathcal{T}$ contained in $\partial_0 M \cup \partial_v M$, and for every type q of normal arc in T, let y_q be the number of normal arcs of type q in a_2. Then we require that

$$x_i + x_j = y_q \cdot \lambda,$$

where x_i and x_j are the variables corresponding to the two types of normal discs that intersect T in an arc of type q.

Consider then the system

$$\begin{cases} \text{matching equations,} \\ \text{consistency equations,} \\ \text{boundary equations,} \\ x_i \geqslant 0 \text{ for every } 1 \leqslant i \leqslant 7t, \\ \lambda \geqslant 0. \end{cases} \tag{A.15}$$

We say that a normal surface S satisfies (A.15) if there exists an integer λ such that $(\mathbf{v}_S, \lambda)$ is a solution of the system. Note that such a λ, if it exists, is necessarily unique; we denote it by $\lambda(S)$.

Bounding fundamental solutions. We can write the system (A.15) in the form described in Proposition A.18, where the matrix $\mathbf{A}$ comes from the matching equations and the boundary equations. Every type of normal disc appears in at most 4 matching or boundary equations, so the ℓ^2 norm of the columns corresponding to the variables x_i is at most 2. A bound for the ℓ^2 norm of the column corresponding to λ is given by

$$\sqrt{\sum_q y_q^2} \leqslant \sum_q y_q = w(a_2) - \frac{1}{2}|\partial a_2| \leqslant w(a_2) \leqslant w(a),$$

where the sums range over all the types q of normal arcs of a.

In conclusion, if $(x_1, \ldots, x_{7t}, \lambda)$ is a fundamental solution of (A.15) and S is the corresponding normal surface – that is, the normal surface such that $\mathbf{v}_S = (x_1, \ldots, x_{7t})$ – then we have the following bound on the weight of S:

$w(S) \leqslant 4(x_1 + \cdots + x_{7t})$	because every normal disc contributes to at most 4 points of $S \cap \mathcal{T}^{(1)}$
$\leqslant 4 \cdot 5t \cdot \max\{x_1, \ldots, x_{7t}\}$	because at most $5t$ coordinates of $\mathbf{v}_S$ are non-zero

$$\leqslant 20t \cdot (7t+1) \cdot 5^{7t/2} \cdot (w(a)^2+1)^{1/2} \quad \text{by Proposition A.18}$$
$$\leqslant 300t^2 \cdot 5^{7t/2} \cdot w(a) \quad \text{by elementary algebra.} \tag{A.16}$$

Showing that A_2 is fundamental. Note that A_2 satisfies (A.15) with $\lambda(A_2) = 1$. We want to show that $(\mathbf{v}_{A_2}, 1)$ is a fundamental solution of the system. Suppose for a contradiction that A_2 can be presented as $A_2 = G_1 + G_2$, with G_1 and G_2 non-empty normal surfaces satisfying (A.15). Since $\lambda(G_1) + \lambda(G_2) = \lambda(A_2) = 1$, we can assume without loss of generality that $\lambda(G_1) = 1$ and $\lambda(G_2) = 0$. In particular, this means that G_1 and A_2 have the same intersection with $\partial_0 M \cup \partial_v M$.

Note that A_2, being admissibly isotopic to A_1, is incompressible and boundary-incompressible, and that no component of A_2 is a sphere, a projective plane, or a disc intersecting Γ at most once. By Theorem A.9, we have that G_1 is incompressible and boundary-incompressible, and has no sphere or clean disc components. Let G_1' be the union of all the components of G_1 that intersect $\partial_0 M \cup \partial_v M$. It immediately follows that G_1' is incompressible and boundary-incompressible, and has no clean disc components. Moreover, we see that G_1' satisfies (A.15) with $\lambda(G_1') = 1$. In particular, note that each component of $G_1' \cap \partial_v M = A_2 \cap \partial_v M$ is a vertical arc. It is then easy to check that the surface G_1' satisfies the conditions of Proposition 3.3, and therefore is vertical in M. But a vertical surface is determined, up to admissible isotopy, by its intersection with $\partial_0 M$; since $G_1' \cap \partial_0 M = A_2 \cap \partial_0 M$, the surfaces G_1' and A_2 are admissibly isotopic. This, however, contradicts the fact that A_2 is least-$w_{\mathcal{E},*}$, since $w_{\mathcal{E}}(G_1') = w_{\mathcal{E}}(A_2)$ and $w(G_1') < w(A_2)$.

Conclusion. We have shown that A_2 is a fundamental solution of the system (A.15). In particular, it is admissibly isotopic to A_1, and has weight bounded by (A.16). The last step of the proof involves applying Proposition A.16 to the surfaces A_1 and A_2, to conclude that A_1 is isotopic fixing Γ to a general position surface A_3 whose weight is bounded by

$$\begin{aligned} w(A_3) &\leqslant 74\,884t^5 \cdot (w(A_2) + w(a)) \\ &\leqslant 74\,884t^5 \cdot (300t^2 \cdot 5^{7t/2} \cdot w(a) + w(a)) && \text{by (A.16)} \\ &\leqslant 2^{9t+40} \cdot w(a) && \text{by elementary algebra.} \end{aligned}$$

The surface A_3 can then be normalised by means of Proposition 2.7, obtaining a normal surface A that is isotopic to A_3 fixing Γ, and has smaller or equal weight. ▢

Bibliography

[1] 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.

[2] Sanjeev Arora and Boaz Barak. *Computational Complexity*. Cambridge University Press, Cambridge, 2009. xxiv+579.

[3] John A. Baldwin and Steven Sivek. "On the Complexity of Torus Knot Recognition". In: *Trans. Amer. Math. Soc.* 371.6 (2019), pp. 3831–3855.

[4] Filippo Baroni. "Uniformly Polynomial-Time Classification of Surface Homeomorphisms". In: *Math. Ann.* (2025).

[5] Mark C. Bell and Richard C. H. Webb. *Polynomial-Time Algorithms for the Curve Graph*. 2016. arXiv: `1609.09392 [math]`. URL: `http://arxiv.org/abs/1609.09392`.

[6] Tobias Holck Colding and David Gabai. "Effective Finiteness of Irreducible Heegaard Splittings of Non-Haken 3-Manifolds". In: *Duke Math. J.* 167.15 (2018), pp. 2793–2832.

[7] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups*. Vol. 49. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 2012. xiv+472.

[8] Albert Fathi, François Laudenbach, and Valentin Poénaru. *Thurston's Work on Surfaces*. Vol. 48. Mathematical Notes. Princeton University Press, Princeton, NJ, 2012. xvi+254.

[9] David Gabai. "Foliations and the Topology of 3-Manifolds". In: *J. Differential Geom.* 18.3 (1983), pp. 445–503.

[10] Vaibhav Gadre and Chia-Yen Tsai. "Minimal Pseudo-Anosov Translation Lengths on the Complex of Curves". In: *Geom. Topol.* 15.3 (2011), pp. 1297–1312.

[11] Robert Haraway III and Neil R. Hoffman. *On the Complexity of Cusped Non-Hyperbolicity*. 2022. arXiv: `1907.01675 [math]`. URL: `http://arxiv.org/abs/1907.01675`.

[12] W. J. Harvey. "Boundary Structure of the Modular Group". In: *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*. Vol. No. 97. Ann. of Math. Stud. Princeton Univ. Press, Princeton, NJ, 1981, pp. 245–251.

[13] 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.

[14] John Hempel. "3-Manifolds as Viewed from the Curve Complex". In: *Topology* 40.3 (2001), pp. 631–657.

[15] Nikolai V. Ivanov. *Subgroups of Teichmüller Modular Groups*. Vol. 115. Translations of Mathematical Monographs. American Mathematical Society, Providence, RI, 1992. xii+127.

[16] S. V. Ivanov. "The Computational Complexity of Basic Decision Problems in 3-Dimensional Topology". In: *Geom. Dedicata* 131 (2008), pp. 1–26.

[17] Adele Jackson. "Recognition of Seifert Fibered Spaces with Boundary Is in NP". In: *Math. Ann.* 391.1 (2025), pp. 309–361.

[18] Klaus Johannson. *Homotopy Equivalences of 3-Manifolds with Boundaries*. Vol. 761. Lecture Notes in Mathematics. Springer, Berlin, 1979. ii+303.

[19] Steven P. Kerckhoff. "The Nielsen Realization Problem". In: *Ann. of Math. (2)* 117.2 (1983), pp. 235–265.

[20] Greg Kuperberg. "Algorithmic Homeomorphism of 3-Manifolds as a Corollary of Geometrization". In: *Pacific J. Math.* 301.1 (2019), pp. 189–241.

[21] Marc Lackenby. *Some Fast Algorithms for Curves in Surfaces*. 2024. arXiv: `2401.16056 [math]`. URL: `http://arxiv.org/abs/2401.16056`.

[22] Marc Lackenby. "The Efficient Certification of Knottedness and Thurston Norm". In: *Adv. Math.* 387 (2021), Paper No. 107796, 142.

[23] Marc Lackenby and Jessica S. Purcell. "The Triangulation Complexity of Fibred 3-Manifolds". In: *Geom. Topol.* 28.4 (2024), pp. 1727–1828.

[24] Marc Lackenby and Saul Schleimer. *Recognising Elliptic Manifolds*. 2025. arXiv: `2205.08802 [math]`. URL: `http://arxiv.org/abs/2205.08802`.

[25] Jason Manning. “Algorithmic Detection and Description of Hyperbolic Structures on Closed 3-Manifolds with Solvable Word Problem”. In: *Geom. Topol.* 6 (2002), pp. 1–25.

[26] Bruno Martelli. *An Introduction to Geometric Topology*. 2023.

[27] Howard A. Masur and Yair N. Minsky. “Geometry of the Complex of Curves. I. Hyperbolicity”. In: *Invent. Math.* 138.1 (1999), pp. 103–149.

[28] Sergei Matveev. *Algorithmic Topology and Classification of 3-Manifolds*. Second. Vol. 9. Algorithms and Computation in Mathematics. Springer, Berlin, 2007. xiv+492.

[29] Jean-Pierre Otal. *The Hyperbolization Theorem for Fibered 3-Manifolds*. Vol. 7. SMF/AMS Texts and Monographs. American Mathematical Society, Providence, RI; Société Mathématique de France, Paris, 2001. xiv+126.

[30] R. C. Penner and J. L. Harer. *Combinatorics of Train Tracks*. Vol. 125. Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1992. xii+216.

[31] Grisha Perelman. *Finite Extinction Time for the Solutions to the Ricci Flow on Certain Three-Manifolds*. 2003. arXiv: `math/0307245`. URL: `http://arxiv.org/abs/math/0307245`.

[32] Grisha Perelman. *Ricci Flow with Surgery on Three-Manifolds*. 2003. arXiv: `math/0303109`. URL: `http://arxiv.org/abs/math/0303109`.

[33] Grisha Perelman. *The Entropy Formula for the Ricci Flow and Its Geometric Applications*. 2002. arXiv: `math/0211159`. URL: `http://arxiv.org/abs/math/0211159`.

[34] Saul Schleimer. *Notes on the Complex of Curves*. URL: `https://sschleimer.warwick.ac.uk/Maths/notes2.pdf`.

[35] Saul Schleimer. “Sphere Recognition Lies in NP”. In: *Low-Dimensional and Symplectic Topology*. Vol. 82. Proc. Sympos. Pure Math. Amer. Math. Soc., Providence, RI, 2011, pp. 183–213.

[36] Saul David Schleimer. “Almost Normal Heegaard Splittings”. PhD thesis. University of California, Berkeley, 2001.

[37] William P. Thurston. “A Norm for the Homology of 3-Manifolds”. In: *Mem. Amer. Math. Soc.* 59.339 (1986), i–vi and 99–130.

[38] William P. Thurston. “Three-Dimensional Manifolds, Kleinian Groups and Hyperbolic Geometry”. In: *Bull. Amer. Math. Soc. (N.S.)* 6.3 (1982), pp. 357–381.

[39] Jeffrey L. Tollefson and Ningyi Wang. “Taut Normal Surfaces”. In: *Topology* 35.1 (1996), pp. 55–75.

[40] Richard Charles Henry Webb. “Effective Geometry of Curve Graphs”. PhD thesis. University of Warwick, 2014.