

THE EFFICIENT CERTIFICATION OF KNOTTEDNESS AND THURSTON NORM

MARC LACKENBY

1. INTRODUCTION

How difficult is it to determine whether a given knot is the unknot? The answer is not known. There might be a polynomial-time algorithm, but so far, this has remained elusive. The complexity of the unknot recognition problem was shown to be in NP by Hass, Lagarias and Pippenger [10]. The main aim of this article is to establish that it is in co-NP. This can be stated equivalently in terms of the KNOTTEDNESS decision problem, which asks whether a given knot diagram represents a non-trivial knot.

Theorem 1.1. *KNOTTEDNESS is in NP.*

In some sense, this result is not new. It was first announced by Agol [1] in 2002, but he has not provided a full published proof. In 2011, Kuperberg gave an alternative proof of Theorem 1.1, but under the extra assumption that the Generalised Riemann Hypothesis is true [18]. In this paper, we provide the first full proof of the unconditional result.

Combined with the theorem of Hass, Lagarias and Pippenger [10], Theorem 1.1 gives the following corollary.

Corollary 1.2. *If either of the decision problems UNKNOT RECOGNITION or KNOTTEDNESS is NP-complete, then $NP = co-NP$.*

This is because if any decision problem in co-NP is NP-complete, then the complexity classes NP and co-NP must be equal. Since this is widely viewed not to be the case (see Section 2.4.3 in [7] for example), then it seems very unlikely that either of these decision problems is NP-complete.

Decision problems that lie in both NP and co-NP are viewed as good potential candidates for being solvable in polynomial time. However, it is very probable that some decision problems in $NP \cap co-NP$ do not lie in P. A notable example is the problem of factorising an integer which, when suitably rephrased as a decision problem, is in $NP \cap co-NP$. It is a very interesting consequence of Theorem 1.1 that UNKNOT RECOGNITION now lies in that list of problems lying in $NP \cap co-NP$ but that is not known to lie in P.

Our proof of Theorem 1.1 follows the outline given by Agol and, as in his argument, we establish more information about the genus of the knot. Recall that the *genus* of a knot K in a 3-manifold is the minimal possible genus of a *Seifert surface* for K , which is a compact orientable embedded surface, with no closed components and with boundary equal to K . When no such surface exists, the genus of K is defined to be infinite. We provide an algorithm to determine the genus of a knot in the 3-sphere. More specifically, we consider the decision problem CLASSICAL KNOT GENUS. This takes, as its input a knot diagram D and a positive integer g (given in binary), and it asks whether the knot specified by D has genus g . We establish the following result.

Theorem 1.3. *CLASSICAL KNOT GENUS is in NP.*

Theorem 1.1 is a consequence of Theorem 1.3 as follows. Given a diagram D of a non-trivial knot K , we need a certificate for its knottedness. Specifically, this certificate needs to be verifiable in polynomial time, as a function of the crossing number $c(D)$ of D . Now, the genus g of K lies between 1 and $(c(D)-1)/2$. The certificate provided by Theorem 1.3 for this diagram D and this genus g is the required certificate for knottedness.

The term ‘classical’ is meant to refer to knots in the 3-sphere. In fact, Theorem 1.3 is unlikely to generalise to knots in arbitrary closed 3-manifolds. Indeed, there is good reason to believe that the problem of determining whether a knot in a closed 3-manifold has genus g is not in NP (when both the knot and the

3-manifold are permitted to vary), because of the following theorem, which is a well known consequence of work of Agol, Hass and Thurston [2].

Theorem 1.4. *If KNOT GENUS IN COMPACT ORIENTABLE 3-MANIFOLDS is in NP, then NP = co-NP.*

Here, the decision problem KNOT GENUS IN COMPACT ORIENTABLE 3-MANIFOLDS takes as its input a triangulation of a compact orientable 3-manifold X , a knot K in X given as a subcomplex of the 1-skeleton and a positive integer g (in binary), and it asks whether the minimal genus of a Seifert surface for K is g . We will give a short proof of this result in Section 14. As a result of Theorem 1.4, it is highly unlikely that KNOT GENUS IN COMPACT ORIENTABLE 3-MANIFOLDS is in NP.

The reason why knots in arbitrary compact orientable 3-manifolds appear to be significantly more complicated than those in the 3-sphere is that when M is the exterior of a knot in the 3-sphere, $H_2(M, \partial M)$ is cyclic, and so there is only one possible homology class (up to sign) that can be represented by a Seifert surface for K . However, when M is the exterior of a knot in a general compact 3-manifold, $H_2(M, \partial M)$ may have rank bigger than 1, and so one must consider many different homology classes that could support a genus-minimising Seifert surface. However, we will see shortly that if we restrict attention to a single homology class, then an NP algorithm is available.

Although Theorem 1.3 is phrased in terms of the genus of a knot, it is not really the genus of a surface S that is its main measure of complexity in this paper. Instead, *Thurston complexity* $\chi_-(S)$ plays this role. Recall that, for a compact connected orientable surface S , $\chi_-(S)$ equals $\max\{-\chi(S), 0\}$. When S is a compact orientable surface with components S_1, \dots, S_n , then $\chi_-(S) = \sum_{i=1}^n \chi_-(S_i)$. So, when S is a Seifert surface for a knot K and $S_1 = S - \text{int}(N(K))$, then S_1 has minimal Thurston complexity in its class in $H_2(S^3 - \text{int}(N(K)), \partial N(K))$ if and only if S has minimal possible genus. However, the same is not true when K is a link with more than one component. In this case, K may have disconnected Seifert surfaces as well as connected ones.

When M is a compact orientable 3-manifold, and z is a class in $H_2(M, \partial M)$, then the *Thurston norm* of z is the minimal Thurston complexity of a compact oriented surface representing z . Theorems 1.1 and 1.3 are special cases of a more general result which allows one to efficiently determine the Thurston norm of a homology class. We define the THURSTON NORM OF A HOMOLOGY CLASS decision problem as follows. The input is a triangulation of a compact orientable irreducible 3-manifold M with boundary a (possibly empty) collection of tori, a simplicial 1-cocycle ϕ representing an element $[\phi]$ in $H^1(M)$ and a non-negative integer n . The problem asks whether the Thurston norm of the Poincaré dual to $[\phi]$ is n . The measure of complexity is the size of the input, which is, up to a bounded linear factor, equal to the sum of the number of the tetrahedra in the triangulation of M , the number of digits of the integer n in binary and the sum of the number of digits of $\phi(e)$, as e ranges over all edges of the triangulation.

The following is the main theorem of this paper.

Theorem 1.5. *THURSTON NORM OF A HOMOLOGY CLASS is in NP.*

Note that the 3-manifold M is required to have boundary a (possibly empty) collection of tori. The algorithm that we supply does not work in the case where M has boundary components of higher genus.

Note also that Theorem 1.3, and hence Theorem 1.1, follow quickly from Theorem 1.5. Given a diagram D for a knot K , one can easily build a triangulation for its exterior M in polynomial time, and where the number of tetrahedra is at most a linear function of the crossing number of D . One can also readily build a simplicial 1-cocycle ϕ representing a generator $[\phi]$ of $H^1(M)$, where the maximal value of $|\phi(e)|$ for each edge e is at most a linear function of the crossing number. Setting $n = \max\{2g - 1, 0\}$, we can apply Theorem 1.5 and thereby obtain an NP algorithm to determine whether the genus of K is g .

Our methods can also be used to certify the irreducibility of a 3-manifold. The decision problem IRREDUCIBILITY OF A COMPACT ORIENTABLE 3-MANIFOLD WITH TOROIDAL BOUNDARY AND $b_1 > 0$ takes, as its input, a triangulation of such a 3-manifold (which is permitted to have empty

boundary) and asks whether it is irreducible.

Theorem 1.6. *The decision problem IRREDUCIBILITY OF A COMPACT ORIENTABLE 3-MANIFOLD WITH TOROIDAL BOUNDARY AND $b_1 > 0$ is in NP.*

This has a consequence for the generalisation of unknot recognition to 3-manifolds other than the 3-sphere. The decision problem KNOTTEDNESS IN IRREDUCIBLE 3-MANIFOLDS takes, as its input, a triangulation of a compact orientable irreducible 3-manifold X with boundary a (possibly empty) union of tori, and a knot K given as a subcomplex of the 1-skeleton, and it asks whether K is knotted. This just means that K does not bound an embedded disc. Since the knottedness of K is closely related to whether $M - \text{int}(N(K))$ is irreducible, Theorems 1.6 and 1.5 can be used to establish the following result, which we prove in Section 15.

Theorem 1.7. *KNOTTEDNESS IN IRREDUCIBLE 3-MANIFOLDS is in $NP \cap co-NP$.*

1.1. OVERVIEW OF THE PROOF

In order to prove our main result, Theorem 1.5, one needs a method for certifying efficiently the Thurston norm of a class in $H_2(M, \partial M)$. The method that both we and Agol use is based on work of Thurston [26] and Gabai [3]. Thurston showed that if M is a compact orientable irreducible 3-manifold with boundary a (possibly empty) collection of tori, and S is a compact leaf of some taut foliation of M , then S minimises Thurston complexity in its class in $H_2(M, \partial M)$. Gabai showed that, conversely, if S is a compact orientable incompressible surface that minimises Thurston complexity in its class and that intersects each component of ∂M in a (possibly empty) collection of coherently oriented essential curves, then there is some taut foliation of M in which S appears as a compact leaf. Gabai's construction used a type of hierarchy for M , known as a sutured manifold hierarchy. A sutured manifold structure on M is a decomposition of ∂M into two subsurfaces $R_-(M)$ and $R_+(M)$, which meet along some simple closed curves γ . It is denoted (M, γ) . The surface $R_-(M)$ is given a transverse orientation pointing into M and $R_+(M)$ is transversely oriented outwards. When M is cut along a transversely oriented properly embedded surface S in general position with respect to γ , the new manifold inherits a sutured manifold structure. A sutured manifold hierarchy is a sequence of decompositions

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} (M_2, \gamma_2) \xrightarrow{S_2} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1}),$$

where (M_{n+1}, γ_{n+1}) is a collection of 3-balls, each of which intersects γ_{n+1} in a single simple closed curve, and which satisfies some mild extra conditions. Gabai used these hierarchies to construct taut foliations on the exteriors of many knots [4, 5], and was thereby able to determine their genus. It was Scharlemann [24] who realised that much of Gabai's theory could work without any reference to taut foliations; just the sutured manifold hierarchies are enough to determine the Thurston norm of a homology class. For example, it is straightforward to verify that the sequence of sutured manifold decompositions given in Figure 1 forms a taut sutured manifold hierarchy, and hence the first decomposing surface S minimises Thurston complexity in its homology class. In particular, the existence of this hierarchy proves the knot 5_1 is not the unknot.

It is sutured manifold hierarchies, such as the one in Figure 1, that we (and Agol) use to certify the Thurston norm of a homology class. The existence of such a hierarchy was proved by Gabai, but crucially, our certificate needs to be verifiable in polynomial time, as a function of the size of the input. Essentially, the sutured manifold hierarchy needs to be efficiently describable. The key to Agol's proof was to achieve this by placing some such hierarchy into 'normal' form with respect to a given triangulation for M . This normalisation procedure was based on work on Gabai [6], who considered the related problem of normalising essential laminations. In the present paper, we follow a similar approach, but instead of using triangulations, we focus on handle structures. The machinery for placing a sutured manifold hierarchy into 'normal' form with respect to a handle structure was developed by the author in [20]. This predates Agol's announcement of Theorems 1.1, 1.3 and 1.5, and plays a central role in this paper.

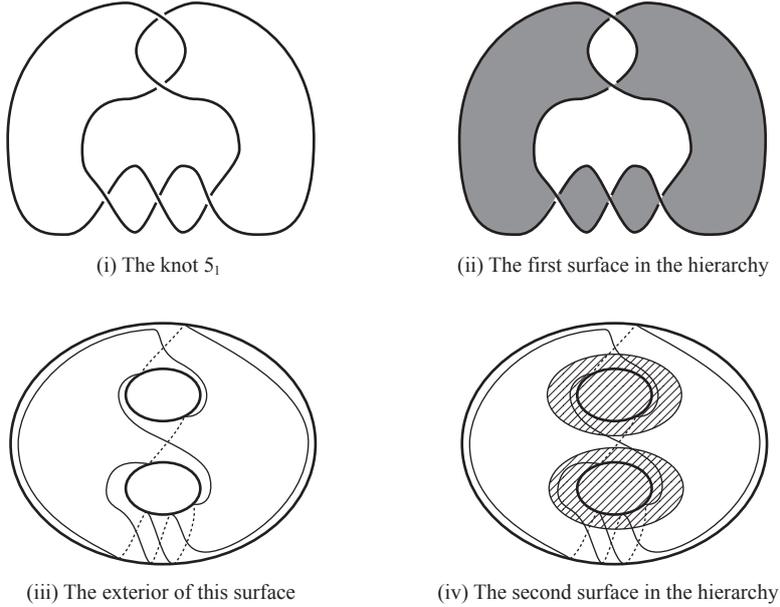


Figure 1: A sutured manifold hierarchy

1.2. NORMALISATION OF SUTURED MANIFOLD HIERARCHIES

Normal surfaces that minimise Thurston complexity in their homology class were studied by Tollefson and Wang [27]. So the first thing that we do is use their theory to realise the first surface S_1 in the hierarchy as a normal surface. However, one difficulty with normal surfaces in triangulated 3-manifolds is that, when one cuts along them, the resulting 3-manifold M_2 does not naturally inherit a triangulation. So we dualise the given triangulation of the initial manifold M to form a handle structure \mathcal{H} . There is a well-established theory of normal surfaces in handle structures [8, 13]. The next 3-manifold M_2 in the hierarchy then inherits a handle structure \mathcal{H}_2 . Unfortunately, this may be much more complicated than \mathcal{H} , in two different senses. For a start, \mathcal{H}_2 may simply have many more handles than \mathcal{H} . This happens if S_1 intersects some handle of \mathcal{H} in many discs, and then this handle of \mathcal{H} is divided into many handles of \mathcal{H}_2 . Fortunately, all but a bounded number of these handles will be very simple copies of $D^2 \times I$, lying between two parallel normal discs. These product regions patch together to form an I -bundle in M_2 , known as its ‘parallelity bundle’. It was shown in [20] how this parallelity bundle may be removed, primarily by decomposing along the annuli that form its intersection with the remainder of M_2 . So, a key part of our argument is the analysis of this I -bundle and an algorithmic method of removing it from M_2 .

However, there is another reason why \mathcal{H}_2 may be more complicated than \mathcal{H} . It is not at all clear that the local structure of the 0-handles of \mathcal{H}_2 is simpler than that of \mathcal{H} . This issue was faced right at the very first use of hierarchies by Haken [8] and Waldhausen [28]. They defined a notion of complexity for a handle structure of a 3-manifold, and showed that, when it is decomposed along a normal surface, then the complexity does not go up. Unfortunately, their notion of complexity does not fit well with sutured manifolds, primarily because it does not take account of the sutures. Fortunately, this issue was resolved in [20]. A variation of normal form more suited to sutured manifolds was introduced. In this paper, we call such surfaces ‘regulated’. Also, a notion of complexity for a handle structure of a sutured manifold was defined in [20]. It was also shown that this does not go up when the manifold is decomposed along a regulated surface. Moreover, it was shown that sutured manifold hierarchies can always be found where each decomposing surface is regulated.

Therefore, it is regulated surfaces that are used in this paper. Unfortunately, they come with their own complications. Although it is the case that we may arrange for the decomposing surfaces to be regulated, they may fail to satisfy one of the key technical requirements for a sutured manifold hierarchy. Some

curves of intersection with the surface R_{\pm} may be simple closed curves bounding discs in R_{\pm} . Such curves are called ‘trivial’. Because surfaces with trivial boundary components are not permitted to be part of a sutured manifold hierarchy, it is not clear that they can be used as part of a certificate for Thurston norm. However, we develop a theory of ‘allowable hierarchies’, where the decomposing surfaces may have trivial boundary curves, but which can nonetheless be used to certify Thurston norm.

So, allowable hierarchies of regulated surfaces will be used as part of our certificate. But these surfaces must be describable in an efficient way. In particular, it is important that our surfaces intersect each handle in at most c^h discs, where c is a universal constant and h is the number of handles in the initial handle structure. To be able to establish such a bound, we use methods from linear algebra, that go back to Haken [8]. We show how regulated surfaces S can be described by means of solution (S) to a system of linear equations, much in the same way that normal surfaces can be. When S , S_1 and S_2 are regulated surfaces and $(S) = (S_1) + (S_2)$, then we say that S is a ‘sum’ of S_1 and S_2 . Just as in the normal surface case, one can place S_1 and S_2 into general position, and then obtain S by resolving the arcs and curves of intersection. It is crucial for our purposes that we may find a decomposing surface that is ‘fundamental’, which means that it cannot then be expressed as a sum of other non-trivial surfaces. This is because fundamental surfaces have a bounded number of discs of intersection with each handle, by methods that go back to Hass and Lagarias [9].

Therefore, we must analyse the case when S is a sum of surfaces S_1 and S_2 . Now, S_1 and S_2 need not inherit orientations from S , and indeed they need not even be orientable. But when they do inherit orientations, then the situation is fairly easy to analyse. It turns out that we can generally show that decomposition along S_1 or S_2 is taut, and so decompose along one of these instead. These are ‘simpler’ surfaces, and so in this way, one may arrange for decomposing surfaces to be fundamental. When S_1 and S_2 do not inherit transverse orientations, then the aim is to show that S was not as simple as it could have been. One can perform an ‘irregular switch’ along one of the arcs or curves of $S_1 \cap S_2$, creating a new transversely oriented surface S' . We show that decomposition along S' is also taut. This is possible when the irregular switch takes place along a curve of $S_1 \cap S_2$. However, the argument does not work in the case of an arc of $S_1 \cap S_2$, because the orientations of R_{\pm} at its endpoints may be problematic. Fortunately, regulated surfaces rescue us here, because their boundaries are very tightly controlled, and in fact, it is possible to show that the regular switch along arcs of $S_1 \cap S_2$ always respect the orientation of the surface.

Thus, we may arrange that the decomposing surfaces are fundamental, and hence have an exponential bound on complexity. But decomposing along these surfaces is not straightforward, because we never want to deal with handle structures having exponentially many handles. They are too unwieldy to be efficiently describable within our certificate. Fortunately, there is technology due to Agol, Hass and Thurston [2] which is applicable. Using their methods, we show that, given a surface S_i with an exponential bound on its complexity, it is possible to determine the topological types of the components of the parallelity bundle for the manifold M_{i+1} obtained by decomposing along S_i . So, we never need to construct the handle structure for M_{i+1} . Instead, we can go directly to the handle structure for the manifold M'_{i+1} that is obtained by removing the parallelity bundle.

Thus, our certificate for Thurston norm is comprised (primarily) of the following pieces of information: handle structures for a sequence of 3-manifolds, and regulated surfaces within these manifolds, expressed as solutions to a system of equations. We verify this certificate by checking that the next manifold is indeed obtained from the previous one by cutting along the regulated surface and then purging the resulting manifold of its parallelity bundle.

1.3. COMPARISON WITH AGOL’S STRATEGY

The strategy of our proof, as explained in Section 1.1, is very similar to the one outlined by Agol [1]. However, the details, as given in Section 1.2, are very different. Instead of using handle structures, Agol used triangulations. Since the result of decomposing a triangulation along a normal surface is not in general a triangulation, Agol had to work hard to build a triangulation for each manifold in the hierarchy.

His technique was based on placing the surface into ‘spun’ normal form. The method of doing this was based on Gabai’s method for normalising taut foliations, as explained in [6], which used ‘sutured manifold evacuation’. In Agol’s argument, the surface still needed to be made ‘fundamental’ in some way. Moreover, an analogue of the removal of parallelity bundles would still need to have been achieved, again by the use of the algorithm of Agol, Hass and Thurston [2]. Agol’s strategy certainly has some advantages, but it seems to us that the extensive use of the established techniques from [20], as followed in this paper, is also very convenient.

1.4. STRUCTURE OF THE PAPER

In Section 2, we recall some of the basic theory of sutured manifolds. In Section 3, we introduce ‘decorated’ sutured manifolds and ‘allowable’ hierarchies. The key result here is Theorem 3.1 which implies that allowable hierarchies can be used to certify Thurston norm. In Section 4, we show that if a surface extends to an allowable hierarchy, then certain modifications can be made to it, and the resulting surface still extends to an allowable hierarchy. In Section 5, we introduce handle structures for sutured manifolds, and regulated surfaces. We also introduce the complexity of a handle structure, and explain how it behaves when a decomposition along a regulated surface is performed. In Section 6, we give some simplifications that can be made to a handle structure. We also explain how we need only to work with a finite universal list of 0-handle types. In Section 7, we develop an algebraic theory of regulated surfaces. We introduce a minor variant, known as ‘boundary-regulated’ surfaces, and show how they may be expressed as solutions to certain equations. We also explain how summation of boundary-regulated surfaces can be interpreted topologically. Finally, we show that decompositions can always be made along fundamental regulated surfaces. This central result is Theorem 7.8. In Section 8, we bound the complexity of fundamental surfaces, using tools from linear algebra. In Section 9, we recall an algorithm of Agol, Hass and Thurston [2], and show how it can be used to determine the parallelity bundle for the manifold obtained by decomposing along a surface. In Section 10, we show how to remove this parallelity bundle algorithmically. In Section 11, we show how it suffices to focus our attention on 3-manifolds that are atoroidal. This convenient hypothesis occurs at several points in the preceding argument. In Section 12, we show how to certify efficiently that a sutured manifold is a product. This is useful because the hierarchies that we use terminate in products rather than 3-balls, for mostly technical reasons. We then go on to prove the main theorem in the special case where the 3-manifold is Seifert fibred. In Section 13, we describe in detail the certificate for Thurston norm, we show why it always exists, and how it may be verified in polynomial time. In Sections 14 and 15, we give the proofs of Theorems 1.4 and 1.7.

2. SUTURED MANIFOLDS

A *sutured manifold* is a compact orientable 3-manifold M , with its boundary decomposed into two compact subsurfaces $R_-(M)$ and $R_+(M)$, in such a way that $R_-(M) \cap R_+(M)$ is a collection of simple closed curves γ . These curves are called *sutures*. The surfaces $R_-(M)$ and $R_+(M)$ are assigned transverse orientations, with $R_-(M)$ pointing into M and $R_+(M)$ pointing outwards. The sutured manifold is usually denoted (M, γ) .

A compact oriented surface S embedded in a 3-manifold M , with ∂S in ∂M , is called *taut* if S is incompressible and it has minimal Thurston complexity in its homology class in $H_2(M, N(\partial S))$, where $N(\partial S)$ is a regular neighbourhood of ∂S in ∂M .

Note that, in this definition, we consider the homology group $H_2(M, N(\partial S))$, not $H_2(M, \partial M)$. Hence, it is possible for S to be taut and yet not have minimal Thurston complexity in its class in $H_2(M, \partial M)$. However, the following simple lemma (which appears as Lemma A.7 in [19]) implies that this phenomenon cannot occur in an important case, when ∂M is a (possibly empty) union of tori. It is for this reason that we restrict, in Theorem 1.5, to manifolds of this form.

Lemma 2.1. *Let M be a compact orientable 3-manifold with boundary a (possibly empty) union of tori. Let S be a compact oriented properly embedded surface, such that no component of ∂S bounds a disc in ∂M . Then S has minimal Thurston complexity in its class in $H_2(M, N(\partial S))$ if and only if it has minimal Thurston complexity in its class in $H_2(M, \partial M)$.*

A sutured manifold (M, γ) is *taut* if M is irreducible and $R_-(M)$ and $R_+(M)$ are both taut.

Suppose that (M, γ) is a sutured manifold and that S is a compact, transversely oriented surface properly embedded in M . Suppose that ∂S intersects γ transversely. Then $M' = \text{cl}(M - N(S))$ inherits a sutured manifold structure, since both the parts of $\partial M'$ arising from ∂M and the parts of $\partial M'$ arising from $N(S)$ have a well-defined transverse orientation. This is called a *sutured manifold decomposition* and is denoted

$$(M, \gamma) \xrightarrow{S} (M', \gamma').$$

It is called *taut* if (M, γ) and (M', γ') are taut. The surface S is often called a *decomposing surface*. Note that in a taut decomposition, we do not require S to be taut, although this will often be the case.

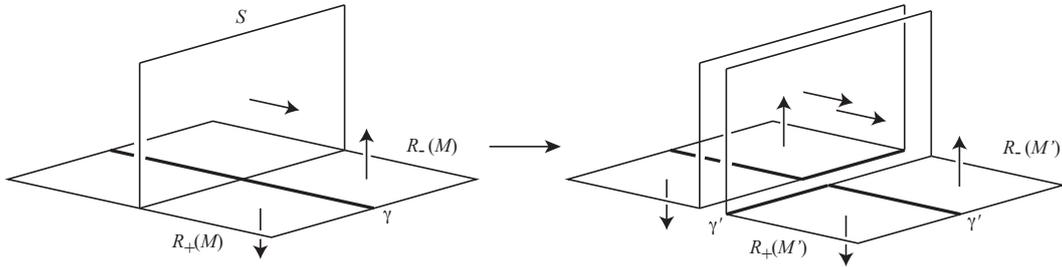


Figure 2: Sutured manifold decomposition

In general, decomposing a taut sutured manifold (M, γ) along a taut surface S does not result in a taut decomposition. This is because the resulting sutured manifold (M', γ') need not be taut. However, there is one important situation where this is the case, which is summarised in the following lemma.

Lemma 2.2. *Let M be a compact orientable irreducible 3-manifold with boundary a (possibly empty) collection of incompressible tori. Give M a sutured manifold structure (M, γ) where $R_+(M) = \partial M$ and $R_-(M) = \emptyset$. Let S be a compact oriented properly embedded surface, such that $S \cap \partial M$ is a (possibly empty) collection of essential simple closed curves. Suppose also that no collection of annular components of S is trivial in $H_2(M, \partial M)$, and that no component of S is a 2-sphere. Then, if S is taut, so is the decomposition*

$$(M, \gamma) \xrightarrow{S} (M', \gamma').$$

Proof. The hypotheses on M imply that (M, γ) is taut. So, we need only show that (M', γ') is taut.

Note first no component of S is a disc, because $S \cap \partial M$ is essential and ∂M is incompressible. Hence, no component of $R_+(M')$ or $R_-(M')$ is a sphere or disc. Therefore,

$$\chi_-(R_+(M')) = -\chi(R_+(M')) = -\chi(S) - \chi(\partial M) = -\chi(S) = \chi_-(S) = -\chi(R_-(M')) = \chi_-(R_-(M')).$$

Hence, $R_+(M')$ has minimal Thurston complexity in its class in $H_2(M', N(\gamma'))$ if and only if $R_-(M')$ does. But $R_-(M')$ is a copy of S . So if $R_-(M')$ was homologous in $H_2(M', N(\gamma'))$ to a surface with smaller Thurston complexity, then S would be homologous in $H_2(M, N(\partial S))$ to a surface with smaller Thurston complexity, contradicting its tautness.

We now show that M' is irreducible. Consider a reducing sphere for M' . This bounds a ball in M . The intersection between this ball and S must be non-empty. Since S is incompressible, it must therefore contain a 2-sphere component, contrary to our assumption.

Now suppose that $R_-(M')$ is compressible in M' . But $R_-(M')$ is parallel in M to S , and so this would imply that S is compressible. Finally suppose that $R_+(M')$ is compressible in M' . Then, since $R_+(M')$ minimises Thurston complexity, this implies that $R_+(M')$ contains a component that is a compressible annulus or a compressible torus. Now, $R_+(M')$ consists of a copy of S , with annuli from ∂M attached, plus possibly some toral components of ∂M . So, if $R_+(M')$ contains a compressible annulus, then ∂M or S compresses in M , which is contrary to assumption. So, $R_+(M')$ contains a compressible torus, which therefore bounds a solid torus, by the irreducibility of M' . This solid torus gives a homology in $H_2(M, \partial M)$ between a collection of annular components of S and a collection of annuli in ∂M . This again is contrary to assumption. \square

One of the most important results in sutured manifold theory is summarised in the phrase ‘tautness usually pulls back’. The precise theorem is as follows. (See Theorem 3.6 in [24]).

Theorem 2.3. *Let*

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

be a sutured manifold decomposition. Suppose that no component of ∂S bounds a disc in ∂M disjoint from γ , and that no component of S is a disc disjoint from γ that forms a compression disc for a solid toral component of (M, γ) with no sutures. Then if (M', γ') is taut, so is (M, γ) and so is S .

As a consequence, if we have a sequence of sutured manifold decompositions, each satisfying the requirements of Theorem 2.3, and we can show that the final sutured manifold is taut, then the entire sequence of decompositions is taut, and each decomposing surface is taut. It is a reasonably straightforward matter to verify that a sutured manifold structure on a collection of 3-balls is taut. Indeed, it is clear that it is taut if and only if each ball contains at most one suture. Thus, this forms the basis for our certificate for Thurston norm. The existence of such a sequence of decompositions is guaranteed by the following central result of Gabai [3] (see also Theorems 2.6 and 4.19 in [24]).

Theorem 2.4. *Let (M, γ) be a taut sutured manifold, and let z be a non-trivial element of $H_2(M, \partial M)$. Then there exists a sequence of taut decompositions*

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1})$$

such that

- (i) *no component of any ∂S_i bounds a disc in ∂M_i disjoint from γ_i ;*
- (ii) *no component of any S_i is a disc disjoint from γ_i ;*
- (iii) $[S_1, \partial S_1] = z \in H_2(M, \partial M)$;
- (iv) (M_{n+1}, γ_{n+1}) *is a collection of taut 3-balls.*

The *length* of this sequence of decompositions is n .

There are two types of decomposing surface that arise frequently. A *product disc* in a sutured manifold (M, γ) is a properly embedded disc that intersects γ transversely at two points. A *product annulus* is a properly embedded annulus A that is disjoint from γ , and that has one component of ∂A in $R_-(M)$ and one component of ∂A in $R_+(M)$. A product annulus A is *trivial* if some component of ∂A bounds a disc in $R_-(M)$ or $R_+(M)$. (Note that this latter definition is a minor variation of the one given in Definition 4.1 of [24].)

Here, we have the stronger form of Theorem 2.3 (see Lemma 4.2 in [24]).

Proposition 2.5. *Let*

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

be a sutured manifold decomposition, where S is either a product disc or a non-trivial product annulus. Then (M, γ) is taut if and only if (M', γ') is taut.

Finally, a *product sutured manifold* is of the form $F \times [-1, 1]$, where F is some compact orientable surface and where $\gamma = \partial F \times \{0\}$. We will frequently refer to the product sutured manifold simply by $F \times [-1, 1]$, with the understanding that $\gamma = \partial F \times \{0\}$. We note that when a decomposition

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

is performed along a product disc or non-trivial product annulus, then (M, γ) is a product sutured manifold if and only if (M', γ') is.

3. DECORATED SUTURED MANIFOLDS

Unfortunately, we must consider decompositions

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

where some curves of ∂S bound discs in ∂M disjoint from γ . This leads to some complications, because Theorem 2.3 does not apply, and so it is not obvious that S can be used as part of a certificate for Thurston norm. To get around this problem, we keep track of the curves of ∂S that bound discs in ∂M disjoint from γ . These give rise to ‘special’ sutures of γ' . We want to be able to distinguish these sutures, and so we now introduce a structure where this is possible, called a decorated sutured manifold.

3.1. DEFINITION

We define a *decorated sutured manifold* to be a sutured manifold (M, γ) where some of the sutures are distinguished. These distinguished sutures are called *u-sutures*, where ‘u’ stands for untouchable.

Given a decorated sutured manifold (M, γ) , its *canonical enlargement*, denoted $E(M, \gamma)$ is an undecorated sutured manifold that is obtained by removing each u-suture and attaching a 2-handle $D^2 \times [0, 1]$ along this curve. One component of $D^2 \times \{0, 1\}$ lies in $R_-(E(M, \gamma))$ and the other lies in $R_+(E(M, \gamma))$. We denote the underlying 3-manifold of this enlargement by $E(M)$.

3.2. PRE-BALLS AND PRE-SPHERICAL PRODUCTS

A *pre-ball* is a decorated sutured manifold of the form $P \times [-1, 1]$, where P is a compact connected planar surface. The sutures are required to be $\partial P \times \{0\}$ and exactly one is not a u-suture. Note that the canonical enlargement of a pre-ball is a taut sutured ball.

A *pre-spherical product manifold* is a decorated sutured manifold of the form $P \times [-1, 1]$, where P is a compact connected planar surface, where $\partial P \times \{0\}$ forms the sutures, and every suture is a u-suture. In this case, its canonical enlargement is a product sutured manifold of the form $S^2 \times [-1, 1]$.

3.3. TRIVIAL CURVES

Let (M, γ) be a decorated sutured manifold, and let C be a simple closed curve disjoint from the sutures. Then C is *trivial* if there is a planar surface P embedded in $R_{\pm}(M)$ such that

- (i) C is a component of ∂P ; and
- (ii) $P \cap \gamma$ is the remaining components of ∂P , and each is a u-suture.

The planar surface P is called the *trivialising planar surface*.

An alternative way of defining a trivial curve is that it is a simple closed curve in $R_{\pm}(M)$ that bounds a disc in the boundary of $E(M)$ disjoint from the sutures of $E(M, \gamma)$. In this way, we see that a trivial

curve cannot bound trivialising planar surfaces on both sides, unless some component of $\partial E(M, \gamma)$ is a sphere with no sutures.

Let S be a surface properly embedded in M . Then a *trivial boundary curve* of S is a boundary component of S that is trivial. Note that, by definition, trivial boundary curves are disjoint from the sutures.

3.4. ALLOWABLE DECOMPOSITIONS AND HIERARCHIES

If (M, γ) is a decorated sutured manifold, then a sutured manifold decomposition

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

is *allowable* provided that

- (i) S is disjoint from the u-sutures of γ , and
- (ii) if C is any trivial boundary curve for S , then its trivialising planar surface inherits a transverse orientation from $R_{\pm}(M)$ that agrees with S near C .

In this case, (M', γ') inherits a decoration, by declaring that a component of γ' is a u-suture if one of the following holds:

- (i) it came from a u-suture of γ , or
- (ii) it came from a trivial boundary curve of S .

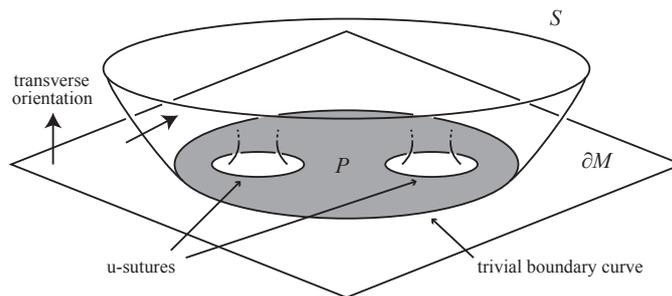


Figure 3: Trivial boundary curve in an allowable decomposition

A sequence of allowable decompositions is called an *allowable hierarchy* if

- (i) the final sutured manifold is a product manifold, no component of which pre-spherical; and
- (ii) no component of any decomposing surface is planar, disjoint from the sutures and where all but at most one of its boundary components is trivial.

We say that the allowable hierarchy is *complete* if the final sutured manifold is a collection of pre-balls.

3.5. THE UTILITY OF ALLOWABLE HIERARCHIES

Before we go any further with the development of the theory of allowable hierarchies, we state a result which explains why they are useful.

Theorem 3.1. *Let M be a compact orientable irreducible 3-manifold with boundary a (possibly empty) union of tori. Give M the sutured manifold structure with $R_+(M) = \partial M$ and $R_-(M) = \emptyset$. Let S be a compact orientable properly embedded surface with no 2-sphere or disc components, such that the intersection between S and each component of ∂M is a (possibly empty) collection of coherently oriented*

essential curves. Then S is incompressible and has minimal Thurston complexity in its class in $H_2(M, \partial M)$ if and only if it is the first surface in allowable hierarchy.

Thus, allowable hierarchies will be the method that we employ to certify Thurston norm. We will prove Theorem 3.1 in Section 3.10.

3.6. THE CANONICAL ENLARGEMENT OF A SUTURED MANIFOLD DECOMPOSITION

Given an allowable decomposition

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

between decorated sutured manifolds, there is an associated decomposition

$$E(M, \gamma) \xrightarrow{E(S)} E(M', \gamma'),$$

called its *canonical enlargement* and which is defined as follows. Consider any trivial simple closed curve of ∂S . Attach its trivialising planar surface P to S , and push this new part of the surface a little into the interior of M . Note that it is possible for two different trivialising planar surfaces P and P' to intersect. But their boundary components that are not u-sutures are disjoint. Hence, P and P' are nested. If P' is contained in P , say, then we push P' a little further into the interior of M . We now extend this surface into the 2-handles that are attached to the u-sutures of M . For each such 2-handle, we insert a collection of discs, so that the resulting surface is properly embedded. The resulting surface is $E(S)$. Then it is clear that decomposing $E(M, \gamma)$ along $E(S)$ gives the same sutured manifold as $E(M', \gamma')$. For a pictorial proof of this fact, see Figure 4.

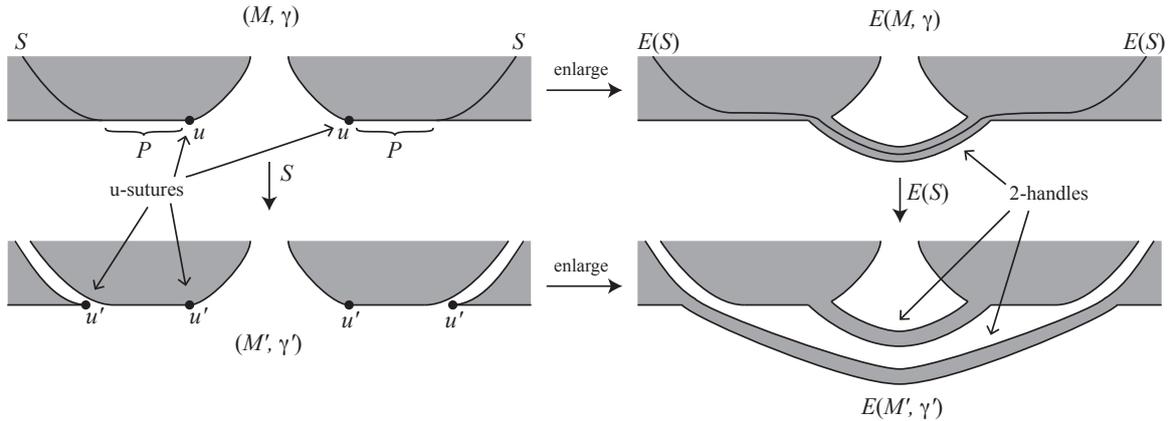


Figure 4: Interchanging the order of enlargement and decomposition

It is partly because of this canonical enlargement that u-sutures are ‘untouchable’, in the sense that decomposing surfaces must avoid them. If, on the contrary, a decomposing surface S were allowed to run over a u-suture, then there would be no good way of defining $E(S)$, because the boundary of S would run over the attaching annulus of the 2-handle.

We now note that $E(S)$ satisfies one of the key hypotheses of Theorem 2.3.

Lemma 3.2. *No boundary curve of $E(S)$ is trivial in $E(M, \gamma)$.*

Proof. Suppose that a boundary curve of $E(S)$ is trivial. It therefore bounds a disc D in $\partial E(M)$ disjoint from the sutures. By choosing the curve of $\partial E(S)$ appropriately, we may ensure that the interior of D is disjoint from $E(S)$. The intersection between this disc D and M is a planar surface P . All but one of its boundary curves are u-sutures of γ . The remaining boundary component is $P \cap \partial S$. The orientations

of P and S must agree locally near $P \cap \partial S$, by the definition of an allowable decomposition. Hence, this boundary curve of ∂S is removed in the construction of $E(S)$, which is a contradiction. \square

3.7. THE CANONICAL ENLARGEMENT OF AN ALLOWABLE HIERARCHY

Lemma 3.3. *The canonical enlargement of an allowable hierarchy of decorated sutured manifolds is taut and each decomposing surface in this canonical enlargement is taut.*

Proof. Let

$$E(M_1, \gamma_1) \xrightarrow{E(S_1)} E(M_2, \gamma_2) \xrightarrow{E(S_2)} \dots \xrightarrow{E(S_n)} E(M_{n+1}, \gamma_{n+1})$$

be this enlargement. By assumption, (M_{n+1}, γ_{n+1}) is a product sutured manifold, no component of which is pre-spherical. Hence, $E(M_{n+1}, \gamma_{n+1})$ is a taut product manifold. By Lemma 3.2, no boundary curve of any $E(S_i)$ bounds a disc in $\partial E(M_i)$ disjoint from the sutures. Also, by (ii) in the definition of an allowable hierarchy, no component of any $E(S_i)$ is a disc disjoint from the sutures. Hence, by Theorem 2.3, each decomposition in the sequence is taut and each decomposing surface is taut. \square

Corollary 3.4. *Let*

$$(M_1, \gamma_1) \xrightarrow{S_1} (M_2, \gamma_2) \xrightarrow{S_2} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1})$$

be an allowable hierarchy. Suppose that (M_1, γ_1) has no u-sutures and that no boundary component of S_1 is trivial in (M_1, γ_1) . Then S_1 is taut in M_1 .

Proof. By Lemma 3.3, the canonical enlargement of this allowable hierarchy is taut and each decomposing surface is taut. But (M_1, γ_1) has no u-sutures and therefore $E(M_1, \gamma_1)$ is just (M_1, γ_1) . Also, since S_1 has no trivial boundary curves, $E(S_1)$ is just S_1 . Hence, S_1 is taut in M_1 . \square

3.8. AN ALTERNATIVE INTERPRETATION OF COMPLETE ALLOWABLE HIERARCHIES

The following lemma provides a useful alternative way of understanding complete allowable hierarchies. Its proof is immediate.

Lemma 3.5. *Let*

$$(M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1})$$

be a sequence of allowable decompositions between decorated sutured manifolds. Then no boundary curve of $E(S_i)$ is trivial in $E(M_i, \gamma_i)$. Furthermore, this forms a complete allowable hierarchy if and only if its canonical enlargement satisfies the following conditions:

- (i) *the final manifold $E(M_{n+1}, \gamma_{n+1})$ is a collection of taut 3-balls, each of the form $D^2 \times [-1, 1]$, where the sutures are $\partial D^2 \times \{0\}$ and where the co-cores of the attached 2-handles are vertical in this product structure;*
- (ii) *no component of $E(S_i)$ is a sphere or disc disjoint from the sutures.*

3.9. ENLARGEMENT AND TAUTNESS

Proposition 3.6. *Let (M, γ) be a decorated sutured manifold. If $E(M, \gamma)$ is taut, then so is (M, γ) .*

Proof. We will prove the following stronger statement. If (M, γ) is a decorated sutured manifold, a single suture is removed, a 2-handle is attached along it and the resulting sutured manifold (M', γ') is taut, then (M, γ) is also taut.

Note first that we may assume that M is connected. For we may restrict attention to the component of M to which the 2-handle is attached.

Consider a 2-sphere S properly embedded in M . Since M' is irreducible, S bounds a ball in M' . This

ball is disjoint from $\partial M'$, and so is disjoint from the attached 2-handle. Therefore, S bounds a ball in M . Hence, M is irreducible.

Suppose now that $R_{\pm}(M)$ is compressible, via a compression disc D . Since $R_{\pm}(M')$ is incompressible, ∂D bounds a disc D' in $\partial M'$ disjoint from the sutures. Note that at least one of the two discs of intersection between the 2-handle and $R_{\pm}(M')$ misses D' , because one of these discs lies in $R_-(M')$ and the other lies in $R_+(M')$. Since M' is irreducible, $D \cup D'$ bounds a ball in M' . This ball cannot intersect the 2-handle, because this would imply that D intersected this 2-handle, whereas D lies in M . Hence, we deduce that the ball lies in M , and in particular, D' lies in M . Therefore, this was not a compression disc for $R_{\pm}(M)$, which is a contradiction.

We now show that $R_{\pm}(M)$ has minimal Thurston complexity in its class in $H_2(M, N(\gamma))$. Consider another surface S in the same class. We may assume that this runs over the attaching annulus of the 2-handle in a single essential curve. Hence, we may extend S to a surface \tilde{S} in M' , by attaching a disc in this 2-handle. Note that \tilde{S} is in the same class as $R_{\pm}(M')$ in $H_2(M', N(\gamma'))$. Since this has minimal Thurston complexity, we deduce that $\chi_-(\tilde{S}) \geq \chi_-(R_{\pm}(M'))$.

Now, we may assume that $R_{\pm}(M')$ has the same number of sphere and disc components as $R_{\pm}(M)$. For if a disc component of $R_{\pm}(M')$ is created, then (M', γ') is a taut 3-ball and so $R_{\pm}(M)$ consists of incompressible annuli, and therefore has minimal Thurston complexity. If a sphere component of $R_{\pm}(M')$ is created, this bounds a ball, by the irreducibility of M' , and hence one of $R_-(M')$ or $R_+(M')$ is empty, which is impossible, because they both have non-empty intersection with the 2-handle. Similarly, \tilde{S} has the same number of sphere and disc components as S . Hence,

$$\chi_-(S) = \chi_-(\tilde{S}) + 1 \geq \chi_-(R_{\pm}(M')) + 1 = \chi_-(R_{\pm}(M))$$

which implies that $R_{\pm}(M)$ does indeed minimise Thurston complexity in its homology class. \square

Corollary 3.7. *An allowable hierarchy of decorated sutured manifolds is taut.*

Proof. By Lemma 3.3, the canonical enlargement of an allowable hierarchy is taut. Hence, by Proposition 3.6, each of the original decorated sutured manifolds is taut. \square

Note that the converse of Proposition 3.6 is false. Consider the case where $E(M, \gamma)$ is a 3-ball with three parallel sutures in its boundary. Then $R_-(E(M, \gamma))$ is an annulus and a disc, as is $R_+(E(M, \gamma))$. In particular, $E(M, \gamma)$ is not taut. However, we may pick a properly embedded arc in $E(M)$, running between the disc component of $R_-(E(M, \gamma))$ and the disc component of $R_+(E(M, \gamma))$. Remove a regular neighbourhood of this arc to form (M, γ) . Then M is irreducible, and $R_{\pm}(M)$ consists of incompressible annuli. Therefore, (M, γ) is taut. This example shows that admitting an allowable hierarchy is a stronger condition than simply being taut.

3.10. ALLOWABLE HIERARCHIES AND TAUTNESS

We are now in a position to be able to prove Theorem 3.1.

Proof. Let M and S be as in the statement of Theorem 3.1. Suppose that S is incompressible and has minimal Thurston complexity in its class in $H_2(M, \partial M)$. Then by Lemma 2.1, S is taut. So, by Lemma 2.2, the decomposition

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

is taut. Note that no collection of annular components of S is trivial in $H_2(M, \partial M)$ by our assumption about the orientations on ∂S , and so Lemma 2.2 does apply. Note also that (M', γ') has no u-sutures. By Theorem 2.4, (M', γ') admits a taut sutured manifold hierarchy, where no decomposing surface has a disc component disjoint from the sutures. Also, at no stage does a boundary curve bound a disc in the boundary of the manifold disjoint from the sutures. So, no u-sutures are created. Thus, this forms an allowable hierarchy.

Conversely, suppose that S is the first surface in an allowable hierarchy. This hierarchy is taut by Corollary 3.7, and in particular the first decomposition is taut. So, by Theorem 2.3, S is taut, and so by Lemma 2.1, S is incompressible and has minimal Thurston complexity in its class in $H_2(M, \partial M)$. \square

3.11. MODIFYING THE DECORATION

Lemma 3.8. *Let (M, γ) be a decorated sutured manifold, and let (M', γ') be the same sutured manifold, but where some of the u-sutures have been replaced by ordinary sutures. Suppose that (M, γ) has an allowable hierarchy. Then so does (M', γ') , with the same reduced length.*

Proof. Let

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1})$$

be the allowable hierarchy for (M, γ) . We can view this as a hierarchy

$$(M', \gamma') = (M'_1, \gamma'_1) \xrightarrow{S'_1} \dots \xrightarrow{S'_n} (M'_{n+1}, \gamma'_{n+1})$$

for (M', γ') . We show by induction the following:

- (i) each (M'_i, γ'_i) inherits a decoration;
- (ii) each u-suture of γ'_i is a u-suture of γ_i ;
- (iii) if a curve in $\partial M'_i - \gamma'_i$ is trivial, then so is the corresponding curve in $\partial M_i - \gamma_i$, with the same trivialising planar surface;
- (iv) decomposition along S'_i is allowable;
- (v) no component of S'_i is planar, disjoint from the sutures and where all but at most one of its boundary components is trivial.

The induction starts with the case $i = 1$, where (i) and (ii) are part of the hypotheses of the lemma. The three remaining claims (iii), (iv) and (v) in the case $i = 1$ are proved in the same way as in the inductive step. So, suppose that (M'_i, γ'_i) satisfies (i) and (ii). Then any trivialising planar surface for a curve in $\partial M'_i - \gamma'_i$ is then a trivialising planar surface in $\partial M_i - \gamma_i$. Hence, we have (iii). Then (ii) and (iii) imply that S'_i is allowable, giving (iv). Also, by (ii), a component of S'_i violating (v) would give a component of S_i that could not be part of an allowable hierarchy. So, we obtain (v). Since decomposition along S'_i is allowable, $(M'_{i+1}, \gamma'_{i+1})$ inherits a decoration, giving (i). The u-sutures of γ'_{i+1} come either from u-sutures of γ'_i or from trivial boundary curves of S'_i . In the former case, (ii) for (M'_i, γ'_i) implies that it is also a u-suture of γ_i . In the latter case, this is also a trivial boundary curve of S_i , by (iii). Hence, we obtain (ii) for $(M'_{i+1}, \gamma'_{i+1})$. This completes the induction.

The final manifold (M_{n+1}, γ_{n+1}) is a product sutured manifold, no component of which is pre-spherical. With the new decoration, some of the u-sutures have become ordinary sutures. So, $(M'_{n+1}, \gamma'_{n+1})$ is still a product sutured manifold, no component of which is pre-spherical. \square

3.12. CANONICAL EXTENSIONS THAT ARE BALLS

Lemma 3.9. *Let (M, γ) be a decorated sutured manifold that admits an allowable hierarchy. Suppose that $E(M)$ is a 3-ball. Then (M, γ) is a pre-ball.*

Proof. Note first that, by Lemma 3.3, $E(M, \gamma)$ is taut, and hence is a product sutured 3-ball B . The manifold (M, γ) is obtained from $E(M, \gamma)$ by removing the attached 2-handles. The co-cores of these 2-handles form a tangle t in B . Thus, the lemma is equivalent to the assertion is that there is an ambient isotopy, keeping the suture in B fixed, after which the tangle respects the product structure on B . Equivalently, there is a collection of disjoint embedded discs D embedded in B , such that

- (i) the intersection between each component D' of D and t is a single component of t in $\partial D'$;

- (ii) the remainder of the boundary of D' lies in ∂B ;
- (iii) the intersection between each component of D and the suture of B is a single point;
- (iv) each component of t lies in a component of D .

Note that there is considerable flexibility in the choice of D . In particular, it may be chosen so that its intersection with $R_+(B)$ is any given collection of disjoint embedded arcs, with the property that each arc starts at a point of $t \cap R_+(B)$ and ends on the suture, and each point of $t \cap R_+(B)$ lies at the endpoint of such an arc.

Now, (M, γ) admits an allowable hierarchy where each surface is connected. We prove the lemma by induction on the length of such a hierarchy. The induction starts trivially, because a sutured manifold that has an allowable hierarchy with length zero is a product, and the only product sutured manifold with canonical extension that is a ball is a pre-ball. So, we consider the inductive step. Let

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

be the first decomposition in the hierarchy. By assumption $E(M, \gamma)$ is a ball. It is taut by Lemma 3.3. The surface $E(S)$ is taut in $E(M, \gamma)$ by Lemma 3.3. In particular, it is incompressible, and is therefore a disc. Therefore, $E(M_S, \gamma_S)$ is two 3-balls B_1 and B_2 . By induction therefore, (M_S, γ_S) is two pre-balls. Let t_1 and t_2 be the tangles in B_1 and B_2 forming the co-cores of the attached 2-handles. The union of t_1 and t_2 forms a tangle t in B that is the co-cores of the 2-handles there. Suppose that the copies of $E(S)$ in B_1 and B_2 lie in $R_-(B_1)$ and $R_+(B_2)$, say. For each point of $t \cap E(S)$, pick an arc in $E(S)$ running from that point to $\partial E(S)$. We may arrange that these arcs α are disjoint, and that they all end on a suture of B_1 , say. These arcs α are arcs in $R_-(B_1)$. By adding in extra arcs disjoint from $E(S)$, extend them to a collection of arcs in $R_-(B_1)$ running from every point of $t \cap R_-(B_1)$ to the suture of B_1 . The arcs α also lie in $R_+(B_2)$, but they do not end on the suture there. Without changing their intersection with $E(S)$, extend them so that they do end on the suture of B_2 . Furthermore, add in extra arcs if necessary so that all point of $t \cap R_+(B_2)$ are at the start of such an arc. We may find a collection of discs D_1 in B_1 satisfying (i)-(iv) above such that D_1 intersects $R_-(B_1)$ in the given collection of arcs. We may find a similar collection of discs D_2 in B_2 . Then $D_1 \cup D_2$ is the required collection of discs in B satisfying (i)-(iv). Hence, (M, γ) is a pre-ball, as required. \square

3.13. SOME CONSEQUENCES OF ATOROIDALITY

Recall that a compact orientable 3-manifold M is *atoroidal* if any properly embedded incompressible torus is boundary parallel. This will be a useful hypothesis at various points in this paper. In this subsection, we collate a few consequences of atoroidality. In Section 11, we will show how the proof of the main theorems of this paper may be reduced to this case.

Lemma 3.10. *Let (M, γ) be a connected sutured manifold with no u-sutures. Let*

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_n, \gamma_n)$$

be an allowable hierarchy. Suppose that M is atoroidal and that its boundary is not a single torus. Then, for any $i > 1$, no component of M_i has boundary a single incompressible torus with no u-sutures.

Proof. Suppose that Y is such a component of M_i . Because Y contains no u-sutures, it is its own canonical enlargement. The decomposing surfaces $E(S_j)$ are all incompressible by Lemma 3.3. Therefore, the inclusion of Y into $E(M_1) = M$ is π_1 -injective. Therefore ∂Y is an incompressible torus in M . It is not boundary parallel, because ∂M would then be a single torus. This contradicts the atoroidality of M . \square

Lemma 3.11. *Let (M, γ) be a taut decorated sutured manifold. Let*

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1})$$

be an allowable hierarchy. Suppose that $E(M, \gamma)$ is atoroidal. Then $E(M_i, \gamma_i)$ is atoroidal for each i . Furthermore, if the only Seifert fibred components of $E(M, \gamma)$ are solid tori and copies of $T^2 \times I$, then the same is true for $E(M_i, \gamma_i)$.

Proof. Consider an incompressible torus T properly embedded in $E(M_i, \gamma_i)$. Since the inclusion of $E(M_i, \gamma_i)$ into $E(M, \gamma)$ is π_1 -injective, by Lemma 3.3, we deduce that T is incompressible in $E(M, \gamma)$. As $E(M, \gamma)$ is atoroidal, T is therefore boundary-parallel in $E(M, \gamma)$. Let $T \times I$ be the product region between T and a boundary component of $E(M, \gamma)$. Now the surfaces $E(S_j)$ are incompressible by Lemma 3.3, and so their intersection with $T \times I$ is boundary-parallel. Hence, the result of decomposing $T \times I$ along the surfaces $E(S_j)$ is to retain a copy of $T \times I$. We deduce that T is boundary-parallel in $E(M_i, \gamma_i)$. So, $E(M_i, \gamma_i)$ is atoroidal. Suppose now that some component of $E(M_i, \gamma_i)$ is Seifert fibred and neither a solid torus nor a copy of $T^2 \times I$. Its boundary tori must be boundary parallel in $E(M, \gamma)$, which implies that this component of $E(M, \gamma)$ is a Seifert fibred space which is neither a solid torus nor a copy of $T^2 \times I$. \square

Lemma 3.12. *Let (M, γ) be a connected taut decorated sutured manifold that admits an allowable hierarchy. Suppose that $E(M, \gamma)$ is atoroidal and not a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Let A_1 and A_2 be properly embedded incompressible annuli in M that intersect in a collection of essential simple closed curves and with non-trivial boundary disjoint from γ . Suppose that there is no homeomorphism $h: M \rightarrow M$ fixed on ∂M such that $h(A_1) \cap A_2$ consists of fewer simple closed curves. Then $A_1 \cap A_2$ consists of at most two curves.*

Proof. Suppose that $A_1 \cap A_2$ consists of at least two curves. Consider an annular component A'_2 of $A_2 - \text{int}(N(A_1))$ with neither boundary component in ∂M . Let A'_1 be the annulus in A_1 bounded by the two curves of $\partial A'_2$.

We first show that the two components of A'_2 emanate from the same side of A_1 . Suppose not. Then consider the torus $T' = A'_1 \cup A'_2$. A Dehn twist h around T' may be chosen so that $|h(A_1) \cap A_2| < |A_1 \cap A_2|$, contradicting our minimality assumption.

Again consider the torus $T' = A'_1 \cup A'_2$. This forms a boundary component of the manifold Y obtained from $E(M, \gamma)$ by cutting along A_1 , then along A'_2 . Since A_1 and A'_2 are incompressible, the inclusion of each component of Y into $E(M, \gamma)$ is π_1 -injective. Let Y' be the component of Y containing T' . So, if T' is incompressible in Y' , then T' is boundary-parallel in $E(M, \gamma)$. If the product region between T' and a component of $\partial E(M)$ contains $A_1 - A'_1$, then we deduce that A_1 is boundary parallel. In this case, the lemma is clear. So, we may assume that this product region Y' does not contain $A_1 - A'_1$. On the other hand, if T' is compressible in Y' , then Y' is a solid torus, because Y' is irreducible. Give this solid torus a Seifert fibration, in such a way that A'_1 and A'_2 are each a union of fibres. Then Y' contains an exceptional fibre. For otherwise, we could isotope A'_2 across Y' (together with any other components of $A_2 \cap Y'$) and thereby reduce the number of curves of $A_1 \cap A_2$.

Now consider an annulus A''_2 of $A_2 - \text{int}(N(A_1))$ adjacent to A'_2 , and also disjoint from ∂M . Such an annulus exists if $A_1 \cap A_2$ consists of at least three curves. Note that A''_2 emanates from the other side of A_1 from A'_2 . Let A''_1 be the annulus in A_1 bounded by $\partial A''_2$. It is possible that A'_1 and A''_1 intersect. Let T'' denote the torus $A''_1 \cup A''_2$. As argued above, we deduce that either T'' is boundary parallel in $E(M, \gamma)$ and that the product region Y'' has interior disjoint from A_1 , or T'' bounds a Seifert fibred solid torus Y'' with an exceptional fibre and with interior disjoint from A_1 .

Let Y''' be a regular neighbourhood of $Y' \cup Y''$. This is Seifert fibred, with planar base space, and where the sum of the number of exceptional fibres and boundary components is at least three. So, $\partial Y'''$ is incompressible in Y''' . All but one of its boundary components is boundary-parallel in $E(M, \gamma)$. The remaining component of $\partial Y'''$ separates off a subset of M with non-empty boundary. So, by the atoroidality of $E(M, \gamma)$, this component of $\partial Y'''$ is boundary-parallel in $E(M, \gamma)$. We deduce that $E(M, \gamma)$ is Seifert fibred, and not a solid torus or a copy of $T^2 \times I$, which is contrary to assumption. \square

A very similar argument, which we omit, also gives the following result.

Lemma 3.13. *Let (M, γ) be a connected taut decorated sutured manifold that admits an allowable hierarchy. Suppose that $E(M, \gamma)$ is atoroidal and not a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Let A be an annulus properly embedded in M with non-trivial boundary disjoint from γ , and let T be a properly embedded incompressible torus. Suppose that there is no homeomorphism $h: M \rightarrow M$ fixed on ∂M such that $h(T) \cap A$ consists of fewer simple closed curves than $T \cap A$. Then $T \cap A$ consists of at most one curve.*

4. SURFACES THAT EXTEND TO AN ALLOWABLE HIERARCHY

In this section, we will consider surfaces S that form the first surface in an allowable hierarchy. We will show that certain modifications can be made to S that preserve this property.

4.1. THE REDUCED LENGTH OF A HIERARCHY

Given a sequence of sutured manifold decompositions

$$(M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1})$$

its *reduced length* is the number of surfaces S_i that are not a union of product discs, annuli disjoint from the sutures and tori.

Lemma 4.1. *Let (M, γ) be a decorated sutured manifold that admits an allowable hierarchy with reduced length zero. Then each component of (M, γ) is either a product sutured manifold, which is not pre-spherical, or a copy of $T^2 \times I$ with no sutures, an orientable I -bundle over the Klein bottle with no sutures, the union of two such I -bundles glued along their boundary or a torus bundle over the circle.*

Proof. We may assume that M is connected. We may also assume that each surface in the allowable hierarchy for (M, γ) is connected. We will prove the lemma by induction on the length of this hierarchy. Let

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

be the first decomposition in the hierarchy. By induction, each component of (M', γ') is either a product sutured manifold which is not pre-spherical or a copy of $T^2 \times I$ with no sutures or an orientable I -bundle over the Klein bottle with no sutures.

Suppose first that S is a product disc or non-trivial product annulus. Then the components of (M', γ') must be products, because they have non-empty intersection with both $R_-(M')$ and $R_+(M')$. So, (M, γ) is also a product, as required.

Suppose that S is an annulus, disjoint from γ , with both boundary components in $R_-(M)$, say. Then a copy of S becomes a component of $R_+(M')$. Let X be the component of M' containing this copy of S . It has non-empty intersection with both $R_-(M')$ and $R_+(M')$, and so must be a product. It is therefore homeomorphic to $S \times I$. There are now two cases to consider.

Suppose that X does not contain the other copy of S . Then S is parallel to an annulus in $R_-(M)$. Hence, decomposing along S simply peels off a copy of $S \times I$. Therefore, each component of (M, γ) has the required form. On the other hand, if X contains both copies of S , then we deduce that S was non-separating in a component of M . This component of M is, up to homeomorphism, obtained from $S \times I$ by identifying $S \times \{0\}$ and $S \times \{1\}$. So, it is either a copy of $T^2 \times I$ or the orientable I -bundle over the Klein bottle. In both cases, it contains no sutures.

Finally suppose that S is a torus. Then each component of M' contains at least one toral boundary component disjoint from the sutures. Hence, inductively, each component of M' is an I -bundle over the Klein bottle or a copy of $T^2 \times I$. Therefore, M is a copy of $T^2 \times I$ with no sutures, an orientable I -bundle over the Klein bottle with no sutures, the union of two such I -bundles glued along their boundary or a torus bundle over the circle. \square

4.2. CONTROLLING TRIVIAL BOUNDARY CURVES

Let S be an allowable decomposing surface for a decorated sutured manifold (M, γ) . We say that a curve C of ∂S is *parallel towards a u-suture* if it is disjoint from γ and there is an annulus A in $R_{\pm}(M)$ such that one component of ∂A is C , the other component of ∂A is a u-suture and the orientations on A and S agree near C . Note that this annulus A is a trivialising planar surface for C .

Lemma 4.2. *Let (M, γ) be a taut decorated sutured manifold that admits an allowable hierarchy. Then it admits an allowable hierarchy in which each trivial boundary curve of each decomposing surface is parallel towards a u-suture. Moreover, we may ensure that each decomposing surface is incompressible. The new hierarchy has the same length as the original hierarchy, the same canonical extension and no greater reduced length. Moreover, if some surface in the original hierarchy is non-separating, then this remains true of the corresponding surface in the new hierarchy.*

Proof. Let

$$(M_i, \gamma_i) \xrightarrow{S_i} (M_{i+1}, \gamma_{i+1})$$

be a decomposition in the allowable hierarchy. Consider a trivial boundary curve C of S_i , with trivialising planar surface P . Suppose that C is not parallel towards a u-suture, and hence that P is not an annulus. We may assume that if there are any curves of $\text{int}(P) \cap S_i$, then these are parallel in P to u-sutures. Let S'_i be the result of attaching P to S_i , and isotoping a little so that it becomes properly embedded and so that each component of $\partial P - C$ is parallel to a u-suture. Decompose (M_i, γ_i) along S'_i instead, giving a new sutured manifold $(M'_{i+1}, \gamma'_{i+1})$. Then $(M'_{i+1}, \gamma'_{i+1})$ is obtained from (M_{i+1}, γ_{i+1}) by attaching a copy of $P \times [0, 1]$ to the u-suture coming from C . All later decomposing surfaces beyond (M_{i+1}, γ_{i+1}) avoid this u-suture. So we can view them as forming a hierarchy, starting at $(M'_{i+1}, \gamma'_{i+1})$ and ending at $(M'_{n+1}, \gamma'_{n+1})$, say. We note $(M'_{n+1}, \gamma'_{n+1})$ is also obtained from (M_{n+1}, γ_{n+1}) by attaching a copy of $P \times [0, 1]$ to the suture that is the copy of C . Hence, $(M'_{n+1}, \gamma'_{n+1})$ is also a product sutured manifold, no component of which is pre-spherical. Note also that, for each j , a curve in ∂M_j disjoint from the sutures is trivial if and only if the corresponding curve in $\partial M'_j$ is trivial. Furthermore, the canonical extensions $E(S_i)$ and $E(S'_i)$ are homeomorphic, via a homeomorphism that respects intersection with the sutures. So no component of $E(S'_i)$ is a disc disjoint from the sutures. Using these observations, it is easy to see that we have found an allowable hierarchy extending S'_i . This surface has fewer trivial boundary curves that are not parallel towards u-sutures. So, repeating this, we end with the required hierarchy. Note that these modifications to the hierarchy do not change its length. Moreover, if a component of a decomposing surface S_i was a product disc, torus or annulus disjoint from γ_i , then this is true of the corresponding component of the new surface S'_i . In the annular case, this follows from the fact that S_i has no component that is planar, disjoint from γ_i and with all but at most one boundary component trivial. So, the reduced length of the new hierarchy is at most that of the original one. Note also there is a one-one correspondence between components of M_i and M'_i , and so if a surface S_i is non-separating, then this remains true of S'_i .

In this way, we may arrange that each trivial boundary curve of each decomposing surface is parallel towards a u-suture. This implies that no component of ∂S_i bounds a disc in ∂M_i disjoint from the sutures. Therefore the hypotheses of Theorem 2.3 hold, and so each decomposing surface is taut and, in particular, incompressible. \square

4.3. SLICING UNDER A DISC OF CONTACT

Suppose that there is a disc D in $R_{\pm}(M)$ with $D \cap S = \partial D$, and with the orientation of D matching that of S near ∂D . Then D is a *disc of contact*. If we attach D to S , and then push it a little into the interior of M so that the resulting surface S' is properly embedded, this is known as *slicing under the disc of contact*.

Lemma 4.3. *Let*

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

be a taut allowable decomposition between decorated sutured manifolds that extends to an allowable

hierarchy. Let S' be obtained from S by slicing under a disc of contact. Then S' also extends to a allowable hierarchy, with reduced length at most that of the one starting with S , and with the same length.

Proof. Let D be the disc of contact. Then ∂D is a trivial boundary curve of S , and so it gives rise to a u-suture u of γ' . Therefore, all future decompositions in the given allowable hierarchy avoid this u-suture. The manifold (M'', γ'') obtained by decomposing along S' is obtained from (M', γ') by removing the u-suture u and attaching a 2-handle. All the decompositions in the hierarchy for (M', γ') avoid u , and so this sequence of decompositions may also be viewed as a hierarchy for (M'', γ'') . Boundary curves of surfaces in this new hierarchy are trivial if and only if the corresponding boundary curves are trivial in the original hierarchy. Hence, there is a one-one correspondence between the u-sutures in the new and original hierarchy. Therefore, all the decompositions in the new hierarchy are allowable. Let (M_{n+1}, γ_{n+1}) be the final manifold in the original hierarchy. The component containing u is a product that is not pre-spherical. The corresponding component in the new hierarchy is therefore also a product that is not pre-spherical. Hence, this forms an allowable hierarchy. The surfaces in the new hierarchy have the same topological type and the same intersection with the sutures as in the original hierarchy, except the first surface S' . Note that S could not have been a union of product discs, annuli disjoint from the sutures and tori. So, the reduced length of the new hierarchy is at most that of the original one. \square

4.4. DECOMPOSITION ALONG A PRODUCT DISC

Decomposition along an allowable product disc is a very common operation. In this section, we show that this preserves the existence of an allowable hierarchy.

Lemma 4.4. *Suppose that (M, γ) and (M', γ') are decorated sutured manifolds that differ by decomposition along an allowable product disc. Then (M, γ) admits an allowable hierarchy if and only if (M', γ') does. Moreover, we may arrange that these hierarchies have the same reduced length.*

Proof. Suppose that there is an allowable decomposition

$$(M, \gamma) \xrightarrow{P} (M', \gamma')$$

where P is a product disc. If (M', γ') admits an allowable hierarchy, then we may place P at the beginning of this hierarchy, and obtain an allowable hierarchy for (M, γ) with the same reduced length.

We now need to show that if (M, γ) admits an allowable hierarchy, then so does (M', γ') , with the same reduced length. We will also show that this hierarchy for (M', γ') can be taken to have the same length as the one for (M, γ) . We will prove this by induction. Our primary measure of complexity for this induction is the minimal reduced length of an allowable hierarchy for (M, γ) . Our secondary measure of complexity is the length of such a hierarchy.

The induction starts as follows. Suppose that (M, γ) has an allowable hierarchy with zero reduced length. Note that $R_+(M)$ and $R_-(M)$ are both non-empty because (M, γ) contains the product disc P . So, by Lemma 4.1, the component of (M, γ) containing P is a product. Moreover, no component is pre-spherical, because $E(M, \gamma)$ is taut by Lemma 3.3. Hence, the new components of (M', γ') are also products, and no component is pre-spherical, because each component created by decomposition along an allowable product disc contains at least one suture that is not a u-suture. So, (M', γ') admits an allowable hierarchy with zero reduced length, and the same length as the hierarchy for (M, γ) .

We now consider the inductive step. So, suppose that (M, γ) admits an allowable hierarchy

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1}).$$

We take this to have minimal reduced length, and subject to this condition, minimal length. By Lemma 4.2, we may assume that if any S_i has a trivial boundary curve, then this is parallel in $R_{\pm}(M_i)$ towards a u-suture. Let S_i be the first surface to intersect P . We may assume that no trivial boundary curve

intersects P . This is because the sutures that P runs over are not u-sutures, and so when we isotope trivial boundary curves of S_i towards u-sutures, we pull these curves away from P .

Suppose first that $P \cap S_i$ contains a simple closed curve that bounds a disc in P . Then, we may find such a curve that is innermost in P , and so bounds a disc D in P with interior disjoint from S_i . Now (M_{i+1}, γ_{i+1}) is taut, by Corollary 3.7, and so ∂D bounds a disc D' in $R_{\pm}(M_{i+1})$. Now, D' consists of parts of $R_{\pm}(M_i)$ and parts of S_i . Let C be the intersection between these two subsurfaces of D' , which is a collection of simple closed curves, because $\partial D'$ is disjoint from $R_{\pm}(M_i)$. We claim that C is empty. If there is such a curve, then an innermost one in D' bounds a disc, which is either a disc component of S_i or a disc in $R_{\pm}(M_i)$. In the former case, this violates our assumption that S_i is part of an allowable hierarchy. In the latter case, the boundary of this disc is a trivial boundary curve of S_i not parallel towards a u-suture, which is contrary to assumption. Hence, D' is a disc in S_i , and we may isotope this disc onto D , and thereby reduce the number of components of $P \cap S_i$. So, we may assume that $P \cap S_i$ contains no simple closed curves.

Suppose now $P \cap S_i$ contains an arc that has endpoints in the same component of $\partial P - \gamma$. Then there is one that is outermost in P , and that separates off a disc D disjoint from γ . Suppose that this disc D is disjoint from γ_{i+1} . Then ∂D bounds a disc D' in $R_{\pm}(M_{i+1})$. Let C be the intersection between $D' \cap R_{\pm}(M_i)$ and $D' \cap S_i$. As argued above, C contains no simple closed curves. Hence, it is a single arc. We may isotope the disc $D' \cap S_i$ onto D , and thereby remove this component of $P \cap S_i$. Suppose now that D is not disjoint from γ_{i+1} . This is therefore a product disc in (M_{i+1}, γ_{i+1}) . We may boundary compress S_i along this disc, giving a new surface \hat{S}_i . Let $(\hat{M}_{i+1}, \hat{\gamma}_{i+1})$ be the result of decomposing (M_i, γ_i) along \hat{S}_i . Then $(\hat{M}_{i+1}, \hat{\gamma}_{i+1})$ and (M_{i+1}, γ_{i+1}) differ by an allowable decomposition along a product disc. Now, the part of the hierarchy after (M_{i+1}, γ_{i+1}) has no greater reduced length and shorter length than the one for (M, γ) . So, by induction, $(\hat{M}_{i+1}, \hat{\gamma}_{i+1})$ extends to an allowable hierarchy with the same reduced length and the same length.

In this way, we may assume that no arc of $P \cap S_i$ has endpoints in the same component of $\partial P - \gamma$. Hence, $P \cap S_i$ consists of a collection of parallel arcs that divide P into a collection of product discs P' . Let (M'_i, γ'_i) be the result of decomposing (M', γ') along S_1, \dots, S_{i-1} . Note that there is a commutative diagram of sutured manifold decompositions

$$\begin{array}{ccc} (M_i, \gamma_i) & \xrightarrow{S_i} & (M_{i+1}, \gamma_{i+1}) \\ \downarrow P & & \downarrow P' \\ (M'_i, \gamma'_i) & \xrightarrow{S_i - \text{int}(N(P))} & (M'_{i+1}, \gamma'_{i+1}). \end{array}$$

The diagram also commutes as allowable decompositions of decorated sutured manifolds, for the following reasons. Decompositions along S_i and P are allowable by assumption. Neither creates any u-sutures that intersect a regular neighbourhood of P . In the case of S_i , this is because the trivial curves of ∂S_i miss P . Hence, the decompositions along P' and $S_i - \text{int}(N(P))$ are allowable. Neither decomposition creates any u-sutures that intersect a regular neighbourhood of P , because each boundary curve of P' and $S_i - \text{int}(N(P))$ that intersects this regular neighbourhood runs over a suture. Away from P , S_i and $S_i - \text{int}(N(P))$ have the same trivial curves. So, both ways around the commutative diagram create the same decorated sutured manifolds.

We need to verify that $S_i - \text{int}(N(P))$ can be part of an allowable hierarchy. Note that no component of $S_i - \text{int}(N(P))$ is planar, disjoint from the sutures and where all but one of its boundary components is trivial. This is because this was true of S_i , and when any component of S_i is cut along P , the resulting components of $S_i - \text{int}(N(P))$ each runs over a suture.

By induction, $(M'_{i+1}, \gamma'_{i+1})$ admits an allowable hierarchy with the same reduced length and the same length as the one starting with (M_{i+1}, γ_{i+1}) . Hence, (M', γ') admits an allowable hierarchy, by starting with S_1, \dots, S_{i-1} , then $S_i - \text{int}(N(P))$, and then the allowable hierarchy for $(M'_{i+1}, \gamma'_{i+1})$. This clearly has the same length as the given hierarchy for (M, γ) . We claim that this has the same reduced length as

this hierarchy for (M, γ) . We must check that each component of S_i is a product disc, torus or annulus disjoint from the sutures if and only if the same is true for each component of $S_i - \text{int}(N(P))$. Note that each arc of $S_i \cap P$ intersects both $R_-(M_i)$ and $R_+(M_i)$. Hence if it lies in an annulus of S_i disjoint from the sutures, then this is a product annulus. So, if S_i is a union of product discs, annuli disjoint from the sutures and tori, then the same is true of $S_i - \text{int}(N(P))$. Conversely, no component of $S_i - \text{int}(N(P))$ that intersects $N(P)$ can be an annulus disjoint from the sutures or a torus, and if each such component of $S_i - \text{int}(N(P))$ is a product disc, then these patch together to form product discs and product annuli of S_i . This proves the claim. Hence, the given hierarchy for (M', γ') does indeed have the same reduced length as the original one for (M, γ) . \square

4.5. TUBING ALONG AN ARC

Suppose that α is an arc in $R_\pm(M)$ with $\alpha \cap S = \partial\alpha$. Then there is an embedding of $\alpha \times [-1, 1]$ in $R_\pm(M)$ with $\alpha \times \{0\} = \alpha$ and $(\alpha \times [-1, 1]) \cap S = \partial\alpha \times [-1, 1]$. Suppose that the orientation that $\alpha \times [-1, 1]$ inherits from $R_\pm(M)$ agrees with the orientation of S near $\partial\alpha \times [-1, 1]$. Then α is called a *tubing arc*.

We can construct a new surface S' as follows. Attach $\alpha \times [-1, 1]$ to S , and push the surface a little into the interior of M , so that it becomes properly embedded. Then S' is obtained from S by *tubing along the arc* α . (See Figure 5.)

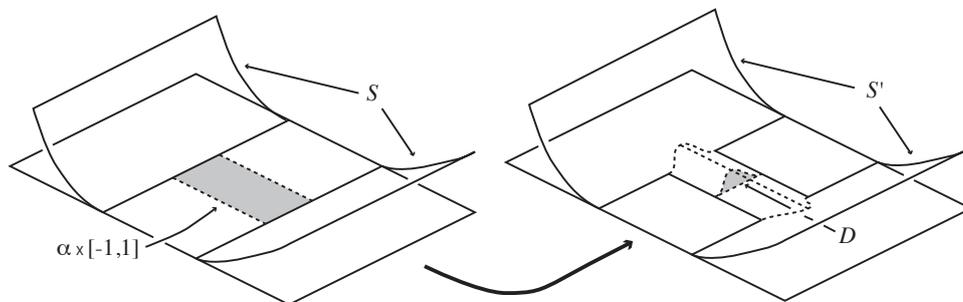


Figure 5: Tubing along an arc

Note that S is obtained from S' by boundary-compressing along a product disc D . This product disc has boundary consisting of an arc in $\alpha \times [-1, 1]$ running from $\alpha \times \{-1\}$ to $\alpha \times \{1\}$, and an arc in S' .

This type of modification played a key role in [20]. So it will be important to understand how it behaves in the context of allowable hierarchies.

Lemma 4.5. *Let S be an allowable decomposing surface for a decorated sutured manifold (M, γ) that extends to an allowable hierarchy. Let S' be obtained from S by tubing along an arc α . Then S' also extends to an allowable hierarchy. Moreover, the reduced length of the hierarchy beyond S' is equal to the reduced length of the hierarchy beyond S .*

Proof. We first show that decomposition along S' is allowable. Since S was disjoint from any u-sutures and so was α , the same is true of S' . Suppose a component C' of $\partial S'$ bounds a trivialising planar surface P' , but where the transverse orientations of P' and S' disagree near C' . Since decomposition along S was allowable, this component C' must run along the new tube. Consider a component C of ∂S that is incident to an endpoint of α . Then C has non-empty intersection with P' . It therefore bounds a planar subsurface P of P' . This forms a trivialising planar surface for C and its transverse orientation disagrees with that of S near C . This contradicts the assumption that S was allowable.

We now check that S' can be part of an allowable hierarchy, by verifying that no component of S' can be a planar surface disjoint from γ with all but at most one boundary curve being trivial. Suppose that F' is such a component. This must contain the new tube, because otherwise it forms a component of S . Thus F' boundary-compresses to form either one or two components of S . Suppose first that we

get two components F_1 and F_2 of S . Then the arc of intersection between the boundary-compression disc and F' must have had endpoints on the same component of $\partial F'$. If $\partial F'$ contains a non-trivial component, it is either this curve or a boundary component of F_1 or F_2 . In each case, at least one of F_1 or F_2 has at most one trivial boundary curve. This contradicts the fact that S was part of an allowable hierarchy. Suppose now that we obtain a single component F of S by boundary-compressing the tube. Then the arc of intersection between the boundary-compression disc and F' had endpoints on distinct components of $\partial F'$. If they were both trivial, then their trivialising planar surfaces combine to form a trivialising planar surface for the boundary component of F at the endpoints of α . Hence, all but at most one boundary component of F is trivial, and again this is a contradiction. On the other hand, if one of the components of $\partial F'$ incident to the boundary-compression disc is non-trivial, then the remaining components of $\partial F'$ are trivial. These end up forming all but one boundary component of F . Again, this is a contradiction.

Thus, we have shown that S' may be part of an allowable hierarchy. We now show that such an allowable hierarchy exists.

Let (M_S, γ_S) and $(M_{S'}, \gamma_{S'})$ be the result of decomposing (M, γ) along S and S' respectively. Let C_1 and C_2 be the curves of ∂S at the endpoints of α . It is possible that $C_1 = C_2$. For each curve C_i , let A_i be the product annulus properly embedded in $(M_{S'}, \gamma_{S'})$ that runs parallel to it. Orient A_i so that its transverse orientation near ∂M is the same as that of ∂S .

We claim that A_i is trivial in $(M_{S'}, \gamma_{S'})$ if and only if α has endpoints on the same trivial curve of ∂S . Suppose first that α has endpoints on the same trivial curve of ∂S . Then $A_1 = A_2$. Removing a regular neighbourhood of α from the trivialising planar surface gives trivialising planar surfaces for the two new components of $\partial S'$. These therefore give rise to u-sutures of $\gamma_{S'}$. The two u-sutures and $A_1 \cap \partial M$ (which equals $A_2 \cap \partial M$) together bound a pair of pants in $\partial M_{S'}$, which forms a trivialising planar surface, and hence A_1 (which equals A_2) is trivial.

Suppose now that some A_i is trivial. Then one of its boundary curves is trivial. It cannot be $A_i \cap S$, because it would then lie in a component of S that is planar, disjoint from the sutures and has all but at most one boundary curve trivial. This contradicts the assumption that S is part of an allowable hierarchy. Hence, $A_i \cap \partial M$ must be the trivial curve. Its trivialising planar surface P must have an orientation that is consistent with that of A_i , as otherwise S is not allowable. Therefore P contains the new components of $\partial S'$ in its boundary. These are therefore u-sutures of $\gamma_{S'}$ and hence trivial boundary curves of $\partial S'$. They must be distinct curves of $\partial S'$, and hence α has endpoints on the same component of ∂S , as otherwise the other component of ∂S would bound a trivialising planar surface with the wrong transverse orientation. The component of ∂S at the endpoints of α is trivial, with trivialising planar surface formed from the union of P and the trivialising planar surfaces for the two components of $\partial S'$. This proves the claim.

The lemma now divides into the cases of the claim. Suppose first that some A_i is trivial. Then α has endpoints on the same trivial curve of ∂S . This curve becomes a u-suture in (M_S, γ_S) , and later surfaces in the hierarchy avoid it. Call this hierarchy

$$(M_S, \gamma_S) = (M_2, \gamma_2) \xrightarrow{S_2} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1}).$$

We have a (non-allowable) decomposition

$$(M_{S'}, \gamma_{S'}) \xrightarrow{D} (M_S, \gamma_S).$$

Hence, we can view M_S as lying in $M_{S'}$. The hierarchy for (M_S, γ_S) gives a sequence of decompositions

$$(M_{S'}, \gamma_{S'}) = (M'_2, \gamma'_2) \xrightarrow{S'_2} \dots \xrightarrow{S'_n} (M'_{n+1}, \gamma'_{n+1}).$$

A simple induction gives that

- (i) each (M'_i, γ'_i) is decorated;

- (ii) apart from the three sutures that intersect a regular neighbourhood of α , the u-sutures of γ_i and γ'_i are equal;
- (iii) a curve in $\partial M_i - \gamma_i$ is trivial if and only if the corresponding curve in $\partial M'_i - \gamma'_i$ is trivial;
- (iv) decomposition along S'_i is allowable;
- (v) no component of S'_i is planar, disjoint from the sutures and where all but at most one boundary component is trivial.

The argument is very similar to the proof of Lemma 3.8 and is omitted. There is a (non-allowable) decomposition

$$(M'_{n+1}, \gamma'_{n+1}) \xrightarrow{D} (M_{n+1}, \gamma_{n+1}).$$

Hence, $(M'_{n+1}, \gamma'_{n+1})$ is also a product sutured manifold, no component of which is pre-spherical. Thus, we have found the required allowable hierarchy for $(M_{S'}, \gamma_{S'})$.

The remaining case is where no A_i is trivial. Then we may decompose $(M_{S'}, \gamma_{S'})$ along these annuli. The resulting decorated sutured manifold is then a copy of (M_S, γ_S) plus a product sutured manifold that is not pre-spherical. Therefore, we may follow this decomposition with the given hierarchy for (M_S, γ_S) . \square

We close this section with a technical lemma that will be useful later.

Lemma 4.6. *Let*

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

be a taut allowable decomposition between decorated sutured manifolds. Let S' be obtained from S by tubing along an arc α . Let C be a trivial curve in ∂M_S disjoint from γ_S and from the discs in S to which the tube is attached. Then C corresponds to a curve trivial C' in $\partial M_{S'}$ disjoint from $\gamma_{S'}$.

Proof. Suppose first that one of the curves of ∂S to which the tube is attached is trivial. It then bounds a trivialising planar surface P . This must contain α . The curve of ∂S at the other endpoint of α cannot lie in the interior of P , because it would then bound a trivialising planar surface with the wrong transverse orientation. We therefore deduce that both endpoints of α lie in the same trivial curve of ∂S . Therefore, this gives rise to two trivial curves of $\partial S'$. Hence, in this case, $E(M_S, \gamma_S)$ is homeomorphic to $E(M_{S'}, \gamma_{S'})$, and the homeomorphism takes C to C' . Thus, C is trivial if and only if C' is trivial.

Suppose now that neither of the curves of ∂S to which the tube is attached is trivial. Then the trivialising planar surface P for C is disjoint from sutures of γ_S corresponding to these curves. Hence, it corresponds to a trivialising planar surface P' for C' . \square

4.6. BOUNDARY COMPRESSING ALONG A PRODUCT DISC

Let S be a surface properly embedded in (M, γ) . Suppose that there is a disc D embedded in M , such that $D \cap S$ is a single arc in ∂D , and $D \cap \partial M$ is a single arc in ∂D and where these two arcs intersect at their endpoints. Suppose $D \cap \gamma$ is empty and the orientations of S and ∂M disagree at ∂D . We will consider the surface S' obtained from S by boundary compressing along D . Thus, S' is obtained from S by tubing along an arc. Let (M_S, γ_S) and $(M_{S'}, \gamma_{S'})$ be the manifolds obtained by decomposing (M, γ) along S and S' respectively. Note that D is a product disc in (M_S, γ_S) , and that decomposing (M_S, γ_S) along D gives a sutured manifold homeomorphic to $(M_{S'}, \gamma_{S'})$.

Proposition 4.7. *Let (M, γ) be a decorated sutured manifold, and let S be a decomposing surface that extends to an allowable hierarchy. Let S' be obtained from S by boundary compressing along a product disc disjoint from γ . Let S'' be obtained from S' by removing any component that is a planar surface disjoint from γ with all but at most one boundary curve that is trivial. Then S'' also extends to an allowable hierarchy. Moreover, the reduced length of this hierarchy beyond S'' is equal to the reduced length beyond S .*

Proof. Let (M_S, γ_S) and $(M_{S'}, \gamma_{S'})$ be the result of decomposing (M, γ) along S and S' respectively. There is a decomposition

$$(M_S, \gamma_S) \xrightarrow{D} (M_{S'}, \gamma_{S'})$$

where D is the product disc. Thus, one might hope to use Lemma 4.4 to prove the proposition. However, there are a number of potential complications. Firstly, S' might not be an allowable decomposing surface, because it might have a trivial boundary curve, where the trivialising planar surface and S' are incompatibly oriented. In this case, we have not even specified which sutures of $\gamma_{S'}$ are to be viewed as u-sutures. Secondly, S' might have a planar component disjoint from γ and with all but at most one boundary component trivial, and hence it might not be possible for S' to be part of an allowable hierarchy. Thirdly, the product disc D may run over a u-suture of γ_S , and so decomposition along D might not be allowable, and therefore Lemma 4.4 might not apply.

Case 1. D does not run over a u-suture of γ_S .

In this case, Lemma 4.4 is applicable. Therefore, we declare that the sutures of $\gamma_{S'}$ incident to D are not u-sutures. Then by Lemma 4.4, $(M_{S'}, \gamma_{S'})$ with this decoration admits an allowable hierarchy with the same reduced length as the one starting from (M_S, γ_S) . However, this does not complete the proof of the lemma, because the first or second problems mentioned above may still hold.

Suppose that S' has a trivial boundary curve C' , where the trivialising planar surface P has transverse orientation differing from that of S' near C' . The planar surface P may contain other boundary curves of S' . These are all trivial, and by choosing C' appropriately, we may assume that the trivialising subsurfaces for the curves $S' \cap \text{int}(P)$ are all correctly oriented. Let P' be the subsurface of P obtained by cutting along the curves $S' \cap \text{int}(P)$ and then taking the component containing C' . The curve C' becomes a suture of $\gamma_{S'}$. The surface P' becomes a planar component of $R_{\pm}(M_{S'})$, and all its boundary curves apart from C' are u-sutures. Hence, it extends to a disc in $E(M_{S'}, \gamma_{S'})$. Now, $E(M_{S'}, \gamma_{S'})$ is taut, by Lemma 3.3, and so the component of $E(M_{S'}, \gamma_{S'})$ is a taut 3-ball. Therefore, the component of S' containing C' is planar, disjoint from the sutures and all but at most one of its boundary curves are trivial. We therefore focus on this case.

Let S'_1 be a planar component of S' , disjoint from γ and with all but at most one boundary component trivial. Then $E(S'_1)$ is a disc disjoint from the sutures of $E(M, \gamma)$. By Lemma 3.3, this is taut. Hence, $E(S'_1)$ is boundary parallel in $E(M, \gamma)$. It therefore separates off a 3-ball component of $E(M_{S'}, \gamma_{S'})$. By Lemma 3.9, the corresponding component of $(M_{S'}, \gamma_{S'})$ is a product sutured manifold. Therefore, $(M_{S'}, \gamma_{S'})$ is the disjoint union of this product sutured manifold and the manifold obtained by decomposing (M, γ) along $S - S'_1$. So, $S - S'_1$ extends to an allowable hierarchy, with the same reduced length. Thus, the proposition is proved in this case.

Case 2. D runs over a u-suture of γ_S .

Case 2A. D runs over two distinct sutures of γ_S .

Call these C_1 and C_2 , where C_1 is trivial. It is also possible that C_2 is trivial. Let C' be the curve of $\partial S'$ incident to D . Then C' is trivial if and only if C_2 is trivial. Furthermore, the trivialising planar surfaces lie on the same side. Thus, C' cannot bound a trivialising planar surface with the wrong orientation, because this would imply that C_2 did also, contradicting the assumption that S was allowable. Hence, decomposition along S' is allowable. Moreover, boundary-compressing S along D does not disconnect this component of S . Hence, if the resulting component of S' is a planar surface disjoint from the sutures and with all but at most one boundary component trivial, then the same is true of the component of S , which is a contradiction. So, S' can be part of an allowable hierarchy.

Now S is obtained from S' by tubing along an arc that starts and ends in C' . We now create a new surface S'' that is obtained from S' by tubing along a different arc. This has the same endpoints but runs along a regular neighbourhood of C_1 . Decomposing along S'' gives a decorated sutured manifold homeomorphic to (M_S, γ_S) . Hence, S'' extends to an allowable hierarchy with the same reduced length as the one starting with S . But S'' has a disc of contact which lies within a regular neighbourhood of C'_1 .

Slicing under this disc of contact gives a surface isotopic (relative to γ) to S' . By Lemma 4.3, S' extends to an allowable hierarchy with the same reduced length as the one starting with S .

Case 2B. D runs over the same u-suture of γ_S twice.

This comes from a boundary component C of ∂S . In this case, the boundary compression gives rise to two distinct boundary components C'_1 and C'_2 of S' . Suppose first that neither C'_1 nor C'_2 is trivial.

Let us change the decoration of (M_S, γ_S) by declaring that the suture of γ_S incident to D is *not* a u-suture. Then, by Lemma 3.8, (M_S, γ_S) with this new decoration still admits an allowable hierarchy. Note that the decomposition

$$(M_S, \gamma_S) \xrightarrow{D} (M_{S'}, \gamma_{S'})$$

is now allowable, because D now does not run over any u-sutures. Moreover, $(M_{S'}, \gamma_{S'})$ has the correct decoration because the curves C'_1 and C'_2 are non-trivial. Hence, Lemma 4.4 establishes the proposition in this case.

Now consider the case where C'_1 or C'_2 is trivial. In fact, if one of them is trivial, then so is the other. Moreover, C'_1 (say) bounds a trivialising planar surface P'_1 with the incorrect orientation, whereas C'_2 bounds a trivialising planar surface P'_2 with the correct orientation. We again change the decoration of (M_S, γ_S) by declaring that the sutures of γ_S incident to D are not u-sutures. Then $(M_{S'}, \gamma_{S'})$ inherits a decoration where neither of the sutures coming from C'_1 or C'_2 is a u-suture. Again, Lemma 4.4 implies that $(M_{S'}, \gamma_{S'})$ admits an allowable hierarchy. Therefore, by Lemma 3.3, $E(M_{S'}, \gamma_{S'})$ is taut. Now, P'_1 extends to a disc in the boundary of $E(M_{S'}, \gamma_{S'})$, and so by the tautness of $E(M_{S'}, \gamma_{S'})$, the component of $E(M_{S'}, \gamma_{S'})$ is a product sutured 3-ball. Hence, as argued in Case 1, the component S'_1 of S' incident to P'_1 is a planar surface, disjoint from γ and that is boundary-parallel via a product region that is a product sutured manifold. Therefore, we may remove this component S'_1 from S' . Let S'' be the union of the remaining components. Let $(M_{S''}, \gamma_{S''})$ be obtained by decomposing (M, γ) along S'' . Note that this decomposition is allowable. Note also that this sutured manifold differs from (M_S, γ_S) by attaching the sutured manifold $S'_1 \times I$ along a product disc. Hence, the allowable sutured manifold hierarchy for (M_S, γ_S) gives an allowable sutured manifold hierarchy for $(M_{S''}, \gamma_{S''})$ with the same reduced length. This is the required hierarchy. \square

4.7. ANNULAR SWAPS AND DECOMPOSITIONS ALONG ANNULI

Let S be a surface properly embedded in M . Suppose that there is an annulus or Möbius band A embedded in the interior of M such that $A \cap S = \partial A$. In the case where A is an annulus, suppose that the orientations on S near $S \cap A$ either both point towards A or both point away from A . Let $N(A)$ be the orientable I -bundle over A , where A is viewed as a section of $N(A)$. Then $N(A)$ embeds in M in such a way that $N(A) \cap S$ is the I -bundle over ∂A . Let S' be the surface that results from removing $N(A) \cap S$ from S and attaching $\text{cl}(\partial N(A) - S)$. Then our assumption about the orientation on S near $S \cap A$ implies that S' can be oriented in such a way that the orientations on S and S' agree on their intersection. We say that S' is obtained from S by an *annular swap* along A .

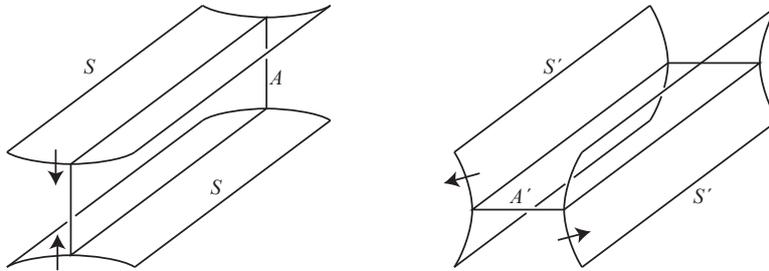


Figure 6: An annular swap

Note that there is an annulus or Möbius band A' embedded in $N(A)$, which is the I -bundle over a core curve of A . So $A' \cap S'$ equals $\partial A'$. Then S is obtained from S' by an annular swap along A' . (See Figure 6.) This modification to a decomposing surface will be important in this paper. In particular, it arises in the context of ‘irregular switches’ in normal surface theory. More details can be found in Sections 7.5 and 7.3.

The following result implies that this modification may be made, while preserving the existence of an allowable hierarchy, under some fairly mild assumptions.

Proposition 4.8. *Let*

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

be a taut allowable decomposition between decorated sutured manifolds that extends to an allowable hierarchy. Suppose that $E(M, \gamma)$ is atoroidal and that no component of $E(M, \gamma)$ is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Suppose also no component of M_S has boundary a single torus with no sutures. Assume that if $E(S)$ contains a component that is a torus or annulus disjoint from γ and from the attached 2-handles of $E(M, \gamma)$, then no other component of $E(S)$ is of this form. Suppose that no component of $A \cap S$ bounds a disc in the canonical extension $E(S)$. Let S' be a surface obtained from S by a swap along an annulus or Möbius band A in M . Orient S' so that the orientations of S' and S agree on their intersection. Then at least one of the following holds:

- (i) S' extends to an allowable hierarchy, or
- (ii) A is an annulus and the two copies of A in S' lie in distinct components of S' , one of these components S'_1 separates off a solid torus with no sutures, and $S' - S'_1$ extends to an allowable hierarchy.

In both cases, the new hierarchy has the same reduced length as the original one. Moreover, if some component of S was not an annulus or torus disjoint from γ , then this remains true of at least one component of the first surface of the new hierarchy.

At the same time, we will prove the following.

Proposition 4.9. *Let (M, γ) be a taut decorated sutured manifold that admits an allowable hierarchy. Suppose that $E(M, \gamma)$ is atoroidal and that no component of $E(M, \gamma)$ is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Let A be an annulus properly embedded in M with boundary disjoint from γ . Suppose that neither curve of ∂A is trivial. Suppose that decomposition of (M, γ) along A does not create a component that is a solid torus disjoint from the sutures. Then A extends to an allowable hierarchy with the same reduced length as the original one.*

Proof. We will prove these propositions simultaneously, by induction on the reduced length of the hierarchy.

We first prove that Proposition 4.9 in the case where the allowable hierarchy for (M, γ) has some reduced length n implies Proposition 4.8 in the case where (M_S, γ_S) has an allowable hierarchy with reduced length n .

So consider the situation of Proposition 4.9. Let $2A$ denote $\text{cl}(\partial N(A) - S)$, the (∂I) -bundle over A . View $2A$ as one or two annuli embedded in (M_S, γ_S) . Orient $2A$ so that the region it bounds containing A inherits all the new sutures. We start to verify that the hypotheses of Proposition 4.9 apply to the components of $2A$ in (M_S, γ_S) . Note that no boundary curve of $2A$ is trivial in (M_S, γ_S) . For if a boundary curve of $2A$ did bound a disc in the boundary of $E(M_S, \gamma_S)$ disjoint from the sutures, then this would imply that either a component of $A \cap S$ bounded a disc in $E(S)$, contrary to assumption, or some component of $E(S)$ is a disc disjoint from γ , contradicting the fact that S is part of an allowable hierarchy, or some boundary curve of $E(S)$ is trivial in $E(M, \gamma)$, contradicting Lemma 3.2. The same argument gives that when $2A$ is disconnected and (M_S, γ_S) is decomposed along one component of $2A$, then the other component of $2A$ remains non-trivial in the resulting manifold. Note also that $E(M_S, \gamma_S)$ is atoroidal, by Lemma 3.11. Also, no component of $E(M_S, \gamma_S)$ is a Seifert fibre space, other than a solid torus or a copy of $T^2 \times I$. Thus, we may assume that all of the hypotheses of Proposition 4.9 apply to the components of $2A$

in (M_S, γ_S) , with the possible exception that this decomposition may create some solid torus components with no sutures. The proof divides into two cases, according to whether this hypothesis also holds.

Case 1. Decomposing (M_S, γ_S) along $2A$ does not create any solid torus components with no sutures.

Then by Proposition 4.9 applied once or twice, decomposing (M_S, γ_S) along $2A$ extends to an allowable hierarchy with reduced length n . But decomposing (M, γ) along S then $2A$ gives the same sutured manifold as decomposing (M, γ) along S' then $2A'$, where $2A' = S - \text{int}(N(S'))$. Both the decomposition along S' and the decomposition along $2A'$ are allowable, and the resulting sutured manifold has the same u-sutures as the manifold obtained by decomposing (M_S, γ_S) along $2A$. This is because the boundary curves of $2A'$ are non-trivial, for the following reason. Suppose that a boundary curve of $2A'$ had a trivialising planar surface P . The intersection $P \cap \partial S$ would give rise to a (possibly empty) collection of u-sutures in (M_S, γ_S) . The component of $P - \text{int}(N(\partial S))$ incident to ∂A would then form a trivialising planar surface for this component of ∂A , which is contrary to hypothesis. Also, no component of $E(S')$ is a disc, because this would imply that $E(S)$ had a disc component disjoint from the sutures or some curve of ∂A bounded a disc in $E(S)$, neither of which is possible. Hence, S' extends to an allowable hierarchy the same reduced length as the one starting with S , as required.

Case 2. Decomposing (M_S, γ_S) along $2A$ does create at least one solid torus component with no sutures.

Case 2A. The same solid torus with no sutures lies on both sides of $2A$.

Let V be the component of (M_S, γ_S) containing this solid torus. Then ∂V consists of tori with no sutures. It cannot be a single torus, by assumption. So it contains two tori. Consider the intersection between $E(S)$ and the boundary of $E(V) = V$. Since $E(S)$ has no trivial boundary curves, this intersection consists of some tori and some essential annuli. Since V has two boundary tori and each has non-empty intersection with $E(S)$, we deduce that $E(S)$ contains at least two components that are tori or annuli disjoint from γ and from the attached 2-handles. This contradicts one of our assumptions.

Case 2B. On both sides of $2A$, there are distinct solid tori disjoint from the sutures.

These solid tori then patch together to give a component of M_S . This has boundary a single torus, that is disjoint from γ_S , contradicting one of our assumptions.

Case 2C. On exactly one side of $2A$, there is a solid torus disjoint from the sutures.

Let Y be this solid torus. If $S \cap Y$ contains any trivial curves, then one of these bounds a disc of contact in M . We may slice under this, and maintain the existence of an allowable hierarchy, with the same reduced length, by Lemma 4.3. In this process, the solid torus Y with no sutures is preserved. So, in this way, we ensure that $S \cap Y$ is essential in ∂Y . We note that the (∂I) -bundle over A must consist of two copies of A in S' lying in distinct components of S' , since one lies in Y and the other does not. Let S'_1 denote the components $Y \cap S'$. We note that these are tori and annuli disjoint from the sutures, because they are essential subsurfaces of ∂Y . By our assumption about annular and toral components of $E(S)$, S'_1 must be connected. We will show that decomposing along $S' - S'_1$ extends an allowable hierarchy with the same reduced length as the one starting with S . Orient A so that when (M_S, γ_S) is decomposed along it, Y inherits some sutures. Then by Proposition 4.9, this decomposition extends to an allowable hierarchy with the same reduced length as the one for (M_S, γ_S) . But decomposing (M, γ) along S then A gives the same manifold as decomposing along $S' - S'_1$ then along the annuli $S'_1 \cap S$. Hence, $S' - S'_1$ does extend to the required allowable hierarchy. This establishes that Proposition 4.8 in the case where (M_S, γ_S) has an allowable hierarchy with reduced length n follows from Proposition 4.9 in the case where (M, γ) has an allowable hierarchy with reduced length n .

So, we will focus on Proposition 4.9. Therefore consider a taut decorated sutured manifold (M, γ) , where $E(M, \gamma)$ is atoroidal and where no component of $E(M, \gamma)$ is Seifert fibred other than a solid torus or a copy of $T^2 \times I$. By hypothesis, there is an allowable hierarchy

$$(M, \gamma) = (M_1, \gamma_1) \xrightarrow{S_1} \dots \xrightarrow{S_n} (M_{n+1}, \gamma_{n+1}).$$

We will make several hypotheses about this hierarchy, the first of which is:

- (1) The reduced length of this allowable hierarchy is as short as possible.

By Lemma 4.2, we may assume our second hypothesis:

- (2) If any S_j has a trivial boundary curve, then this is parallel in $R_{\pm}(M_j)$ towards a u-suture.

We note that, in general, we cannot assume that the surfaces in the hierarchy are connected, because to break a surface down into its components may increase the reduced length of the hierarchy. However, decomposing along tori and annuli disjoint from the sutures separately is possible:

- (3) If any surface S_j contains a component that is a torus or an annulus disjoint from the sutures, then this is all of S_j .

Our final hypothesis is:

- (4) Assuming (1) - (3), the length of the hierarchy is as short as possible.

We will prove the proposition by induction. We will consider only allowable hierarchies satisfying (1) - (4). Our primary measure of complexity for such a hierarchy will be its reduced length. Our secondary measure of complexity will be the length of the hierarchy.

The induction starts with the case where the hierarchy has reduced length zero. Let

$$(M, \gamma) \xrightarrow{A} (M', \gamma')$$

be the given decomposition. Then by Lemma 4.1, the component of (M, γ) containing A is a product sutured manifold, a copy of $T^2 \times I$ with no sutures or an orientable I -bundle over the Klein bottle with no sutures. Note that the other possibilities are ruled out because A has non-empty boundary. Since A is non-trivial, it is incompressible in M . So, in all cases, A is either a union of I -fibres or is boundary parallel. In both cases, (M', γ') has an allowable hierarchy with reduced length zero, as required.

Before we embark on the main flow of the argument, we deal with the case where the component of M containing A is Seifert fibred. We first show that this component of M has no u-sutures. For if it has a u-suture, then the annular components of $R_{\pm}(M)$ that are adjacent to it extend to discs or spheres in the boundary of $E(M, \gamma)$. But $E(M, \gamma)$ is taut, by Lemma 3.3, and so this component of $E(M, \gamma)$ is a 3-ball with a single suture. We deduce that this component of M had a single torus boundary component, containing two sutures, exactly one of which is a u-suture. But in this case, each component of ∂A is parallel to the u-suture and hence is trivial, which is contrary to assumption. So, this component of M does indeed have no u-sutures, and so this component of (M, γ) is its own canonical extension. By our assumption about Seifert fibred components of $E(M, \gamma)$, we deduce that it is a solid torus or a copy of $T^2 \times I$. It is easy to check that in this case, the manifold obtained by decomposing this component of (M, γ) along A has an admissible hierarchy with reduced length at most one. Moreover, it admits an admissible hierarchy with reduced length zero if and only if the same is true of (M, γ) . Thus, Proposition 4.9 is proved in this case.

We now consider the general inductive step. We will consider various possibilities for $S_1 \cap A$.

Case 1. S_1 is disjoint from A .

Then we have a commutative diagram

$$\begin{array}{ccc} (M_1, \gamma_1) & \xrightarrow{S_1} & (M_2, \gamma_2) \\ \downarrow A & & \downarrow A \\ (M'_1, \gamma'_1) & \xrightarrow{S_1} & (M'_2, \gamma'_2). \end{array}$$

Note that (M_2, γ_2) admits an allowable hierarchy satisfying (1)-(4), with smaller complexity than the given hierarchy for (M_1, γ_1) . Hence, we may apply induction to the decomposition of (M_2, γ_2) along A . However,

we need to check that the hypotheses of the proposition hold. Note that $E(M_2, \gamma_2)$ is atoroidal by Lemma 3.11, and no component of $E(M_2, \gamma_2)$ is Seifert fibred, unless it is a solid torus or a copy of $T^2 \times I$. Note also that A has non-trivial boundary in M_2 . If (M'_2, γ'_2) has no solid toral components disjoint from the sutures, then inductively, (M'_2, γ'_2) admits an allowable hierarchy with the same reduced length as the one for (M_2, γ_2) starting with S_2 . Hence, in this case, the proposition is proved.

So, suppose that there is a solid torus component Y of M'_2 disjoint from γ'_2 . Then Y has non-empty intersection with S_1 because otherwise decomposing (M_1, γ_1) along A would also create Y , contrary to a hypothesis of the proposition. So, S_1 is an annulus disjoint from the sutures, by hypotheses (2) and (3). Give Y a Seifert fibration so that A and S_1 are a union of fibres. This must have an exceptional fibre. For otherwise, we can isotope S_1 across this solid torus. It is then easy to see that decomposing (M_2, γ_2) along A does not then create a solid torus with no sutures. For if it did, this solid torus would have to meet both A and S_1 , and we could then deduce that M is Seifert fibred, and we have dealt with this case already. Hence, in this case, the proposition is proved by induction, as above. So, Y has an exceptional fibre.

Note that S_1 is not parallel to an annulus in (M_1, γ_1) disjoint from the sutures. For if it were, then decomposition along S_1 just creates a copy of (M_1, γ_1) and a product sutured manifold. In this case, we could remove S_1 from the start of the hierarchy, contradicting our assumption (4). Similarly, we may assume that A is not parallel to an annulus in (M_1, γ_1) disjoint from the sutures, because in this case, Proposition 4.9 is trivial.

Let $\overline{S_1}$ be a parallel copy of S_1 but with opposite orientation, and arranged so that it does not lie in Y . Define \overline{A} similarly. Then decomposing (M_2, γ_2) along $\overline{S_1} \cup \overline{A}$ does not create any solid tori disjoint from the sutures, because this would imply either that S_1 or A is parallel to an annulus in $R_{\pm}(M_1)$, or that $E(M, \gamma)$ is toroidal, or that $E(M, \gamma)$ has a Seifert fibred component other than a solid torus or a copy of $T^2 \times I$. No boundary curve of $\overline{S_1} \cup \overline{A}$ is trivial in (M_2, γ_2) . So, inductively, decomposing (M_2, γ_2) along $\overline{S_1} \cup \overline{A}$ extends to an allowable hierarchy with the required reduced length. In other words, decomposing (M_1, γ_1) along $\overline{S_1} \cup S_1 \cup \overline{A}$ extends to such an allowable hierarchy. Let (M''_2, γ''_2) be the result of making this decomposition. This is homeomorphic to the manifold (M'''_2, γ'''_2) obtained from (M_1, γ_1) by decomposing along $\overline{S_1} \cup A \cup \overline{A}$. So, (M'''_2, γ'''_2) also admits an allowable hierarchy with the required reduced length. But then we see that A does indeed extend to the hierarchy that we are looking for.

Thus, we may assume that $S_1 \cap A$ is non-empty. Exactly as argued in the proof of Lemma 4.4, we may ensure that $A \cap S_1$ contains no simple closed curves or arcs that are inessential in A .

Case 2. $A \cap S_1$ is a collection of essential simple closed curves in A , and S_1 is neither an annulus disjoint from γ_1 nor a torus.

Note first that no component of $E(S_1)$ is a torus or annulus disjoint from γ_1 and from the attached 2-handles. This is because, in this situation, S_1 would either contain a trivial curve bounding a disc of contact, contradicting hypothesis (2) or would itself be a torus or annulus disjoint from γ_1 and from the attached 2-handles, which we are assuming is not the case here.

We give the curves $A \cap S_1$ the transverse orientation coming from S_1 . We also transversely orient ∂A by using the transverse orientation of $R_{\pm}(M_1)$.

Suppose that two curves of $A \cap S_1$ are adjacent in A and incoherently oriented. They then bound an annulus A_1 such that $A_1 \cap S_1 = \partial A_1$. Let S'_1 be the result of performing an annular swap along A_1 . We now check that the conditions of Proposition 4.8 hold in this case. Note that $E(M, \gamma)$ is atoroidal, and that no component of $E(M, \gamma)$ is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Suppose that a component of M_2 has boundary a single torus with no sutures. At least one component of S_1 would be an essential subsurface of ∂M_2 , and therefore S_1 would have a component that is a disc, annulus or torus. The former case is impossible since S_1 intersects A in essential curves. So S_1 contains a component that is an annulus or torus disjoint from the sutures. But we assuming that this is not the case here. We verified above that no component of $E(S_1)$ is a torus or annulus disjoint from the sutures and the attached 2-handles. Neither curve of $A_1 \cap S_1$ bounds a disc in $E(S_1)$. This is because we could extend this

to form a compression disc for a curve of ∂A . Since $E(M_1, \gamma_1)$ is taut, this curve would bound a disc in $\partial E(M_1)$ disjoint from the sutures, and so A would be trivial, which is contrary to hypothesis. Note that, because S_1 is neither an annulus disjoint from γ_1 nor a torus, the reduced length of the hierarchy beyond (M_2, γ_2) is strictly less than the one for (M_1, γ_1) . Therefore, by induction, we deduce that, possibly after removing one of its components, S'_1 extends to an allowable hierarchy with the same reduced length as the one starting with S_1 . Note that S'_1 has fewer curves of intersection with A .

Suppose now all the curves of $A \cap S_1$ are coherently oriented, but that this orientation disagrees with at least one component of ∂A . Then we may find such a curve that is outermost in A , which therefore separates off an annulus A' , and so that the transverse orientations on $\partial A'$ either both point towards A' or both point away from it. We now add an annular component A_+ to S_1 . This annulus A_+ is boundary parallel in M_1 , and runs along a regular neighbourhood of $\partial A' \cap \partial M_1$. We orient A_+ so that the product region between it and ∂M_1 inherits two sutures. Clearly, $S_1 \cup A_+$ extends to an allowable hierarchy with the same reduced length as the original one. This is because by cutting along $S_1 \cup A_+$ rather than S_1 , the resulting sutured 3-manifold inherits a new component, which is a taut product solid torus. By adding A_+ , we have introduced a new curve of intersection with A . But we may then perform an annular swap to remove this curve and also the other curve of $\partial A'$. Let S'_1 be the result of performing this annular swap. It is easy to check that the conditions of Proposition 4.8 are satisfied. Note in particular that $E(S_1)$ contains no components that are annuli or tori disjoint from the sutures and the attached 2-handles, and so only one component of $E(S_1 \cup A_+)$ has this form. Hence, possibly after removing a component, S'_1 extends to an allowable hierarchy with the same reduced length.

In this way, we may replace the first surface S_1 of the allowable hierarchy with a surface S'_1 , such that $A \cap S'_1$ is a (possibly empty) collection of coherently oriented essential curves, and so that these transverse orientations agree with both components of ∂A . In doing so, we have not increased the reduced length of the hierarchy. We note that at least one component of S'_1 is not an annulus disjoint from γ_1 , a torus or product disc. Let S''_1 be the result of removing from S'_1 the components that are tori and annuli disjoint from the sutures. Then S''_1 also extends to an allowable hierarchy with the same reduced length. The surface S''_1 satisfies (2) and (3). We may further ensure that the later surfaces also satisfy (2) and (3), by Lemma 4.2. The reduced length of this new hierarchy is at most that of the original one. We do not make any claim about the length of the new hierarchy, however. Note that $A' = A - \text{int}(N(S'_1))$ is a collection of product annuli. We have a commutative diagram

$$\begin{array}{ccc} (M_1, \gamma_1) & \xrightarrow{S'_1} & (M''_2, \gamma''_2) \\ \downarrow A & & \downarrow A' \\ (M'_1, \gamma'_1) & \xrightarrow{S''_1 - \text{int}(N(A))} & (M'_2, \gamma'_2). \end{array}$$

The decomposition along A' does not create any solid tori with no sutures, because A' consists of product annuli. No boundary curve of A' is trivial. Note also that $E(M''_2, \gamma''_2)$ is atoroidal by Lemma 3.11, and no component is a Seifert fibre space, other than a solid torus or a copy of $T^2 \times I$. The reduced length of the hierarchy for (M''_2, γ''_2) is strictly less than that for (M_1, γ_1) . Hence, inductively, (M'_2, γ'_2) extends to an allowable hierarchy with the same reduced length as that for (M''_2, γ''_2) . So, (M'_1, γ'_1) has the required allowable hierarchy. It has the same reduced length as the original one.

Case 3. $A \cap S_1$ is a collection of essential simple closed curves in A , and S_1 is a torus.

By Lemma 3.13, we may assume that $A \cap S_1$ is a single curve. Let $A' = A - \text{int}(N(S_1))$. We have a commutative diagram of allowable decompositions

$$\begin{array}{ccc} (M_1, \gamma_1) & \xrightarrow{S_1} & (M_2, \gamma_2) \\ \downarrow A & & \downarrow A' \\ (M'_1, \gamma'_1) & \xrightarrow{S_1 - \text{int}(N(A))} & (M'_2, \gamma'_2). \end{array}$$

It is easy to check that the hypotheses of Proposition 4.9 apply to A' . In particular, decomposition along A' does not create a solid torus component disjoint from the sutures. So, inductively, (M'_2, γ'_2) extends to an allowable hierarchy with the same reduced length as the one starting with (M_2, γ_2) . Therefore, (M'_1, γ'_1) admits an allowable hierarchy with the required reduced length.

Case 4. $A \cap S_1$ is a collection of essential simple closed curves in A , and S_1 is an annulus disjoint from γ_1 .

We may assume that $A \cap S_1$ consists of as few curves as possible. By Lemma 3.12, this number is therefore one or two. Now, (M_2, γ_2) admits a hierarchy satisfying all the requirements (1) - (4). Its reduced length is at most that of the original hierarchy, and it has shorter length. So, we may attempt to apply induction to the decomposition of (M_2, γ_2) along $A - \text{int}(N(S_1))$. However, this decomposition may create a solid torus disjoint from the sutures. Suppose that it does not. Then we have a commutative diagram

$$\begin{array}{ccc} (M_1, \gamma_1) & \xrightarrow{S_1} & (M_2, \gamma_2) \\ \downarrow A & & \downarrow A - \text{int}(N(S_1)) \\ (M'_1, \gamma'_1) & \xrightarrow{S_1 - \text{int}(N(A))} & (M'_2, \gamma'_2). \end{array}$$

By induction, (M'_2, γ'_2) extends to a hierarchy with reduced length equal to the one starting with (M_2, γ_2) . Hence, we obtain a hierarchy for (M'_1, γ'_1) , starting with $S_1 - \text{int}(N(A))$, then the hierarchy for (M'_2, γ'_2) . Its reduced length is equal to the one for (M, γ) .

So, suppose that decomposing (M_2, γ_2) along $A - \text{int}(N(S_1))$ does create a solid torus Y disjoint from the sutures. Consider first the case where $A \cap S_1$ is a single curve. The solid torus Y has boundary consisting of a component of $A - \text{int}(N(S_1))$, one component of $S_1 - \text{int}(N(A))$, and an annulus in ∂M . We give Y a Seifert fibration with at most one exceptional fibre, so that $A \cap S_1$ is a fibre. In fact, Y must contain an exceptional fibre, because otherwise, we may isotope S_1 off A , and we have dealt with this case already. Hence, by the atoroidality of $E(M, \gamma)$ and the assumption that no component of $E(M, \gamma)$ is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$, there is just one component of $M - \text{int}(N(A \cup S_1))$ that is a solid torus disjoint from the sutures. It is incident to just one of the four sides of $A \cap S_1$, because of the orientations on S_1 and A .

Consider the annuli A' shown in Figure 7, and oriented as shown. Note that decomposing (M_2, γ_2) along A' does not create a solid torus with no sutures. For if it did, this solid torus could not support an exceptional fibre in its Seifert fibration, and we could therefore isotope S_1 off A . Hence, by induction, this decomposition extends to an allowable hierarchy with the same reduced length as the one starting with (M_2, γ_2) . But instead of decomposing along S_1 then A' , we may start with (M_1, γ_1) , decompose along A , and then A' . This gives the same decorated sutured manifold, up to homeomorphism. Hence, A does extend to an allowable hierarchy, with the required reduced length, as claimed.

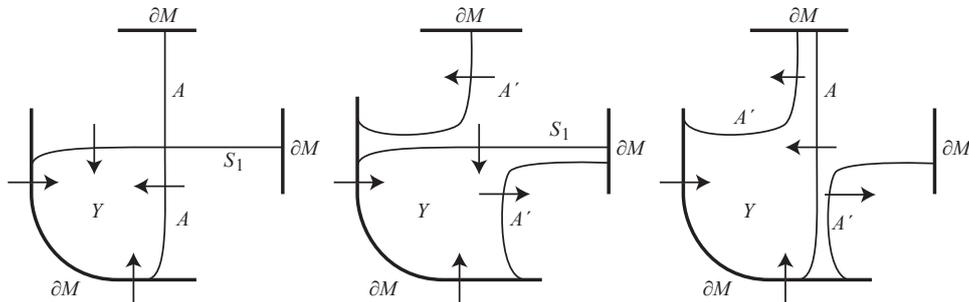


Figure 7: When $A \cap S_1$ is a single curve

Now suppose that $A \cap S_1$ consists of two simple closed curves. Consider the annulus of $A - \text{int}(N(S_1))$ disjoint from ∂M and the annulus of $S_1 - \text{int}(N(A))$ disjoint from ∂M . Their union is a torus T . We claim that T bounds a solid torus with interior disjoint from $S_1 \cup A$. Let Y be the component of $E(M, \gamma) -$

$\text{int}(N(S_1 \cup A))$ with T in its boundary. The inclusion of Y into $E(M, \gamma)$ is π_1 -injective. So, if Y is not a solid torus, then T is boundary parallel in $E(M, \gamma)$. But it is then clear that the component of $E(M, \gamma) - \text{int}(N(S_1 \cup A))$ that is a solid torus with no sutures cannot be Seifert fibred with an exceptional fibre. Hence, in this case, there is an isotopy reducing $|S_1 \cap A|$, and we have dealt with this already. So, Y is indeed a solid torus. It must be Seifert fibred with an exceptional fibre, for otherwise, we can isotope S_1 off A . No other component of $E(M, \gamma) - \text{int}(N(S_1 \cup A))$ can be an exceptionally fibred solid torus, because this would imply that $E(M, \gamma)$ itself was Seifert fibred and not a solid torus or a copy of $T^2 \times I$, by atoroidality. Therefore, Y is the solid torus component of $M - \text{int}(N(S_1 \cup A))$ with no sutures, which are assuming to exist. Therefore, the orientations on S_1 and A near Y are as shown in Figure 8 (up to reversing all of the transverse orientations). The three different rows in that figure illustrate the possible transverse orientations on ∂M near ∂A .

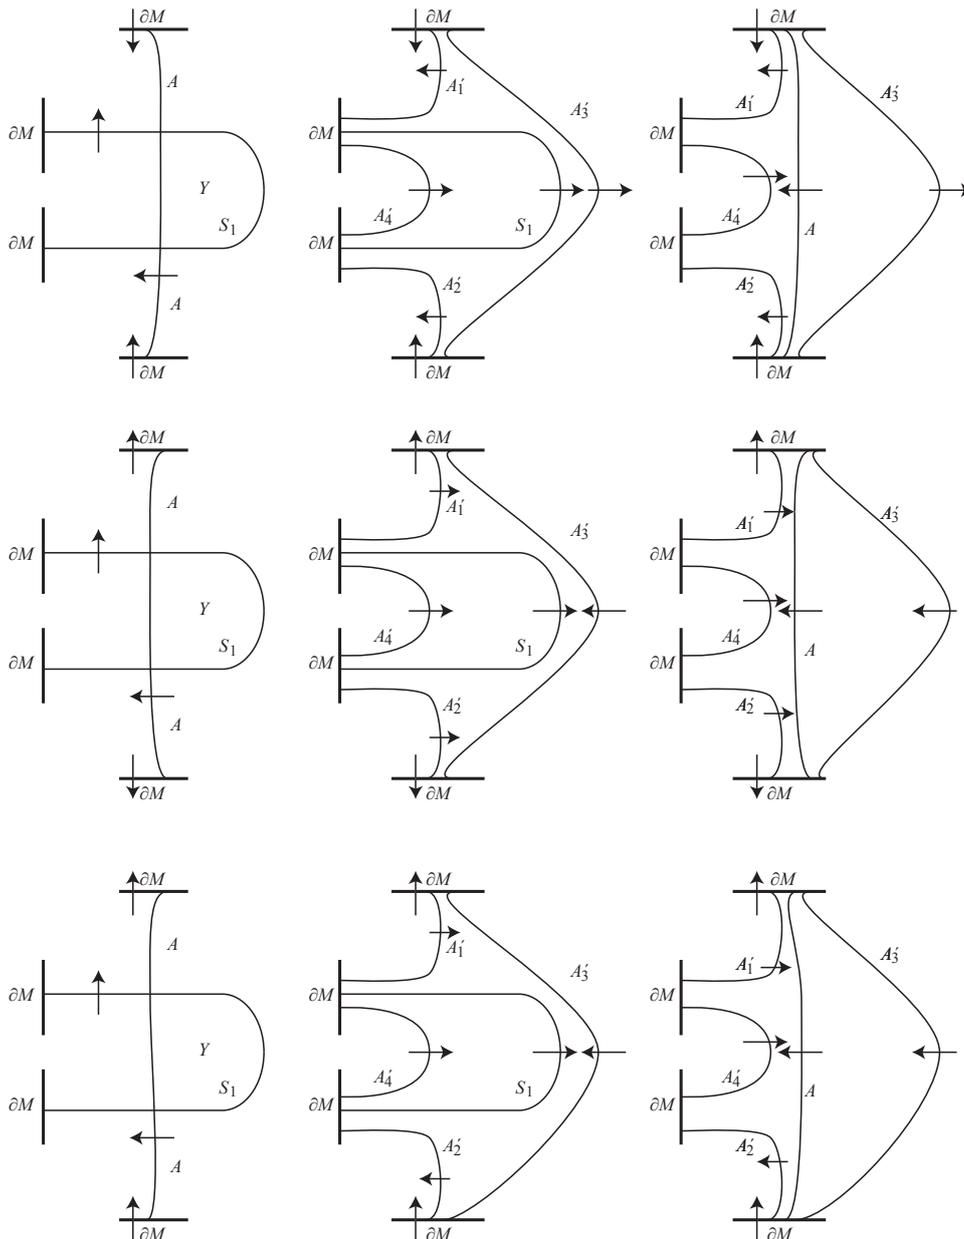


Figure 8: When $A \cap S_1$ is two curves

In the figure are also shown various annuli A'_1, A'_2, A'_3 and A'_4 . Note that none of A'_1, A'_2 or A'_4 separate off a solid torus with no sutures, because in this case, either we could isotope S_1 to reduce $|S_1 \cap A|$ or decomposition along S_1 would create a solid torus with no sutures. However, A'_3 might separate off a solid torus with no sutures, in which case it is boundary parallel. Let A' denote the union of A'_1, A'_2, A'_4 and possibly also A'_3 . Include A'_3 if and only if it does not separate off a solid torus with no sutures. Then decomposing (M_2, γ_2) along A' does not create any solid tori with no sutures. Hence, inductively, this extends to an allowable hierarchy with same reduced length as the one starting with S_2 . Let (M'_2, γ'_2) be the manifold that results from decomposing (M_2, γ_2) along A' . Similarly, let (M''_2, γ''_2) be the manifold that results from decomposing (M_1, γ_1) along $A' \cup A$. Then, aside from two components in each, (M'_2, γ'_2) and (M''_2, γ''_2) are equal. Furthermore, it is easy to check that they have admissible hierarchies with the same reduced length. Hence, A' extends to the required admissible hierarchy, with the required reduced length.

Case 5. A is not a product annulus, $A \cap S_1$ consists of essential arcs in A , and these arcs are not all coherently oriented.

Then we may find a component of $A - \text{int}(N(S_1))$ that is a disc D disjoint from γ_2 . Using the fact that (M_2, γ_2) is taut, ∂D bounds a disc D' in $R_\pm(M_2)$. Let C be the intersection between $S_1 \cap D'$ and $R_\pm(M_1) \cap D'$. As argued in the proof of Lemma 4.4, C is a collection of arcs. In this case there are two of them. There are two cases. If the region in D' between these two arcs lies in S_1 , then we may isotope this part of S_1 onto D , and so remove two curves of $A \cap S_1$. On the other hand, if the region in D' lies in $R_\pm(M_1)$, then we deduce that A is boundary-compressible in M_1 along a boundary compression disc disjoint from γ_1 . Therefore, A is parallel to an annulus in $R_\pm(M_1)$. Hence, decomposing (M_1, γ_1) creates a solid torus with no sutures, or a solid torus with a product sutured manifold structure. The former case is excluded by hypothesis, and in the latter case, the proposition holds trivially.

Case 6. A is a product annulus, or $A \cap S_1$ consists of arcs that are essential in A and coherently oriented.

Hence, $A - \text{int}(N(S_1))$ consists of a collection of product discs P embedded in (M_2, γ_2) . Note that we have a commutative diagram

$$\begin{array}{ccc} (M_1, \gamma_1) & \xrightarrow{S_1} & (M_2, \gamma_2) \\ \downarrow A & & \downarrow P \\ (M'_1, \gamma'_1) & \xrightarrow{S_1 - \text{int}(N(A))} & (M'_2, \gamma'_2). \end{array}$$

By Lemma 4.4, (M'_2, γ'_2) admits an allowable hierarchy. But we may also obtain (M'_2, γ'_2) by first decomposing (M, γ) along A , then along $S_1 - \text{int}(N(A))$. Each of these surfaces meets the conditions of being part of an allowable hierarchy, except possibly $S_1 - \text{int}(N(A))$. Its canonical extension may contain some discs disjoint from the sutures. Suppose that there is such a disc. Then it is not boundary parallel in M'_1 , because this would imply that A is not a product annulus and that not all the arcs of $A \cap S_1$ are coherently oriented. Therefore we deduce that $R_\pm(E(M'_1))$ is compressible. We deduce that at least one of the hypotheses of Theorem 2.3 does not apply to the decomposition

$$E(M'_1, \gamma'_1) \xrightarrow{E(S_1 - \text{int}(N(A)))} E(M'_2, \gamma'_2)$$

because $E(M'_2, \gamma'_2)$ is taut, whereas $E(M'_1, \gamma'_1)$ is not. We see that the only possibility is that $E(M'_1, \gamma'_1)$ contains a component Y that is a solid torus disjoint from the sutures, and which is compressed by some disc components of $E(S_1 - \text{int}(N(A)))$. But then (M'_1, γ'_1) contains a solid torus disjoint from the sutures, and we are assuming that this is not the case.

We have therefore found the required hierarchy extending A . \square

4.8. PRODUCT-SEPARATING SURFACES

Let $(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$ be an allowable decomposition between decorated sutured manifolds. A connected surface S is known as *product-separating* if it separates the component of M that contains

it, and one of the components of $E(M_S, \gamma_S)$ obtained by decomposing the corresponding component of $E(M, \gamma)$ along $E(S)$ is a product sutured manifold $F \times I$, such that

- (i) F is homeomorphic to $E(S)$;
- (ii) any attached 2-handles in $F \times I$ are of the form $D^2 \times I$, for a disc D^2 in the interior of F ;
- (iii) each such 2-handle has non-empty intersection with the copy of $E(S)$ in the boundary of $F \times I$.

We say that a 2-handle satisfying (ii) and (iii) is *strongly vertical* in $F \times I$. A disconnected surface S is *product-separating* if each of its components is.

Note that this does not imply that S is boundary parallel, even when S is connected. For example, let F be a connected surface properly embedded in M , that is parallel to an essential subsurface of $R_{\pm}(M)$ and that has no trivial boundary curves. Orient F so that the region between it and the subsurface of $R_{\pm}(M)$ is a product sutured manifold. Now pick an embedded arc such that $\alpha \cap F$ is an endpoint of α lying in the interior of F , and such that $\alpha \cap \partial M$ is the other endpoint of α . Choose α so that it is disjoint from the product region. Also arrange that at one endpoint, F points towards α and at the other endpoint, ∂M points away from α . Now remove a regular neighbourhood of $\alpha \cap F$ from F , and attach the annulus that encircles α . Let S be the new surface. Then $E(S)$ is homeomorphic to F , and one component of $E(M_S, \gamma_S)$ is a copy of $F \times I$. This product manifold contains none of the attached 2-handles of $E(M_S, \gamma_S)$. So, S is product-separating. However, S need not be boundary-parallel. This example shows that $|\chi(S)|$ may be greater than $|\chi(F)|$, where $F \times I$ is the product sutured manifold.

A related example is where α is chosen to lie within the product sutured manifold between F and the subsurface of $R_{\pm}(M)$. Again, the resulting surface S is product-separating. However, if α is knotted, then S need not be boundary-parallel.

A third example again starts with F as above. We now isotope F , keeping its boundary in ∂M , so that it approaches a component of γ that is not a u-suture. We then slide ∂F over this component of γ so that afterwards, a sub-arc of ∂F and a sub-arc of γ together bound a disc in $R_{\pm}(M)$. The resulting surface S continues to have non-trivial boundary. Hence $E(S) = S$. Moreover, $E(S)$ separates off a product sutured manifold $F \times I$, where F is homeomorphic to $E(S)$ and that is disjoint from the 2-handles of $E(M_S, \gamma_S)$. So S is product-separating. In this case, S is boundary-parallel. Note, however, that in this case S is not the entirety of a component of $R_{\pm}(F \times I)$.

The following lemma asserts that the examples given above essentially exhaust the possibilities for a product-separating surface.

Lemma 4.10. *Let S be a connected product-separating surface in the sutured manifold (M, γ) . Then, after slicing under discs of contact, S becomes a surface \bar{S} that is parallel to a subsurface G in ∂M . Moreover, one of $G \cap R_-(M)$ or $G \cap R_+(M)$ is a connected subsurface of G that is homeomorphic to G . The remainder of G is a collection of discs, each of which intersects ∂G in a single arc, and annuli, each of which intersects ∂G in a single closed curve. Moreover, decomposing (M, γ) along \bar{S} creates a copy of (M, γ) and a component that is a product sutured manifold homeomorphic to $\bar{S} \times I$.*

Proof. By assumption, S is product-separating. Modify S by slicing under discs of contact as many times as possible. This does not change $E(S)$ or $E(M_S, \gamma_S)$. Thus, the new surface \bar{S} is still product-separating. So, one component of $E(M_S, \gamma_S)$ is a product sutured manifold $F \times I$. By (i) in the definition of product-separating, F is homeomorphic to $E(S)$. Now, a copy of $E(S)$ is a subsurface of $R_{\pm}(F \times I)$, $R_+(F \times I)$ say. But by (i), $R_+(F \times I)$ is homeomorphic to $E(S)$. Furthermore, $\text{cl}(R_+(F \times I) - E(S))$ can contain no discs disjoint from γ_S because such a disc would form a disc of contact for \bar{S} , which we have removed. Therefore, $\text{cl}(R_+(F \times I) - E(S))$ consists of discs that intersect γ_S in a single arc and annuli that intersect γ_S in a simple closed curve. So, $E(S)$ is boundary parallel in $E(M, \gamma)$. Let G' be the subsurface of $\partial E(M, \gamma)$ to which it is parallel. The surface $\text{cl}(R_+(F \times I) - E(S))$ is disjoint from the attached 2-handles, because if a 2-handle did intersect $\text{cl}(R_+(F \times I) - E(S))$, it would be disjoint from $E(S)$, contradicting (iii) in the definition of product-separating. Hence, each component of $\text{cl}(R_+(F \times I) - E(S))$ is a copy of a component

of $\text{cl}(G' - R_-(E(M, \gamma)))$. In the product region between $E(S)$ and G' , any attached 2-handles are strongly vertical. Thus, when we form \bar{S} by removing from $E(S)$ any trivialising planar surfaces for trivial boundary curves of \bar{S} , we see that \bar{S} is parallel to a subsurface G of ∂M with the required properties. \square

The following result asserts that we may avoid product-separating surfaces in an allowable hierarchy.

Proposition 4.11. *Let (M, γ) be a decorated sutured manifold that admits an allowable hierarchy. Then (M, γ) admits an allowable hierarchy where each decomposing surface S is connected and not product-separating.*

Proof. We may arrange that each decomposing surface is connected, because we may replace a decomposition along a disconnected surface by a sequence of decompositions along its components. Consider a shortest allowable hierarchy for (M, γ) where each decomposing surface is connected. Then we claim that the first surface S cannot be product-separating. This will imply that every surface in the hierarchy is not product-separating, by induction. Suppose S is product-separating. Then we modify the hierarchy to a shorter one, as follows. First arrange, using Lemma 4.2, for each trivial boundary curve of each decomposing surface to be parallel towards to a u-suture. This does change the fact that the first surface (which we will still call S) is product-separating. Since S now has no discs of contact, Lemma 4.10 implies that S is now boundary-parallel. Moreover, decomposing (M, γ) along S gives a copy of (M, γ) and a product sutured manifold. So, we may discard S , and discard the product sutured manifold. The later surfaces in the hierarchy therefore form an admissible hierarchy for (M, γ) with shorter length. \square

5. SUTURED MANIFOLDS AND HANDLE STRUCTURES

5.1. HANDLES STRUCTURES AND THEIR COMPLEXITY

In this paper, handle structures on 3-manifolds will play a central role. If \mathcal{H} is a handle structure, then \mathcal{H}^i denotes the union of the i -handles. We will always insist that our handle structures satisfy the following requirements:

- (i) The intersection between each i -handle $D^i \times D^{3-i}$ and $\bigcup_{j < i} \mathcal{H}^j$ is $\partial D^i \times D^{3-i}$.
- (ii) Any two i -handles are disjoint.
- (iii) The intersection between any 2-handle $D^2 \times D^1$ and any 1-handle $D^1 \times D^2$ is $\alpha \times D^1$ in the 2-handle, where α is a collection of disjoint arcs in ∂D^2 , and of the form $D^1 \times \beta$ in the 1-handle, where β is a collection of disjoint arcs in ∂D^2 .
- (iv) Every 2-handle runs over at least one 1-handle.

When (M, γ) is a sutured manifold, we will also insist that a handle structure \mathcal{H} for M satisfies the following conditions:

- (i) The intersection between γ and any 0-handle is a (possibly empty) collection of arcs and simple closed curves.
- (ii) The intersection between γ and any 1-handle $D^1 \times D^2$ is a (possibly empty) collection of arcs, each of the form $D^1 \times \{*\}$, where $*$ is a point in ∂D^2 .
- (iii) The sutures γ are disjoint from the 2-handles and 3-handles.

In a handle structure \mathcal{H} , the surface $\mathcal{F} = \mathcal{H}^0 \cap (\mathcal{H}^1 \cup \mathcal{H}^2)$ will be important. We view it as a subsurface of $\partial \mathcal{H}^0$. It inherits a handle structure, where each 0-handle is a component of $\mathcal{H}^0 \cap \mathcal{H}^1$, and each 1-handle is a component of $\mathcal{H}^0 \cap \mathcal{H}^2$. We denote the union of these 0-handles of \mathcal{F} by \mathcal{F}^0 , and the union of the 1-handles of \mathcal{F} by \mathcal{F}^1 .

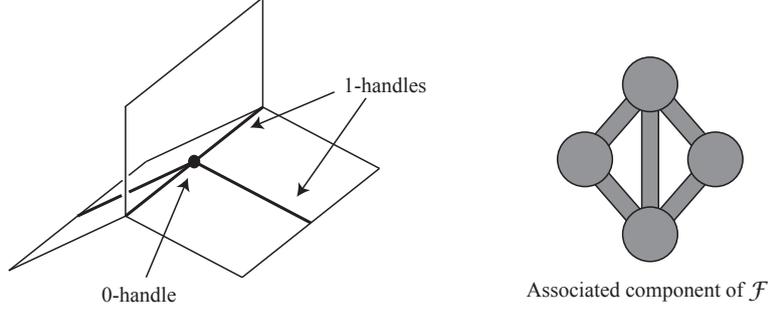


Figure 9: Handles near a 0-handle

In [20], the author defined a measure of complexity for a handle structure of a sutured manifold. This will form a foundation for this paper. We recall it now.

For a component F of \mathcal{F} , its *index* $I(F)$ is

$$I(F) = -2\chi(F) + |F \cap \gamma|.$$

Three integers were defined:

$$\begin{aligned} C_1(F) &= |F \cap \mathcal{F}^1| + 1, \\ C_2(F) &= I(F), \\ C_3(F) &= |\partial F|. \end{aligned}$$

The \mathcal{F} -complexity set $C_{\mathcal{F}}(\mathcal{H})$ is the collection of ordered triples

$$\{(C_1(F), C_2(F), C_3(F)) : F \text{ is a component of } \mathcal{F} \text{ with } I(F) > 0\},$$

where repetitions are retained.

We order the triples within an \mathcal{F} -complexity set, as follows. We declare that

$$(C_1(F), C_2(F), C_3(F)) \leq (C_1(F'), C_2(F'), C_3(F'))$$

if and only if

- (i) $C_1(F) \leq C_1(F')$, or
- (ii) $C_1(F) = C_1(F')$ and $C_2(F) \leq C_2(F')$, or
- (iii) $C_1(F) = C_1(F')$ and $C_2(F) = C_2(F')$ and $C_3(F) \leq C_3(F')$.

We use this to place a total ordering on \mathcal{F} -complexity sets, as follows. Given two \mathcal{F} -complexity sets $C_{\mathcal{F}}(\mathcal{H})$ and $C_{\mathcal{F}}(\mathcal{H}')$, for handle structures \mathcal{H} and \mathcal{H}' , we first consider the largest triple in each set. If one is greater than the other, then we declare that the \mathcal{F} -complexity set containing it is greater. If the two triples are equal, then we pass to second-largest triples in each set, and compare these, and so on. If at some stage, we run out of triples in one of the sets, then the other \mathcal{F} -complexity set is declared to be the greater one. This allows us to compare any two \mathcal{F} -complexity sets. It is clear that this is in fact a well-ordering on \mathcal{F} -complexity sets (see Lemma 5.3 in [20]).

We can use this ordering on \mathcal{F} -complexity sets to define a notion of complexity for a handle structure \mathcal{H} , as follows. The *complexity* $C(\mathcal{H})$ of a handle structure \mathcal{H} is the ordered pair $(C_{\mathcal{F}}(\mathcal{H}), n(\mathcal{H}))$, where $n(\mathcal{H})$ is the number of 0-handles of \mathcal{H} containing a component of \mathcal{F} with positive index. We compare the complexity of handle structures \mathcal{H} and \mathcal{H}' by asserting that $C(\mathcal{H}) > C(\mathcal{H}')$ if and only if

- (i) $C_{\mathcal{F}}(\mathcal{H}) > C_{\mathcal{F}}(\mathcal{H}')$, or
- (ii) $C_{\mathcal{F}}(\mathcal{H}) = C_{\mathcal{F}}(\mathcal{H}')$ and $n(\mathcal{H}) < n(\mathcal{H}')$.

It is shown in Lemma 5.3 in [20] that this ordering on complexity of handle structures is a well-ordering.

The index $I(F)$ of a component F of \mathcal{F} may be computed as follows. If F_0 is a 0-handle of \mathcal{F} , then its *index* is defined to be

$$I(F_0) = -2 + |F_0 \cap \mathcal{F}^1| + |F_0 \cap \gamma|.$$

Then it is easy to check that

$$I(F) = \sum_{F_0} I(F_0),$$

where F_0 runs over every 0-handle of F .

In many circumstances, it will be useful to focus on handle structures \mathcal{H} with some constraints. We say that \mathcal{H} is *positive* if each 0-handle of \mathcal{F} has positive index, and, for each 0-handle H_0 of \mathcal{H} , $H_0 \cap (\mathcal{F} \cup \gamma)$ is connected.

5.2. STANDARD AND REGULATED SURFACES

A surface S properly embedded in a 3-manifold M with a handle structure \mathcal{H} is *standard* if

- (i) it intersects each 0-handle in a collection of properly embedded disjoint discs;
- (ii) its intersection with any 1-handle $D^1 \times D^2$ is $D^1 \times \alpha$, where α is a collection of arcs properly embedded in D^2 ;
- (iii) it intersects each 2-handle $D^2 \times D^1$ in discs of the form $D^2 \times \{*\}$, where $*$ is a point in the interior of D^1 ;
- (iv) it is disjoint from the 3-handles.

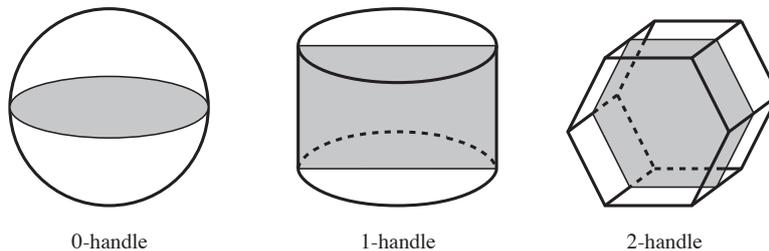


Figure 10: Standard surface

Let (M, γ) be a sutured manifold with a handle structure \mathcal{H} . Let S be a transversely oriented, standard surface properly embedded in M , with ∂S transverse to γ . Then, the sutured manifold (M', γ') obtained by decomposing along S inherits a handle structure. Each component of intersection between M' and an i -handle of \mathcal{H} becomes an i -handle for M' . Let \mathcal{H}' denote this handle structure.

We now explain how, in certain circumstances, the complexity of \mathcal{H}' must be at most that of \mathcal{H} . In order to guarantee this, the standard surface S must satisfy some conditions. These were explained in detail in Section 9 of [20]. We recall them here.

Condition 1. Each curve of $S \cap \partial\mathcal{H}^0$ meets any 1-handle of \mathcal{F} in at most one arc.

This crucial condition is also one that appears in the theory of normal surfaces in handle structures.

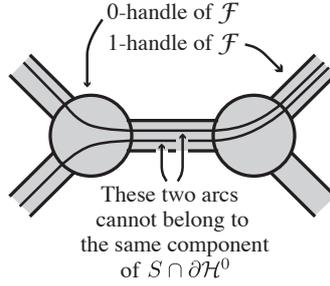


Figure 11: Condition 1

Condition 2. If α is an arc of $(\mathcal{F}^0 \cap \partial\mathcal{F}) - S$, and both its endpoints lie in S , then the transverse orientations on α and on $S \cap \mathcal{F}$ near $\partial\alpha$ cannot all agree.

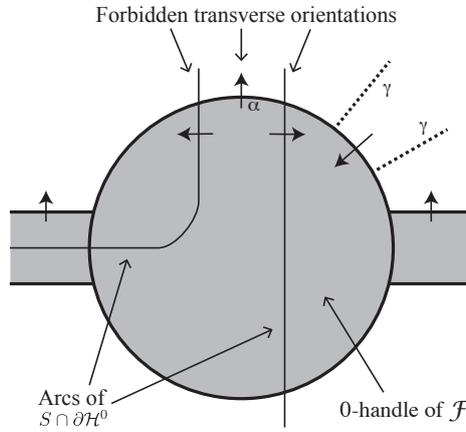


Figure 12: Condition 2

This condition does not arise in normal surface theory. It was one of the distinctive features of the surfaces considered in [20]. It imposes considerable constraints on the possible transverse orientations on the discs of $S \cap \mathcal{H}^0$ that intersect ∂M .

In [20], the arc α is known as a *tubing arc*, and we use the same terminology here.

Condition 3'. Suppose that D is a disc in \mathcal{F}^0 with ∂D the union of two arcs α and β , where $\alpha = S \cap \partial D$ and $\beta = D \cap \partial\mathcal{F}^0$. Suppose that one endpoint of α lies in $R_{\pm}(M)$ and one endpoint lies in \mathcal{F}^1 . Then at least one of the following must hold:

- (i) the interior of β contains a component of $\mathcal{F}^0 \cap \mathcal{F}^1$;
- (ii) β has non-empty intersection with γ .

This is very close to Condition 3 of [20], except it is somewhat weaker. In Condition 3 of [20], another configuration was ruled out, where β had a single intersection with γ , and S was oriented in a certain way.

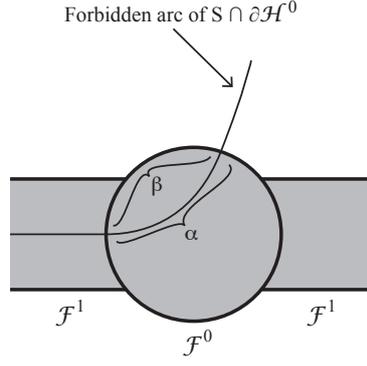


Figure 13: Condition 3'

Condition 4. Each component of $S \cap \partial \mathcal{H}^0$ meets any component of $R_{\pm}(M) \cap \partial \mathcal{H}^0$ in at most one arc.

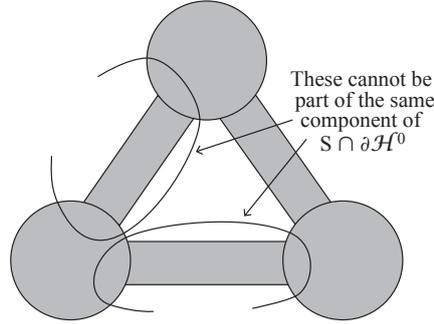


Figure 14: Condition 4

Condition 5'. If β is an arc of $S \cap \mathcal{F}^0$ with both endpoints in $R_{\pm}(M)$, then each of the two arcs in $\partial \mathcal{F}^0$ joining $\partial \beta$ must either contain a component of $\mathcal{F}^0 \cap \mathcal{F}^1$ or hit γ .

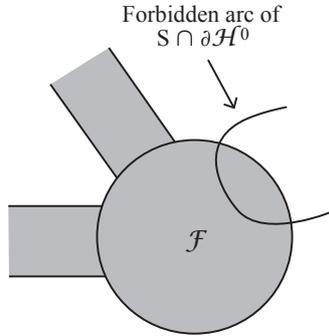


Figure 15: Condition 5'

Condition 5' is very similar to Condition 5 in [20], except it is weaker. In Condition 5 of [20], other configurations were also ruled out, where one of the arcs in $\partial \mathcal{F}^0$ joining $\partial \beta$ had at most two intersections with γ .

We say that a standard surface satisfying Conditions 1, 2, 3', 4 and 5' is *regulated*.

5.3. COMPARISON WITH NORMAL SURFACES

There is obviously a close connection between standard surfaces, regulated surfaces and the more well-known normal surfaces. In this subsection, we briefly recall the definition of normal surfaces and

compare them with standard and regulated surfaces.

Normal surfaces also come in various flavours. In the setting of triangulated 3-manifolds, a properly embedded surface is *normal* if it intersects each tetrahedron in a collection of triangles and squares. However, although triangulations will be used briefly in this paper, it is handle structures that will be our main concern.

So, let M be a compact 3-manifold with a handle structure \mathcal{H} . For simplicity, we will focus on the case where ∂M is not decorated in any way. In particular, we will disregard any sutured manifold structure. Then a *normal surface* in M is a standard surface satisfying Conditions 1 and 4 from the previous subsection and in addition, the following constraints:

- (i) No component of $S \cap \mathcal{F}^0$ has endpoints lying in the same component of $\partial\mathcal{F}^0 \cap \mathcal{F}^1$, or in the same component of $\partial\mathcal{F}^0 - \mathcal{F}^1$, or in adjacent components of $\partial\mathcal{F}^0 \cap \mathcal{F}^1$ and $\partial\mathcal{F}^0 - \mathcal{F}^1$.
- (ii) No arc of intersection between S and $\partial\mathcal{H}^0 - \mathcal{F}$ has endpoints in the same component of $\partial\mathcal{F}^0 - \mathcal{F}^1$.

5.4. ELEMENTARY DISC TYPES

Let \mathcal{H} be a handle structure of a 3-manifold, and let H be one of its handles. Let D be a disc properly embedded in H . Then a *normal isotopy* of D is an ambient isotopy that preserves each of the handles of \mathcal{H} and also γ (when M is a sutured manifold). When two discs properly embedded in H are normally isotopic, they are said to be of the same *type*.

Typically, we will be interested in discs D that are a component of $S \cap H$, for some handle H , where S is a standard, normal or regulated surface. Such discs are called *elementary discs*.

One of the limitations of standard surfaces is that there is, in general, no upper bound on the number of elementary disc types in a handle. However, it is clear that there is usually such an upper bound for regulated and for normal surfaces. In fact, for this to be true, we need to make the hypothesis that, for each 0-handle H of \mathcal{H} , $\partial H \cap (\gamma \cup \mathcal{F})$ is connected. This is equivalent to the statement that $\partial H \cap R_{\pm}(M)$ consists of discs or a sphere. Under this hypothesis, we can then show that for each handle H of \mathcal{H} , H contains only finitely many normal or regulated disc types, as follows.

It suffices to prove this for elementary discs in a 0-handle, because once these are controlled, then the elementary discs in 1-handles are determined. Furthermore, in each 2-handle, there is always only one elementary disc type. An elementary disc D properly embedded in a 0-handle H is determined, up to normal isotopy, by its boundary curve. This curve is determined by its intersection with \mathcal{F}^1 , with \mathcal{F}^0 and with $\partial H \cap R_{\pm}(M)$. It can run over each 1-handle of \mathcal{F} at most once by Condition 1 and it can run over each component of $\partial H \cap R_{\pm}(M)$ at most once by Condition 4. Thus, there are only finitely many possibilities for its intersection with \mathcal{F}^1 and $\partial H \cap R_{\pm}(M)$, because $\partial H \cap R_{\pm}(M)$ consists of discs. Hence, there are only finitely many possibilities for its intersection with $\partial\mathcal{F}^0$, and so there are only finitely many possibilities for its intersection with \mathcal{F}^0 . So, up to normal isotopy, there are only finitely many possibilities for D . Moreover, it is clear that once one is given the handle structure on $\mathcal{F} \cap \partial H$ and the arcs $\gamma \cap \partial H$, the possible elementary disc types are readily computable.

5.5. TRIVIAL AND WEAKLY TRIVIAL MODIFICATIONS

When a sutured manifold (M, γ) with a handle structure \mathcal{H} is decomposed along certain surfaces, then sometimes very little happens to (M, γ) and \mathcal{H} . In [20], decompositions of this sort were considered, and the resulting changes to \mathcal{H} were termed ‘trivial modifications’. In this subsection, we briefly recall this term in this context, and then introduce a variant, known as a ‘weakly trivial modification’.

Let S be a standard, transversely oriented surface properly embedded in (M, γ) . Let H_0 be a 0-handle of \mathcal{H} . Then decomposition along S results in a *trivial modification* to H_0 if

- (i) each component of $S \cap H_0$ is parallel to a disc in ∂H_0 that is disjoint from γ ;
- (ii) this component of $S \cap H_0$ has boundary entirely in \mathcal{F} ;
- (iii) the disc in ∂H_0 intersects \mathcal{F} in an annulus A , and each 0-handle of A has index zero.

We say that decomposition along S results in a *weakly trivial modification* to H_0 if

- (i) each component of $S \cap \mathcal{F}$ is parallel in \mathcal{F} to an arc or circle in $\partial \mathcal{F}$;
- (ii) in the resulting product region, each 0-handle of $\mathcal{F} - \text{int}(N(S))$ has index zero;
- (iii) no component of $S \cap \partial H_0$ separates components of $\mathcal{F} - \text{int}(N(S))$ with positive index.

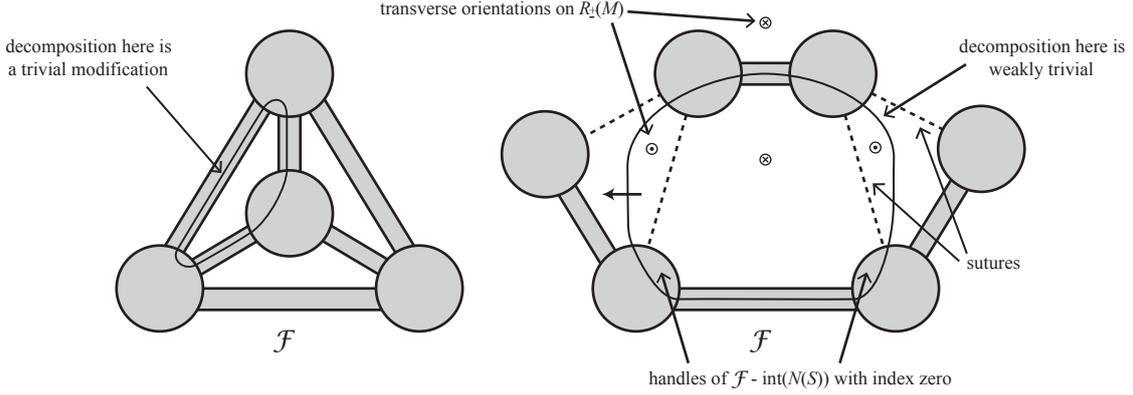


Figure 16: Trivial and weakly trivial modifications

Proposition 5.1. *Let \mathcal{H} be a handle structure of a sutured manifold (M, γ) . Suppose that \mathcal{H} is positive. Let*

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

be a taut decomposition, where S is a standard surface properly in M . Suppose that this results in a weakly trivial modification to each 0-handle of \mathcal{H} . Then S is product-separating.

Proof. We will show that each component S'' of S satisfies the following:

- (i) S'' is separating in the component of M that contains it;
- (ii) one of the components of the sutured manifold obtained by decomposing this component of (M, γ) along S'' is a product sutured manifold $F'' \times I$;
- (iii) the inclusion of S'' into $R_+(F'' \times I)$ or $R_-(F'' \times I)$ is a homotopy equivalence.

This will imply that S'' is product-separating, for the following reasons. Note that (iii) implies that any u-suture in $F'' \times I$ must be parallel to a trivial component of $\partial S''$. So when we form $E(S'')$, the effect is to add a disc along this trivial curve, and when we form $E(F'' \times I)$, the effect is to add a 2-handle along the corresponding suture. Hence, $E(F'' \times I)$ is still a product sutured manifold $F' \times I$, and moreover F' is homeomorphic to $E(S'')$. This verifies condition (i) in the definition of product-separating. We also verified that any attached 2-handles in $E(F'' \times I)$ respect its product structure and intersect $E(S'')$. Therefore, (ii) and (iii) in the definition of product-separating are also verified.

Our aim is now to verify (i), (ii) and (iii) above. We first modify S , forming a new surface S' . Consider any 0-handle H_0 of \mathcal{H} , and let F be $\mathcal{F} \cap \partial H_0$. Consider any arc component of $S \cap F$. By definition of a weakly trivial modification, each such arc is boundary parallel in F . Hence, we may find an arc that is outermost in F . It separates off an index zero subdisc of F which intersects ∂F in a single arc α . Push

α a little away from F . Then α forms part of the boundary of a product disc in M_S . The rest of the boundary of this product disc lies in the interior of a disc of $S \cap H_0$. Boundary compress S along this disc. Now repeat this process for all arcs of $S \cap F$, and for all 0-handles of \mathcal{H} . Note that this new surface S' will, in general, run over u-sutures, and so decomposition along S' need not be allowable, but it is taut by Proposition 2.5. Let $(M_{S'}, \gamma_{S'})$ be the sutured manifold that results from this decomposition.

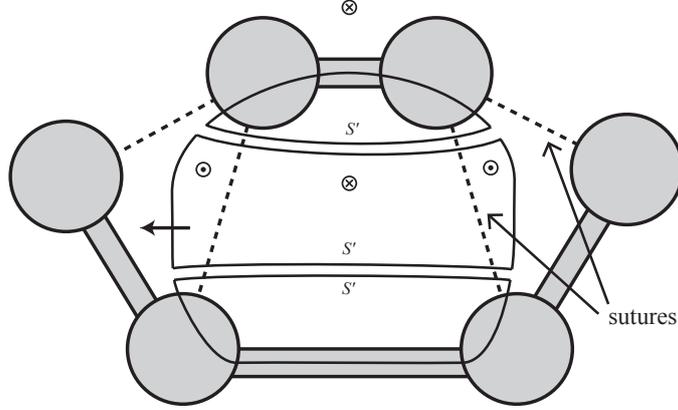


Figure 17: The surface S' that results from boundary-compressing S

We will now show that every component of S'' of S' satisfies (i), (ii) and (iii) above. Our proof follows Proposition 10.7 in [20] closely. Let C be the collection of arcs and simple closed curves of $S' \cap \mathcal{F}$ that are incident to components of $\mathcal{F} - \text{int}(N(S'))$ with positive index. Then C separates off from \mathcal{F} a collection of annuli A disjoint from γ and product discs P (which may contain other arcs and curves of $S' \cap \mathcal{F}$). Let D be the collection of discs of $S' \cap \mathcal{H}^0$ which contain a component of C in their boundary. Then $(A \cup P) \cup D$ is a collection of discs properly embedded in M . These are parallel to discs in ∂M via balls B_0 . The existence of such a disc in ∂M is guaranteed, in the case where a component of C lies in \mathcal{F} , by the fact that $H_0 \cap (\mathcal{F} \cup \gamma)$ is connected for each 0-handle H_0 of \mathcal{H} . When a component of C is an arc in \mathcal{F} , then the existence of the disc in ∂M follows from the construction of S' . Each of these balls B_0 may be given a product structure as $D^2 \times [-1, 1]$ so that $D^2 \times \{1\}$ is a component of $S' \cap B_0$.

For each 1-handle $H_1 = D^2 \times [0, 1]$, the discs $D^2 \times \{0\}$ and $D^2 \times \{1\}$ are each divided up by the decomposition along S' . For each $i \in \{0, 1\}$, all but one 0-handle of $D^2 \times \{i\} - \text{int}(N(S'))$ has index zero. The remaining component has index equal to the index of $D^2 \times \{i\}$. We are assuming that the index of $D^2 \times \{i\}$ is positive. Hence, the product structure on H_1 matches $(A \cup P) \cap (D^2 \times \{0\})$ with $(A \cup P) \cap (D^2 \times \{1\})$. We may therefore unambiguously define $(A \cup P) \cap D^2$. Let B_1 be the union, over all 1-handles, of the balls $((A \cup P) \cap D^2) \times [0, 1]$. Similarly, we may find a collection B_2 of balls in the 2-handles of M such that $B_2 \cap \mathcal{H}^0 = \mathcal{H}^2 \cap (A \cup P)$. Again, each component of B_1 and B_2 may be given a product structure of the form $D^2 \times [-1, 1]$ so that $S' \cap (D^2 \times [-1, 1]) \supseteq D^2 \times \{1\}$. Moreover, the product structures on B_0 , B_1 , and B_2 agree on their intersection. So, $B_0 \cup B_1 \cup B_2$ is an $[-1, 1]$ -bundle over a surface F , in which the sutures form the zero-section over ∂F . So, $B_0 \cup B_1 \cup B_2$ is a product sutured manifold that forms a product region between some components of S' and a subsurface of ∂M . Remove these components of S' and repeat. We eventually deduce that every component of S' that intersects \mathcal{F} is boundary parallel.

We claim also that the components of S' that are disjoint from \mathcal{F} are also boundary parallel. Their boundary curves come from components of $\partial S \cap \partial \mathcal{H}^0$, possibly isotoped a little so that they become disjoint from \mathcal{F} . Since no disc of $S \cap \mathcal{H}^0$ has 0-handles of M_S with positive index on both sides of it, we deduce that the discs of $S' \cap \mathcal{H}^0$ that are disjoint from \mathcal{F} cannot separate components of \mathcal{F} , which proves the claim. These discs therefore separate off components of $(M_{S'}, \gamma_{S'})$ which are balls. These balls are product sutured manifolds, because $(M_{S'}, \gamma_{S'})$ is taut.

We have therefore shown that each component S'' of S' satisfies (i), (ii) and (iii) above. Cutting M along S' gives a copy Y of M , together with some product sutured manifolds. We divide S' into subsurfaces S'_1, \dots, S'_n , as follows. Define the distance of a component of S' from Y to be the minimal number of intersection points between an arc and S' , where the arc starts on that component and ends in Y . Define S'_i to be the union of components with distance i . So, S'_1 is the union of components of S' incident to Y , and so on.

Consider S'_1 , the union of the components of S' closest to Y . These separate off product regions W in (M, γ) with interior disjoint from S . Consider any of the product discs that was used to boundary compress S , and that is incident to S'_1 . By construction, these have interior disjoint from W . Hence, when we reverse any of these boundary compressions, the resulting surface is still boundary parallel, and still separates off a product sutured manifold. This new surface satisfies (i), (ii) and (iii) above. Repeat this until all boundary compressions incident to S'_1 have been reversed. The resulting surface S_1 is a union of components of S . Continue with S'_2 , and so on.

In this way, we see that each component of S satisfies (i), (ii) and (iii). So, S is product-separating. \square

5.6. THE BEHAVIOUR OF COMPLEXITY UNDER DECOMPOSITION

For the purposes of this paper, it is critically important that when a sutured manifold (M, γ) with a handle structure \mathcal{H} is decomposed along a surface S , the resulting handle structure \mathcal{H}' is not any more complex than \mathcal{H} . One of the main reasons for introducing regulated surfaces is that this holds in this case, under some mild hypotheses. This was essentially shown in [20]. However, there the surfaces satisfied the slightly different Conditions 1-5. (See Propositions 10.6 and 10.7 in [20]). In this subsection, we explain how a similar result holds for regulated surfaces.

Theorem 5.2. *Let \mathcal{H} be a positive handle structure for a sutured manifold (M, γ) . Let*

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

be a taut decomposition along a regulated surface S , and let \mathcal{H}' be the handle structure obtained by decomposing \mathcal{H} along S . Then $C(H_0 \cap \mathcal{H}') \leq C(H_0)$ for each 0-handle H_0 of \mathcal{H} . If this inequality is an equality for some 0-handle H_0 , then $H_0 \cap \mathcal{H}'$ is obtained from H_0 by a weakly trivial modification. If this inequality is an equality for every 0-handle H_0 , then S is product-separating.

Throughout this subsection, we will use the following terminology. We denote the handle structures of M and M' by \mathcal{H} and \mathcal{H}' . As usual, \mathcal{H}^i will denote the union of i -handles of \mathcal{H} , and \mathcal{F} denotes $\mathcal{H}^0 \cap (\mathcal{H}^1 \cup \mathcal{H}^2)$. We define \mathcal{F}' similarly, but with \mathcal{H}' in place of \mathcal{H} .

The following is a version of Lemma 10.5 in [20]. The difference is that, in Lemma 10.5 of [20], Conditions 1, 2 and 3 were assumed to hold. However, we note that only the weaker Condition 3' was used in the proof.

Lemma 5.3. *Let S be a standard surface satisfying Conditions 1, 2 and 3'. Let D be a component of \mathcal{F}' with a negative index 0-handle. Then there is a 1-handle of \mathcal{F} which lies entirely in D .*

The following is a version of Propositions 10.6 and 10.7 in [20].

Proposition 5.4. *Let \mathcal{H} be a positive handle structure. Let H_0 be a 0-handle of \mathcal{H} , and let H'_0 denote the 0-handles $H_0 \cap \mathcal{H}'$. If S is a standard surface satisfying Conditions 1, 2 and 3', then $C(H'_0) \leq C(H_0)$. Also, if we have equality, then H'_0 is obtained from H_0 by a weakly trivial modification.*

The difference here is that only Conditions 1, 2 and 3' are assumed to hold, whereas in Proposition 10.6 in [20], Conditions 1 - 5 were hypothesised. However, we obtain a slightly weaker conclusion. We deduce in the case where $C(H'_0) = C(H_0)$ that H'_0 is obtained from H_0 by a weakly trivial modification. On the other hand, in Proposition 10.6 in [20], we could deduce that a trivial modification had been performed.

Proof. Let F be a component of $\mathcal{F} \cap H_0$, and let F' be $F \cap \mathcal{F}'$. The first paragraph in the proof of Proposition 10.6 in [20] gives that $C_{\mathcal{F}}(F') \leq C_{\mathcal{F}}(F)$. Since this holds for each component F of $\mathcal{F} \cap H_0$, we deduce that $C_{\mathcal{F}}(H'_0) \leq C_{\mathcal{F}}(H_0)$. Note also that $n(H_0) = 1$ and $n(H'_0) \geq 1$. Therefore, $C(H'_0) \leq C(H_0)$.

Suppose that this is an equality. This implies that if F is any component of $\mathcal{F} \cap H_0$, and $F' = F \cap \mathcal{F}'$, then $C_{\mathcal{F}}(F') = C_{\mathcal{F}}(F)$. As argued in the first paragraph of the proof of Proposition 10.6 in [20], one component X of F' has $C_1(X) = C_1(F')$ and $I(X) = I(F')$. As argued in the second paragraph of the proof of Proposition 10.6 in [20], $|\partial X| = |\partial F'|$, and so F' is obtained from F by cutting along arcs and circles which are boundary parallel in F . Each component of $F' - X$ has index zero. Moreover, every 0-handle of F' not lying in X has index zero. We therefore have verified conditions (i) and (ii) in the definition of a weakly trivial modification.

Therefore, each component F of $\mathcal{F} \cap H_0$ gives rise to a corresponding component of $\mathcal{F}' \cap H_0$ also with positive index. All of these components of \mathcal{F}' must lie in the same component of H'_0 , as otherwise $n(H'_0) > 1 = n(H_0)$. Thus, we have verified condition (iii) in the definition of a weakly trivial modification. \square

Proof of Theorem 5.2. According to Proposition 5.4, decomposition along the regulated surface does not increase the complexity of any 0-handle. Also, if the complexity of a 0-handle is unchanged, then a weakly trivial modification is performed there. If a weakly trivial modification occurs at every 0-handle, then according to Proposition 5.1, S is product-separating. \square

5.7. THE COMPLEXITY AND WEIGHT OF STANDARD SURFACES

Many arguments in this field aim to modify surfaces until they are of some required type. For example, one might first make a surface standard and then perform further modifications to make it normal. In this second stage, one typically aims to reduce some measure of complexity for the surface.

In this paper, we define the *complexity* of a standard surface S to be the ordered pair of integers $(|S \cap \mathcal{H}^2|, |\partial S \cap \mathcal{H}^1|)$. We compare complexities in the usual way, using lexicographical ordering. This notion of complexity is one of the most commonly used ones. In particular, it is the same as the one employed by the author in [20].

The first integer in the pair, $|S \cap \mathcal{H}^2|$, is called the *weight* of the surface, and is also denoted $w(S)$. Similarly, when S is a normal surface properly embedded in a triangulated 3-manifold, its weight $w(S)$ is its number of points of intersection with the 1-skeleton.

5.8. THE EXISTENCE OF REGULATED SURFACES

In Section 9 of [20], it was established that a standard decomposing surface may be often upgraded to a regulated one. In this and the next subsection, we will prove a version of this result, which is as follows.

Theorem 5.5. *Let*

$$(M, \gamma) \xrightarrow{S} (M', \gamma')$$

be an allowable decomposition between taut decorated sutured manifolds that extends to an allowable hierarchy. Suppose that $E(M, \gamma)$ is atoroidal and that no component of $E(M, \gamma)$ is a Seifert fibre space, other than a solid torus or a copy of $T^2 \times I$. Assume that (M, γ) is not a product sutured manifold. Let \mathcal{H} be a positive handle structure for (M, γ) . Suppose that S is standard in \mathcal{H} . Then there is another allowable sutured manifold decomposition

$$(M, \gamma) \xrightarrow{S'} (M'', \gamma'')$$

such that

- (i) S' is a regulated surface in \mathcal{H} ;

- (ii) *it is connected;*
- (iii) *it is not product-separating;*
- (iv) *S' extends to an allowable hierarchy;*
- (v) *(M'', γ'') is taut;*
- (vi) *if S is non-separating, then S' is also.*

This is proved using the methods developed in Section 9 of [20]. There, various modifications were made to the standard surface S . These modifications do not increase the complexity of the surface at any stage. See the discussion after Lemma 9.9 in [20] where this assertion is explicitly made. Indeed the complexity of S is defined the way it is in [20] to ensure that it is decreased by these modifications, and hence that these modifications are guaranteed to terminate. The modifications are as follows.

Modification 1. Tubing along an arc α

This was described in detail in Section 4.5. We will always require that the two arcs of $\alpha \times \{-1, 1\}$ do not lie in the same trivial boundary curve of the resulting surface S' . Note that, according to Lemma 4.5, S' extends to an allowable hierarchy.

Modification 2. Slicing under a non-trivial annulus.

Let A be an annulus in $R_{\pm}(M)$ which has non-trivial boundary curves and has $A \cap S = \partial A$. Suppose that the orientation on S agrees with orientation on A . Then we can construct a new surface S' by attaching A to S , and then isotoping so that it becomes properly embedded. Let (M_S, γ_S) and $(M_{S'}, \gamma_{S'})$ be the sutured manifolds obtained from (M, γ) by decomposing along S and S' respectively. Note that there is a non-trivial product annulus P properly embedded in $(M_{S'}, \gamma_{S'})$, with one boundary component being a core curve of A , and the other boundary component being a core curve of the copy of A in S' . Moreover, decomposing $(M_{S'}, \gamma_{S'})$ along P gives a decorated sutured manifold homeomorphic to (M_S, γ_S) . Therefore, by Proposition 4.9, S' extends to an allowable hierarchy if and only if S does, under the hypotheses of Theorem 5.5.

Modification 3. Sliding across γ .

We will not describe this move in detail here, because we will avoid using it in this paper. The reason is that it creates new intersection points between ∂S and γ . In [20], this was fine, but in this paper, ∂S is not permitted to intersect a u-suture. This is why we avoid this move.

Modification 4. Slicing under a disc of contact.

This modification was described in detail in Section 4.3. According to Lemma 4.3, the resulting surface S' also extends to an allowable hierarchy.

Modification 5. Boundary compressing along a product disc disjoint from γ , and then possibly removing a planar component.

This is described in Section 4.6. If the surface obtained by boundary-compressing has a component that is planar, disjoint from γ and with all but at most one boundary component trivial, then this component is removed. By Proposition 4.7, if S' is obtained from S by this process, and S extends to an allowable hierarchy, then so does S' .

Modification 6. Removal of a product region.

Suppose that a component S_1 of S is parallel into $R_{\pm}(M)$, and that the interior of the product region between S_1 and the subsurface of $R_{\pm}(M)$ is disjoint from S . This modification is the removal of S_1 from S . Let S' be the resulting surface. Note that the sutured manifold obtained by decomposing along S is equal to the disjoint union of the manifold obtained decomposing along S and a product sutured manifold. Hence, S' extends to an allowable hierarchy if and only if S does.

Lemma 5.6. *If a surface is non-separating in M , then this remains the case for at least one component of the surface that results from Modifications 1, 2, 4, 5 and 6.*

Proof. If one starts with a surface S , then each of these modifications does not change $[S, \partial S] \in H_2(M, \partial M)$. If S is non-separating, then $[S, \partial S]$ is non-trivial and so the resulting surface S' also has $[S', \partial S']$ non-trivial. Hence, some component is homologically non-trivial, and therefore non-separating. \square

Proposition 5.7. *Suppose that (M, γ) contains no non-separating product disc or non-separating product annulus. Let S be a connected surface in (M, γ) that is not product-separating and that extends to an allowable hierarchy. Then this remains the case for at least one component of the surface that results from Modifications 1, 2, 4, 5 and 6.*

Proof. We may assume that M is connected, by focusing on the component containing S . Note also that the statement of the lemma is empty for Modification 6, because S is connected and not product-separating, and so Modification 6 cannot be applied to it.

If S is non-separating, the lemma follows from Lemma 5.6. So, suppose that S is separating. We may also assume that each component of S' is separating, because a non-separating surface is automatically not product-separating.

We now consider each of the modifications in turn. Suppose that S' is obtained from S by Modification 1. As explained in Section 4.5, there is a decomposition

$$(M_{S'}, \gamma_{S'}) \xrightarrow{D} (M_S, \gamma_S)$$

where D is a product disc. We are assuming that the two arcs $\alpha \times \{-1, 1\}$ do not lie in the same trivial boundary curves of S' . Hence, there are a number of other possible cases to consider.

Suppose first that at least one arc of $\alpha \times \{-1, 1\}$ lies in a trivial boundary curve of S' . We will focus on the case where exactly one arc of $\alpha \times \{-1, 1\}$ lies in a trivial curve of $\partial S'$, as the other case is very similar. Then $E(S')$ is homeomorphic to $E(S)$. If S' is product-separating, then $E(S')$ separates off a product manifold homeomorphic to $E(S') \times I$ in which each attached 2-handle is vertical and intersects $E(S')$. The manifold $E(M_{S'}, \gamma_{S'})$ is homeomorphic to $E(M_S, \gamma_S)$. In fact, $E(M_{S'}, \gamma_{S'})$ is obtained from $E(M_S, \gamma_S)$ by attaching a copy of $D^2 \times I$ along a sub-arc of γ_S . This region $D^2 \times I$ is made up of attached 2-handles and $P \times I$, where P is the trivialising planar surface for the trivial boundary curve of $\partial S'$ incident to $\alpha \times \{-1, 1\}$. So, $E(M_S, \gamma_S)$ contains a component homeomorphic to $E(S) \times I$ and the attached 2-handles are vertical in this product structure and intersect $E(S)$. So, S is product-separating.

Suppose now that neither arc of $\alpha \times \{-1, 1\}$ lies in a trivial boundary curve of S' . Then there is a decomposition

$$E(M_{S'}, \gamma_{S'}) \xrightarrow{D} E(M_S, \gamma_S).$$

One component of $E(M_{S'}, \gamma_{S'})$ is, by assumption, a copy of $E(S') \times I$. If this component contains D , then we decompose along it to obtain a component of $E(M_S, \gamma_S)$ that is a copy of $E(S) \times I$. Hence, in this case, S is product-separating. So, suppose that the product manifold $F \times I$ does not contain D . Then it remains a product manifold component of $E(M_S, \gamma_S)$. However, it may not be of the required form: it may not be homeomorphic to $E(S) \times I$. So, we use an alternative argument. Consider an arc of $\alpha \times \{-1, 1\}$. This is part of the boundary of S' . Consider the arc in $R_{\pm}(E(M_S, \gamma_S))$ that is the copy of this arc of $\alpha \times \{-1, 1\}$, and that is not part of a suture. Say it lies in $R_+(E(M_S, \gamma_S))$. This arc lies in the product sutured manifold $F \times I$. Since each component of $\text{cl}(R_+(F \times I) - E(S'))$ intersects γ by Lemma 4.10, we may pick an arc running in $R_+(M)$ from $\alpha \times \{-1, 1\}$ to γ . We do this for each of the two components of $\alpha \times \{-1, 1\}$. Join these two arcs by an arc running over the tube. The result is an arc β in the product sutured manifold. This is part of a product disc in $F \times I$. Now enlarge this product disc across the tube to form a product disc in (M, γ) . It is non-separating, since one may find a closed curve intersecting it once, consisting of an arc in the connected surface S and an arc running across the tube. This contradicts a hypothesis of the lemma.

The argument in the case of Modification 2 is similar but simpler. In this case, there is a decomposition

$$E(M_{S'}, \gamma_{S'}) \xrightarrow{P} E(M_S, \gamma_S),$$

where P is a product annulus. By assumption, some component of $E(M_{S'}, \gamma_{S'})$ is a product sutured manifold. If P lies in it, then we deduce that S is also product-separating. On the other hand, if P does not lie in the product manifold, then the union of P and a vertical annulus in the product structure forms a product annulus in (M, γ) . It is non-separating, because one may find a closed curve that intersects it once, as in the previous case. This contradicts one of the hypotheses of the lemma.

Modification 4 is particularly easy to handle. If S' is obtained from S by slicing under a disc of contact, then $E(S')$ and $E(S)$ are homeomorphic, as are $E(M_S, \gamma_S)$ and $E(M_{S'}, \gamma_{S'})$. Moreover, in the latter case, this homeomorphism sends the attached 2-handles to the attached 2-handles. So, S' is product-separating if and only if S is.

Now consider the case of Modification 5. The surface S' may be connected or disconnected. Suppose first that S' is connected. Suppose also that the tubing arc α is disjoint from the product sutured manifold. There are a number of possibilities. If neither curve of $\alpha \times \{-1, 1\}$ lies in a trivial boundary curve of S , then we obtain a decomposition

$$(E(S) \times [-1, 1], \partial E(S) \times \{0\}) \xrightarrow{D} (E(S') \times [-1, 1], \partial E(S) \times \{0\}).$$

Hence, S is product-separating. If just one curve of $\alpha \times \{-1, 1\}$ lies in a trivial boundary curve of S , then $E(S)$ is homeomorphic to $E(S')$, and $E(M_S, \gamma_S)$ is homeomorphic to $E(M_{S'}, \gamma_{S'})$. Hence, as argued in the case of Modification 1, $E(M_S, \gamma_S)$ contains a component homeomorphic to $E(S) \times I$ and any attached 2-handles are vertical in this product structure. So, again S is product-separating. If the arcs of $\alpha \times \{-1, 1\}$ lie in distinct trivial boundary curves of S , then again $E(S)$ is homeomorphic to $E(S')$, and $E(M_S, \gamma_S)$ is homeomorphic to $E(M_{S'}, \gamma_{S'})$, and so S is again product-separating. Finally, suppose that the arcs $\alpha \times \{-1, 1\}$ lie in the same trivial boundary curve of S . Then we obtain $E(S)$ from $E(S')$ by attaching an annulus, and $E(S') \times [-1, 1]$ is obtained from a component of $E(M_S, \gamma_S)$ by decomposing along a product annulus. Therefore, this component of $E(M_S, \gamma_S)$ is homeomorphic to $E(S) \times I$, as required.

Now consider the case where S' is connected, and the tubing arc α lies in the product sutured manifold. Then, α lies in the closure of a component of $R_{\pm}(M) - \partial S'$. By Lemma 4.10, this is a disc that intersects $\partial S'$ in a single arc or an annulus that intersects $\partial S'$ in a simple closed curve. Therefore, exactly one arc of $\alpha \times \{-1, 1\}$ lies in a trivial component of ∂S . We have already seen in this case that $E(S)$ is homeomorphic to $E(S')$, and $E(M_S, \gamma_S)$ is homeomorphic to $E(M_{S'}, \gamma_{S'})$. Hence, S is product-separating.

Now suppose that S' is disconnected with components S'_1 and S'_2 , that are both product-separating. If neither of the resulting product manifolds contains the tubing arc, then the argument above applies. So, suppose that at least one of the product manifolds, $E(S'_1) \times I$, say, contains the tubing arc. Then α lies in the closure of a component of $R_{\pm}(M) - \partial S'_1$. By Lemma 4.10, this is a disc D that intersects $\partial S'_1$ in a single arc or an annulus that intersects $\partial S'_1$ in a simple closed curve. Consider the disc case first. The other end of α lies in S'_2 . This creates a suture of $\gamma_{S'}$ which lies in D . Since decomposition along S' is taut, we deduce that $E(S'_2)$ is a disc, which separates off a product ball. Hence, there is a component of $E(M_S, \gamma_S)$ homeomorphic to $E(S'_1) \times I$, and this homeomorphism takes 2-handles to 2-handles. Also, $E(S)$ is homeomorphic to $E(S'_1)$. So, S is product-separating.

Suppose now that α lies in an annulus A that is the closure of a component of $R_{\pm}(M) - \partial S'_1$. We may assume that the intersection between $\partial S'_2$ and A does contain any arcs. For if it did, an outermost arc in A would separate off a disc, and we could then argue as in the previous case. Therefore, we may assume that $\partial S'_2$ intersects A in at least one core curve. Hence, decomposing (M, γ) along S'_2 gives an annular component of R_{\pm} . Suppose that this lies in the product sutured manifold that $E(S'_2)$ separates off. Then $E(S'_2)$ is an annulus. Now, the tubing arc α lies in an annular component of $R_{\pm}(M) - \partial S'$. Hence, the two arcs $\alpha \times \{-1, 1\}$ lie in the same trivial curve of ∂S . This bounds a disc of contact. So, we deduce

that $E(S)$ is homeomorphic to $E(S'_1)$, and $E(M_S, \gamma_S)$ is homeomorphic to $E(M_{S'_1}, \gamma_{S'_1})$. Therefore, S must have been product-separating.

There is one final case that we must consider. Suppose that S' is disconnected with components S'_1 and S'_2 , both of which are product-separating. Suppose that both of the product manifolds contain the tubing arc α . Assume also that α lies in an annular component of $R_\pm(M) - S'_1$ and an annular component of $R_\pm(M) - S'_2$. Then, α lies in an annular components A of $R_\pm(M)$, and it runs between a component of $\partial S'_1$ and a component of $\partial S'_2$ that are essential curves in this annulus. Let $F_1 \times I$ be the product manifold that $E(S'_1)$ separates off. Then $E(S'_2)$ lies in $F_1 \times I$. It is an incompressible surface in this product, and hence vertical or parallel to a subsurface of $F_1 \times \partial I$. In the vertical case, it is actually parallel to a regular neighbourhood of a component of ∂A . In the boundary parallel case, it separates off a product submanifold. This has an annular component of R_\pm , that is a subset of A . Hence, $E(S'_2)$ must be an annulus. It is therefore again parallel to a regular neighbourhood of a component of ∂A . Now, when we tube along α to form S , we create a curve of ∂S that bounds a disc of contact in A . Hence, $E(S)$ is actually homeomorphic to $E(S'_1)$. Also, a component of (M_S, γ_S) is homeomorphic to $F_1 \times I$. So, S is product-separating. \square

5.9. ENSURING CONDITIONS 1, 2, 3', 4 AND 5'

We follow the procedure given in Section 9 of [20], where Conditions 1 - 5 were introduced. There, it was shown that if the surface S did not satisfy these conditions, then Modifications 1 - 6 could be performed to it which reduced its complexity. These modifications were performed in a specific order, and we follow a similar strategy here.

First, if a surface fails to satisfy Condition 1, 2 or 3', then moves described in Section 9 of [20] are performed which reduces its complexity. We note that when Condition 3' is violated, then Modification 3 is not required. Note also that Modification 4 might be required when making the new surface standard, as in Lemma 9.9 in [20]. We therefore obtain the following lemma.

Lemma 5.8. *Let*

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

be an allowable decomposition between taut decorated sutured manifolds. Suppose that S is standard in the handle structure \mathcal{H} , but fails to satisfy one of Conditions 1, 2 or 3'. Then a sequence of Modifications 1, 2 and 4 can be applied to create a standard surface with smaller complexity.

It is shown in [20] that if S does not satisfy Conditions 4 or 5', then Modifications 4, 5 or 6 may be applied to take it to a standard surface with smaller complexity. So, we obtain the following.

Lemma 5.9. *Let*

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

be an allowable decomposition between taut decorated sutured manifolds. Then a sequence of Modifications 1, 2, 4, 5 and 6 can be applied to create a regulated surface with no greater complexity.

Proof of Theorem 5.5. By Proposition 4.11, (M, γ) admits an allowable hierarchy, where the first surface S is connected and is not product-separating. Suppose first (M, γ) admits such a hierarchy where the first surface is non-separating. By Lemma 5.9, we may apply a sequence of Modifications 1, 2, 4, 5 and 6 to create a regulated surface S'' with smaller complexity. This also extends to an allowable hierarchy. By Lemma 5.6, at least one component of S'' is non-separating. Call this S' . Then S' is automatically not product-separating.

Suppose now that (M, γ) does not admit an allowable hierarchy where the first surface S is connected and non-separating. In particular, (M, γ) does not contain a non-separating product disc or a non-separating product annulus by Lemma 4.4 and Proposition 4.9. Then, by Lemma 5.9, a sequence of Modifications 1, 2, 4, 5 and 6 can be applied to create a regulated surface with no greater complexity. This also extends to an allowable hierarchy. At each stage, one of the components of the surface is not

product-separating, by Proposition 5.7. At each stage, we focus on this component, and discard the rest. At the final stage, we end with the required surface S' . \square

6. NEW HANDLE STRUCTURES FROM OLD ONES

6.1. SIMPLIFICATION MODIFICATIONS

According to Theorem 5.5, if a taut sutured manifold admits an allowable hierarchy, then one may find such a hierarchy where the first decomposing surface is regulated. However, the theorem has various technical hypotheses, including that the initial handle structure \mathcal{H} must be positive. Recall from Section 5.1 that this means that each 0-handle of \mathcal{F} must have positive index and, for each 0-handle H_0 of \mathcal{H} , $H_0 \cap (\mathcal{F} \cup \gamma)$ is connected. In this subsection, we explain some of the modifications that one can make to a handle structure of a sutured manifold to ensure that it becomes positive. These procedures were first introduced in Sections 7 and 8 of [20].

Procedure 1. Slicing a 0-handle along a disc.

Suppose that D is a disc properly embedded in some 0-handle H_0 , and that ∂D lies in $R_{\pm}(M)$, is disjoint from \mathcal{F} , and separates $\mathcal{F} \cap H_0$. Assuming that $R_{\pm}(M)$ is incompressible, then ∂D bounds a disc D' in $R_{\pm}(M)$. Then Procedure 1 is decomposition along D , oriented so that a suture appears along $D \cap D'$. If (M, γ) is taut, then this decomposition is taut.

Procedure 2. Collapsing a 2-handle and a 1-handle disjoint from γ .

Suppose that H_1 is a 1-handle of M disjoint from γ and that intersects \mathcal{H}^2 in a single disc. Let H_2 be the 2-handle containing this disc. Then Procedure 2 is the removal of H_1 and H_2 . It is also the enlargement of \mathcal{H}^3 if these handles are incident to \mathcal{H}^3 .

Procedure 3. Collapsing a 2-handle and a 1-handle containing an arc of γ .

Let H_1 be a 1-handle of M that intersects \mathcal{H}^2 in a single disc and intersects γ in a single arc. Let H_2 be the 2-handle incident to H_1 . Then Procedure 3 is the removal of H_1 and H_2 . The arc $\gamma \cap H_1$ is replaced by an arc that runs along the part of the attaching circle of H_2 that is disjoint from H_1 , and also through the two attaching discs of H_1 .

Procedure 4. Decomposition along a product disc, then sliding γ .

Let H_1 be a 1-handle disjoint from \mathcal{H}^2 and that intersects γ in two arcs. Then Procedure 4 is the removal of H_1 . For each of the two discs that form $H_1 \cap \mathcal{H}^0$, the two points of intersection between this disc and γ are joined by new sutures. This is achieved by a decomposition along the product disc which is the co-core of H_1 . In this paper, we will only perform this move when neither of the arcs of $\gamma \cap H_1$ lie in a u-suture. This guarantees that decomposition along the product disc is allowable.

Procedure 5. Collapsing a 3-ball disjoint from γ .

Suppose that a component of M is a 3-ball disjoint from γ , and consists of two 0-handles joined by a 1-handle. Then Procedure 5 is the removal of this 1-handle and one of the 0-handles.

Procedure 6. Collapsing a 2-handle and 3-handle.

Let H_2 be a 2-handle that intersects \mathcal{H}^3 in a single disc. Let H_3 be the incident 3-handle. Then Procedure 6 is the removal of H_2 and H_3 .

Procedure 7. Introduction of a 2-handle.

Suppose that M' is a subset of M with the following properties:

- (i) M' is a union of handles of \mathcal{H} ;
- (ii) if any j -handle of \mathcal{H} is incident to an i -handle of M' , where $j > i$, then the j -handle also lies in M' ;
- (iii) M' is homeomorphic to $D^2 \times I$ (and we will make this identification);
- (iv) $(D^2 \times I) \cap \partial M = D^2 \times \partial I$;
- (v) M' is disjoint from γ ;
- (vi) the only handles of M' incident to $\partial D^2 \times I$ are 1-handles and 2-handles;
- (vii) each 1-handle of M' incident to $\partial D^2 \times I$ intersects the 2-handles in exactly two discs.

Then in Procedure 7, M' is removed from M , and is replaced by a single 2-handle. Note that by (i) and (ii), $\text{cl}(M - M')$ inherits a handle structure. By (iii) and (iv), attaching the 2-handle to $\text{cl}(M - M')$ results in a manifold homeomorphic to M . By (vi) and (vii), the attaching locus of this 2-handle is well-defined. By (v), M' inherits a sutured manifold structure homeomorphic to (M, γ) .

6.2. THE UNIVERSAL COLLECTION OF 0-HANDLE TYPES

Handle structures are very general objects and, in particular, there are infinitely many possibilities for the way that a handle may intersect the union of its neighbours. We will need to restrict our attention to handle structures where the 0-handles come in finitely many possible local types, as follows.

Define a 0-handle H_0 as *tetrahedral* if $H_0 \cap \mathcal{F}$ has four 0-handles and six 1-handles, so that any two 0-handles of \mathcal{F} are joined by a 1-handle. (The boundary of a tetrahedral 0-handle is shown in the left of Figure 16.) Tetrahedral 0-handles are very common. For example, if one starts with any triangulation of a closed 3-manifold, and one forms the dual handle structure, then each 0-handle is tetrahedral.

Define a 0-handle H_0 as *subtetrahedral* if $H_0 \cap \mathcal{F}$ is obtained from the tetrahedral case by removing some handles. These arise, for example, when one starts with a triangulation of a compact 3-manifold with boundary. Then, the dual handle structure has an i -handle for each $(3 - i)$ -simplex that does not lie entirely in ∂M . It is easy to see that each 0-handle is subtetrahedral.

We will decompose handle structures along regulated surfaces, to form new handle structures. The following result, which is explained in Section 11 of [20] and in the proof of Theorem 1.4 of [20], implies that, in the resulting handle structures, there are only finitely many possible types of 0-handle. It is proved by induction on the complexity of the handle structure, as defined in Section 5.1.

Theorem 6.1. *There is a universal computable constant N and a universal computable collection of 0-handle types H_0^1, \dots, H_0^k , with following property. Suppose that one starts with a subtetrahedral 0-handle disjoint from any sutures. Suppose then that a sequence of simplifying procedures as in Section 6.1 and decompositions along regulated surfaces are performed. Assume also that, whenever a decomposition along a regulated surface is made, then the handle structure is positive. Then, each of resulting 0-handles is one of the universal types H_0^i , and at most N of these have positive index.*

We say that a handle structure of a sutured manifold is *of uniform type* if each 0-handle is one of the types given in Theorem 6.1. An algorithm to produce all the 0-handle types in Theorem 6.1 is given in Section 11 of [20].

7. BOUNDARY-REGULATED SURFACES

A key feature of normal surfaces is that they can be encoded by a finite list of numbers, which form a vector. This introduction of linear algebra is very useful. Unfortunately, it does not seem to be possible to create such a theory for regulated surfaces. The main difficulty is that regulated surfaces are required to have a transverse orientation, which is hard to incorporate algebraically.

Therefore, in this section, we introduce a new type of surface, which we call boundary-regulated. These are weaker than regulated surfaces, but they have the advantage of admitting a more algebraic theory, somewhat akin to normal surface theory.

7.1. DEFINITION

Let (M, γ) be a sutured manifold with a handle structure \mathcal{H} . Let S be a standard surface properly embedded in M . We say that S is *boundary-oriented* if each elementary disc D of S that intersects ∂M is assigned a transverse orientation, and, at each point in ∂S where two elementary discs intersect, their transverse orientations agree.

We note that the Conditions 1, 2, 3', 4 and 5' in Section 5.2 only depend on a transverse orientation on S near ∂S . Hence, they continue to make sense for a boundary-oriented surface S , even when S does not admit a global transverse orientation. We say that a boundary-oriented standard surface satisfying Conditions 1, 2, 3', 4 and 5' is *boundary-regulated*. Thus, the only difference between regulated and boundary-regulated surfaces is that regulated surfaces have a global transverse orientation, whereas boundary-regulated surfaces have a transverse orientation specified only over a subset of the surface.

In this section, we will develop a theory of boundary-regulated surfaces. But first, we recall some of the main points of classical normal surface theory.

7.2. NORMAL SURFACE VECTORS

Let \mathcal{H} be a handle structure for a 3-manifold M . It is well known that normal surfaces properly embedded in M may be described as solutions to a system of linear equations, as follows.

There is one variable for each type of elementary normal disc in the 0-handles of \mathcal{H} . So, for a normal surface S , one obtains a list of non-negative integers, which count the number of elementary normal discs of S of each type in each 0-handle. This is the *normal surface vector* associated with S , and is denoted by (S) .

There is one equation for each type of elementary normal disc in the 1-handles of \mathcal{H} . When an elementary normal disc in a 0-handle runs over a 0-handle of \mathcal{F} , it intersects this 0-handle in a collection of arcs. These arcs extend in a unique way to elementary normal discs in the associated 1-handle of \mathcal{H} . The *matching equations* express the fact that, for each 1-handle H_1 of \mathcal{H} that is attached to 0-handles H_0 and H'_0 , and for each type of elementary normal disc D in H_1 , the numbers of elementary normal discs in H_0 and H'_0 which touch elementary discs of type D are equal.

There are some other further constraints on normal surface vectors. We say that two types of normal disc in a 0-handle of \mathcal{H} are *incompatible* if there is no normal isotopy that makes them disjoint. Clearly, a properly embedded normal surface cannot contain incompatible normal discs. The *compatibility constraints* assert that for each pair of incompatible normal disc types in a 0-handle of \mathcal{H} , one of the associated variables is zero.

Haken showed [8] that there is a one-one correspondence between normal surfaces, up normal isotopy, and non-negative integer solutions to the matching equations that satisfy the compatibility constraints.

7.3. SUMMATION OF NORMAL SURFACES

Given that normal surfaces have such a close connection with linear algebra, it makes sense to exploit this algebraic structure. Haken therefore introduced the following definition [8].

A properly embedded normal surface S is said to be the *sum* of two other normal surfaces S_1 and S_2 if $(S) = (S_1) + (S_2)$.

In this situation, there is a well-known topological interpretation of this summation. Indeed, suppose merely that $(S) = v_1 + v_2$ for non-negative integral vectors v_1 and v_2 satisfying the matching equations. Then v_1 and v_2 also satisfy the compatibility constraints, and hence correspond to normal surfaces S_1 and S_2 . Consider the elementary normal discs of S in each 0-handle of \mathcal{H} . Since $(S) = (S_1) + (S_2)$, we may partition these discs into two collections, so that the number of elementary normal discs of each type in the first (respectively, second) collection is equal to the number of elementary normal discs of that type in S_1 (respectively, S_2). Since S_1 satisfies the matching equations, it is possible to build S_1 starting from these elementary normal discs by inserting discs into the 1-handles and 2-handles. Specifically, we consider, for each 1-handle $D^1 \times D^2$, the arcs of $S_1 \cap (\partial D^1 \times D^2)$ of each type. Because (S_1) satisfies the matching equations, there are the same number of arcs of this type in each component of $\partial D^1 \times D^2$. Hence, we may insert discs into $D^1 \times D^2$ interpolating between these arcs. We may do this for each disc type in the 1-handles. We may similarly insert discs into the 2-handles. The result is not quite a standard surface, because it does not respect the product structure on the 1-handles and 2-handles. Specifically, for any 1-handle $D^1 \times D^2$ and each point p in D^1 , the intersection between $\{p\} \times D^2$ is a collection of properly embedded arcs, but these arcs vary as p varies. See Figure 18. Clearly, one may ambient isotope this surface to form a normal surface, normally isotopic to S_1 , but we will not do so.

We may perform the same construction for S_2 . Then S_1 and S_2 may intersect in the 1-handles and 2-handles, and after a small isotopy, this intersection will be a collection of simple closed curves and properly embedded arcs. These are called *trace curves*.

Each of these trace curves may be resolved in two possible ways. One way is called a *regular switch* if, when we perform the resolution along such a trace curve, the resulting discs in the 1-handles are still normally isotopic to standard discs. When the trace curves are all resolved in this way, the result is a surface normally isotopic to S . The other way of resolving a trace curve is called an *irregular switch*. If a regular switch is performed along all but one trace curve, and an irregular switch is performed there, then the resulting surface is not normal. In fact, there is an ambient isotopy that reduces its complexity.

If one removes an open regular neighbourhood of the trace curves from S_1 and S_2 , the components of the resulting compact surface are called *patches*.

The trace curves $S_1 \cap S_2$ give rise to discs, annuli and Möbius bands properly embedded in the exterior of S , as follows. Since S is obtained from $S_1 \cup S_2$ by resolving the intersections $S_1 \cap S_2$, an I -bundle over $S_1 \cap S_2$ lies in the manifold, with the ∂I -bundle forming precisely the intersection with S . This I -bundle forms the *trace discs, annuli and Möbius bands*.

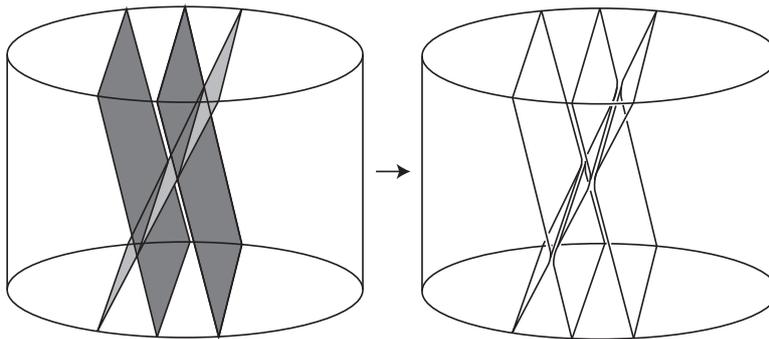


Figure 18: Summation of normal surfaces within a 1-handle

7.4. A VECTOR FOR BOUNDARY-REGULATED SURFACES

We now develop an analogue of the results in the previous two subsections for boundary-regulated surfaces.

For each type of elementary disc D in a 0-handle of \mathcal{H} , we introduce either one or two variables, depending on whether D is disjoint from the boundary of M . When D intersects ∂M , the two different variables corresponding to D relate to the two possible transverse orientations on D .

When S is a boundary-regulated surface, the associated *boundary-regulated vector* $(S)_{\partial r}$ counts the number of elementary discs of each type, except that when elementary discs intersect ∂M , their orientations are taken into account, and therefore one obtains two non-negative integers for this disc type.

The boundary-regulated vector satisfies a collection of equations, which we again term the *matching equations*. These fall into two types. Equations of the first type are just the classical matching equations, viewing S as an unoriented standard surface. Equations of the second type are concerned with the transverse orientations on the intersection arcs between ∂S and the handles. For each component of $\mathcal{H}^1 \cap \partial M$, the curves ∂S intersect this component in a collection of parallel arcs. These arcs have two possible transverse orientations. So, for this component of $\mathcal{H}^1 \cap \partial M$, we obtain two equations, one for each of these two transverse orientations, as follows. In the two adjacent 0-handles of \mathcal{H} , the elementary discs of S that are incident to this component of $\mathcal{H}^1 \cap \partial M$ are transversely oriented. So, the two equations assert that the number of arcs with some transverse orientation, arising from elementary discs in each of the two adjacent 0-handles, agree.

There are also *compatibility constraints*. Again, these come in two types. The constraints of the first type are just like those from classical normal surface theory. They assert that, for each pair of elementary disc types in a 0-handle, only one of these two types can occur if two discs of these types inevitably intersect. The second type of compatibility constraint is concerned with transverse orientations. For each 0-handle of \mathcal{H} and for each pair of transversely oriented elementary disc types that intersect ∂M and that lie within that 0-handle, the co-ordinate of one of these is forced to be zero if any two disjoint discs of these types inevitably give rise to a tubing arc α , as in Condition 2 of Section 5.2.

Let v be a vector with non-negative integer entries satisfying the boundary-regulated matching equations and compatibility constraints. Then we may construct a boundary-regulated surface S with $(S)_{\partial r} = v$, as follows.

For each elementary disc type in the 0-handles, insert into the 0-handle as many elementary discs of that type as specified by the vector v . At this stage, we do not attempt to orient them. Since the vector v satisfies the first type of compatibility constraint, there is a normal isotopy that makes any two of these elementary discs disjoint. Moreover, all of the discs may be made disjoint simultaneously, and there is a unique position, up to normal isotopy, for their union. We now impose the transverse orientations on the elementary discs in the 0-handles that intersect ∂M , as specified by the vector v . For each unoriented disc type, there is at most one way of doing this without creating a tubing arc. More precisely, there can be at most one switch of transverse orientations between parallel discs of the same type. Moreover, the switch can occur in just one way. For example, suppose that the discs of type D run over $R_-(M) \cap \mathcal{H}^1$. Then, at a switch of orientations between two discs of this type, their transverse orientations point away from each other. Since v satisfies the compatibility constraints, there is therefore a unique way of imposing these transverse orientations without creating a tubing arc.

We now insert the elementary discs into the 1-handles. The method of doing this is described in Section 7.3. Since v satisfies the unoriented matching equations, one may insert unoriented discs into each 1-handle $D^1 \times D^2$ to interpolate between the elementary discs in the incident 0-handles. As in Section 7.3, each of these discs intersects $\{p\} \times D^2$ in a properly embedded arc, for each point p in D^1 . If any of these elementary discs in a 1-handle intersects ∂M , then a transverse orientation may be imposed upon it, so that, near ∂M , this transverse orientation agrees with the transverse orientations on the two elementary discs in the 0-handles to which it is incident. This is because v satisfies the second type of equation, which is concerned with the transverse orientation on arcs in $\mathcal{H}^1 \cap \partial M$.

Finally, elementary discs may be inserted into the 2-handles. For each 2-handle $D^2 \times D^1$, the surface that has been constructed so far intersects the 2-handle in a collection of circles, each isotopic to $D^2 \times \{p\}$ for some point p in D^1 . Hence, these curves may be extended to properly embedded disjoint discs in the

2-handle. No transverse orientations are imposed upon these discs, because they are disjoint from ∂M .

This procedure creates a surface S . As explained in Section 7.3, this is not yet boundary-regulated, simply because it is not standard. However, there is a normal isotopy, after which S respects the product structures of the 1-handles and 2-handles, and is therefore a standard surface. This is the boundary-regulated surface that is required. We have therefore proved the following lemma.

Lemma 7.1. *Let v be a non-negative integral solution to the boundary-regulated equations, satisfying the compatibility constraints. Then there is some properly embedded, boundary-regulated surface S such that $(S)_{\partial r} = v$.*

7.5. SUMMATION OF BOUNDARY-REGULATED SURFACES

In this subsection, we consider the case where S is a boundary-regulated surface, and its vector $(S)_{\partial r}$ can be expressed as a sum $v_1 + v_2$, where v_1 and v_2 are non-negative and integral and satisfy the boundary-regulated matching equations. Our aim is first to show that there are boundary regulated surfaces S_1 and S_2 such that $v_1 = (S_1)_{\partial r}$ and $v_2 = (S_2)_{\partial r}$. Then we will show how this summation may be interpreted topologically.

Note that v_1 and v_2 also satisfy the boundary-regulated compatibility constraints, for the following reasons. If v_1 , say, contains two disc types in a 0-handle that inevitably intersect, then the same is true for $(S)_{\partial r}$, which is impossible. Suppose that there is a tubing arc α joining two elementary discs that are present in v_1 . Then these discs are also present in S . So, some subset of α forms a tubing arc for S , which contradicts the fact that it boundary-regulated. Hence, by Lemma 7.1, v_1 and v_2 do correspond to boundary-regulated surfaces S_1 and S_2 . Of course, $(S)_{\partial r} = (S_1)_{\partial r} + (S_2)_{\partial r}$.

We wish to develop a geometric interpretation for this summation. Again consider the elementary discs of S in each 0-handle. Since $(S)_{\partial r} = (S_1)_{\partial r} + (S_2)_{\partial r}$, we may partition these discs into two collections, so that the number of elementary discs of each type in the first (respectively, second) collection is equal to the number of elementary normal discs of that type in S_1 (respectively, S_2). Since S_1 satisfies the matching equations, it is possible to build S_1 starting from these elementary discs by inserting discs into the 1-handles and 2-handles. We build S_2 similarly.

Note that, exactly as in Section 7.3, S is obtained, up to normal isotopy, by cutting $S_1 \cup S_2$ along $S_1 \cap S_2$ and then resolving these intersections. We borrow the terminology from the normal case, by using the phrases *regular switch*, *irregular switch*, *trace curves* and *patches* in the obvious way.

Note that the boundary-regulated matching equations take account of the transverse orientations of the elementary discs that intersect ∂M . Hence, we deduce that, in each component D of $\partial M \cap \mathcal{H}^1$, the arcs $D \cap S_1$ and $D \cap S_2$ are compatibly oriented whenever they intersect. (See Figure 19). We therefore deduce the following result.

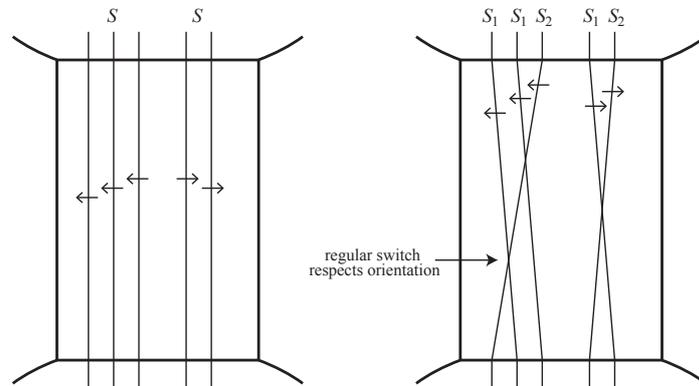


Figure 19: A component of $\mathcal{H}^1 \cap \partial M$

Lemma 7.2. *Let S be a properly embedded, boundary-regulated surface. Let $(S)_{\partial r}$ be its boundary-regulated vector. Suppose that $(S)_{\partial r} = (S_1)_{\partial r} + (S_2)_{\partial r}$ for boundary-regulated surfaces S_1 and S_2 . Then for each properly embedded arc of $S_1 \cap S_2$, the regular switch along this arc respects the transverse orientations of ∂S_1 and ∂S_2 near that arc.*

Note that the lemma only refers to arcs of $S_1 \cap S_2$, not simple closed curves.

A properly embedded boundary-regulated surface S is *fundamental* if it cannot be written as a sum of two non-empty boundary-regulated surfaces.

7.6. AVOIDING TRIVIAL TRACE ANNULI

A large part of Section 4 was concerned with decompositions along annuli disjoint from the sutures. One of the reasons for this is that trace annuli are of this form. However, it will be important that these annuli are non-trivial. In the next few sections, we will prove the following proposition, which implies that trivial trace annuli may be avoided.

Proposition 7.3. *Let \mathcal{H} be a handle decomposition of a compact orientable sutured manifold (M, γ) . Let S be a connected regulated surface properly embedded in M . Suppose that S extends to an allowable hierarchy and is not product-separating. Assume that at least one of the following holds:*

- (i) S is non-separating in the component of M that contains it, or
- (ii) (M, γ) contains a non-separating product disc or a non-separating product annulus.

Suppose also that S has smallest possible complexity among connected standard surfaces with these properties. Let (M_S, γ_S) be the result of decomposing along S . Suppose that S is a sum of non-empty boundary-regulated surfaces S_1 and S_2 . Suppose also that $S_1 \cap S_2$ has fewest components, among all ways of expressing S as such a sum. Then no trace annulus for this summation is trivial in (M_S, γ_S) .

7.7. GENERALISED SUMMATION

We say that a properly embedded surface S in an orientable 3-manifold is a *generalised sum* of two properly embedded surfaces S_1 and S_2 if S_1 and S_2 are in general position, and hence $S_1 \cap S_2$ is a collection of properly embedded arcs and simple closed curves, and S is obtained from $S_1 \cup S_2$ by resolving these intersections in some way. We write $S = S_1 + S_2$, although there is more than one surface that can be obtained from S_1 and S_2 by generalised summation.

Generalised summation arises in several possible ways. Summation of normal surfaces, as described in Section 7.3, is a type of generalised summation. Similarly, summation of boundary-regulated surfaces is also. Another natural method of performing generalised summation is when S_1 and S_2 are oriented and the resolution of $S_1 \cap S_2$ occurs in the way that respects this orientation.

We again borrow terminology from normal surface theory, by speaking of patches, trace curves, and trace discs, annuli and Möbius bands. A trace curve is *two-sided* if the I -bundle over it with (∂I) -bundle in S is orientable; otherwise it is *one-sided*. Thus, trace annuli and trace discs arise from two-sided trace curves, and trace Möbius bands arise from one-sided trace curves.

Consider a trace arc or curve a of $S_1 \cap S_2$. Suppose the four patches incident to a are all distinct. We say that two of these patches are *opposing along a* if one lies in S_1 , the other lies in S_2 , and the two copies of these patches in S do not intersect along a . Thus, the four patches incident to a decompose into two opposing pairs.

7.8. REDUCIBLE SUMMATION

The summation $S = S_1 + S_2$ is *reducible* if there are subsurfaces $P_1 \subset S_1$ and $P_2 \subset S_2$ which are each a union of patches, such that no two patches of P_1 are incident in S_1 , no two patches of P_2 are incident in S_2 , and along every arc and curve of $\partial P_1 - \partial S_1$ and along every arc and curve of $\partial P_2 - \partial S_2$, P_1 and P_2 are opposed.

Suppose that the summation is reducible. Then one may attach P_2 to $S_1 - P_1$, and attach P_1 to $S_2 - P_2$, giving two surfaces S'_1 and S'_2 . Then $S = S'_1 + S'_2$, but $|S'_1 \cap S'_2| < |S_1 \cap S_2|$.

Note that if the summation $S_1 + S_2$ arose from the summation of normal or boundary-regulated surfaces, then so does the summation $S'_1 + S'_2$. So, if one considers an expression of S as a normal or boundary-regulated sum $S_1 + S_2$ of non-empty surfaces, and $|S_1 \cap S_2|$ is minimised, then the summation is not reducible.

7.9. ALTERNATIVE SUMMANDS

Let S be a generalised sum $S_1 + S_2$. Then an *alternative summand* for S is a surface F that is a union of patches of $S_1 \cup S_2$ such that, for each trace arc and curve of $S_1 \cap S_2$, exactly two of the four patches emanating from this curve lie in F , and these are not opposing.

The union of the patches that do not lie in F forms another surface F' . Note that S is also the generalised sum $F + F'$. Moreover, each of the trace arcs and curves for $F + F'$ is a trace arc or curve for $S_1 + S_2$, and the way that these are resolved is the same. Hence if $S = S_1 + S_2$ is a normal (or boundary-regulated) sum of normal (or boundary-regulated) surfaces, then F and F' are normal (or boundary-regulated) and the summation is normal (or boundary-regulated).

7.10. IRREGULAR SWITCHES

Let S be a generalised sum $S_1 + S_2$. Then a surface is obtained from $S_1 \cup S_2$ by making some *irregular switches* if it is obtained from $S_1 \cup S_2$ by resolving the intersections in some way, and at least one of these arcs or curves, the way that it resolved is different from that of S .

7.11. TUBING AND SUMMATION

Let S be a generalised sum $S_1 + S_2$. Suppose that S_1 and S_2 have been given a transverse orientation in a regular neighbourhood of ∂M and that the summation respects this orientation. Then S inherits a well-defined transverse orientation in a regular neighbourhood of ∂M . The main case where this hypothesis holds is in the case of summation of boundary-regulated surfaces.

Let α be the closure of an arc component of $\partial S_1 - \partial S_2$. Suppose that α is a tubing arc for S_2 . Thus, α is disjoint from the sutures γ , and the transverse orientations of S_2 at the endpoints of α and the transverse orientation of $R_{\pm}(M)$ near α all agree.

Let S_2^+ be the result of attaching a tube to S_2 along α . Then $S_1 \cap S_2^+$ is obtained from $S_1 \cap S_2$ by modifying it near α , in the following way. The arcs of $S_1 \cap S_2$ at the endpoints of α are truncated just before ∂M , and then a parallel copy of α is attached to them.

One may perform a corresponding sum $S_1 + S_2^+$. For the arcs and curves of $S_1 \cap S_2^+$ that are equal to an arc or curve of $S_1 \cap S_2$, the way that these are resolved is unchanged. For the new arc or curve of $S_1 \cap S_2^+$, the resolution is the one that agrees with the resolution of the arcs of $S_1 \cap S_2$ at the endpoints of α . The resulting surface S^+ is obtained from S by attaching a tube that runs parallel to α .

Note that S_1 and S_2^+ also have a well-defined transverse orientation in a regular neighbourhood of ∂M and the summation $S_1 + S_2^+$ respects this orientation.

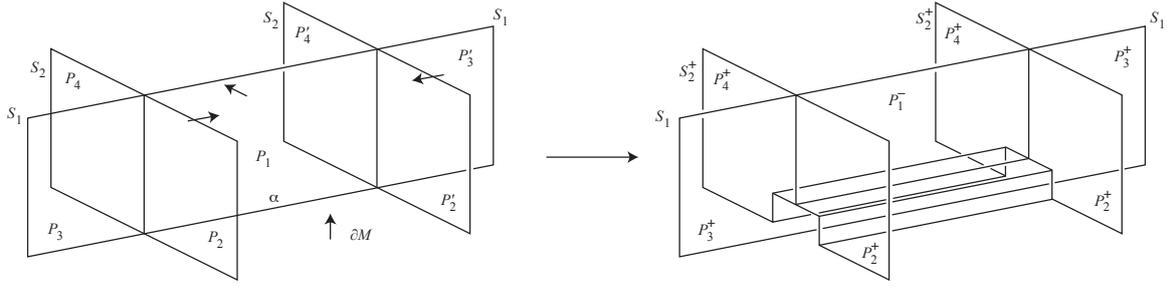


Figure 20: Attaching a tube and summation

Lemma 7.4. *If the summation $S_1 + S_2^+$ is reducible, then so was the summation $S_1 + S_2$.*

Proof. Figure 20 shows the patches of $S_1 + S_2$ near α . These are labelled $P_1, P_2, P_3, P_4, P'_2, P'_3$ and P'_4 . It is possible that some of these patches may actually be the same patch. Up to changing all transverse orientations near ∂M and up to symmetry, we may assume that the transverse orientations are also as shown in Figure 20. Thus, the following form opposing pairs: P_1 and P_2 ; P_3 and P_4 ; P_1 and P'_2 ; P'_3 and P'_4 .

The patches of $S_1 + S_2^+$ can also be read off from Figure 20. One patch P_3^+ is obtained from $P_3 \cup P'_3$ by attaching a band along α . This opposes a similar patch P_4^+ obtained from $P_4 \cup P'_4$. There is a patch P_1^- obtained from P_1 by removing a regular neighbourhood of α . This opposes a patch P_2^+ obtained from $P_2 \cup P'_2$ obtained by attach a band along α .

Suppose that $S_1 + S_2^+$ is reducible. If none of the patches near α form the specified subsurfaces of S_1 and S_2^+ , then the corresponding subsurfaces of S_1 and S_2 give that $S_1 + S_2$ is reducible. If the patch P_2^+ forms part of one of the subsurfaces, then P_1^- must form part of the other subsurface, and then we obtain similar subsurfaces of S_1 and S_2 giving reducibility there. A similar argument holds when P_3^+ and P_4^+ form part of the subsurfaces. \square

Lemma 7.5. *If the surface S^+ has an alternative summand F^+ , then S has a corresponding alternative summand F . Moreover, F is either isotopic to F^+ or is obtained from F^+ by performing a boundary-compression along a product disc disjoint from γ .*

Proof. Consider the right-hand diagram in Figure 20. Along the arc or curve of $S_1 \cap S_2^+$ shown, exactly two patches lie in F^+ . The possibilities are P_1^- and P_3^+ ; P_2^+ and P_4^+ ; P_2^+ and P_3^+ ; P_1^- and P_4^+ . In each case, we may define F to be the union of the corresponding patches for $S_1 \cup S_2$. For example, in the former case, we take P_1, P_3 and P'_3 together with all the other patches of F^+ . In this case, F is equal to F^+ . But in two of the other three cases, F is obtained from F^+ by boundary-compressing along the obvious product disc disjoint from γ . \square

A very similar argument gives the following result. Its proof is omitted.

Lemma 7.6. *If the surface T^+ is obtained from S^+ by performing irregular switches, then there is a corresponding surface T obtained from S by performing irregular switches. If T^+ and S^+ can be transversely oriented so that the transverse orientations agree on each patch of S^+ and T^+ and agree with the given transverse orientations near ∂M , then T and S can also be transversely oriented so that the transverse orientations agree on each patch of S and T and agree with the given transverse orientations near ∂M . Moreover, in this case, T is obtained from T^+ by boundary-compressing along a product disc disjoint from γ .*

7.12. TRIVIAL PATCHES

Let $S = S_1 + S_2$ be a generalised summation. Then a patch P is *trivial* if it is a planar surface, with one boundary curve forming a trace curve of $S_1 \cap S_2$ and the remaining boundary curves being trivial boundary curves of S .

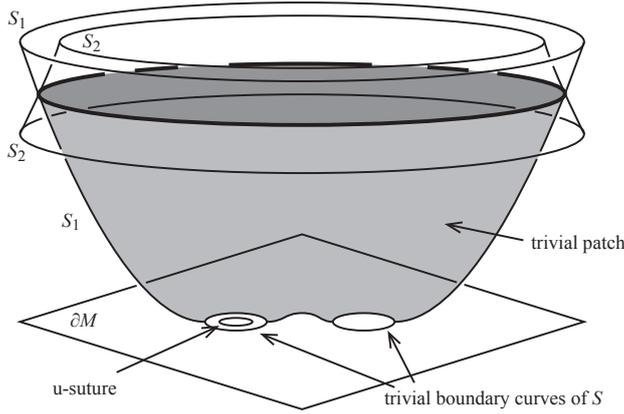


Figure 21: Trivial patch

Suppose that

$$(M, \gamma) \xrightarrow{S} (M_S, \gamma_S)$$

is a taut allowable decomposition between decorated sutured manifolds. Suppose also that S is connected. One would like to say that if there is a trace annulus A that is trivial in (M_S, γ_S) , then there has to be a trivial patch. However, this need not obviously be the case. Since A is trivial, one of its boundary curves bounds a disc in $E(M_S, \gamma_S)$. The restriction of this disc to M_S is a planar surface, and the intersection between this and S is a planar surface \tilde{P} . The complication is that \tilde{P} may contain trace curves and trace arcs in its interior. The possible presence of trace arcs in \tilde{P} is problematic, because an ‘innermost’ patch P in \tilde{P} does not necessarily form a trivial patch. An example is given in Figure 22.

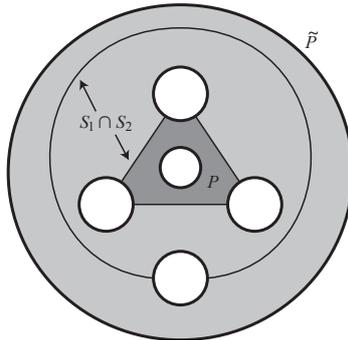


Figure 22: The ‘innermost’ planar surface P is not a trivial patch

7.13. PROOF OF PROPOSITION 7.3.

We now embark on this proof. We use the terminology of the proposition. We suppose for contradiction that (M_S, γ_S) contains a trivial trace annulus.

We are assuming that S is a sum of non-empty boundary-regulated surfaces S_1 and S_2 , and where $|S_1 \cap S_2|$ is minimal. Therefore, the summation is not reducible.

The plan of the proof is to follow the argument of Lemma 2.1 in [13] where, in the case of normal summation, disc patches in the interior of M were excluded. However, the difficulties described in Section 7.12 relating to innermost patches cause complications. To circumvent these, we introduce a new surface S' which will be a generalised sum $S_1 + S'_2$. The problematic innermost patches, as described in Figure 22, will not arise in this summation. We can then apply the argument of Lemma 2.1 in [13]. We then use the lemmas in Section 7.11 to provide information about the original summation $S = S_1 + S_2$. This will lead to a contradiction.

The new properly embedded surface S' will be obtained from S by attaching tubes along some arcs of $\partial S_1 - \partial S_2$. These will be tubing arcs. Because of Lemma 7.2, the summation at the endpoints of these arcs satisfies the conditions of Section 7.11. Hence, this new surface S' will be a generalised summation $S_1 + S'_2$. By Lemma 7.4, this new summation will not be reducible. By Lemma 4.5, S' will extend to an allowable hierarchy. Also, by Lemma 5.6 or Proposition 5.7, S' will not be product-separating. We denote the manifold obtained by decomposing along S' by $(M_{S'}, \gamma_{S'})$. Since this manifold is at the start of an allowable hierarchy, it is taut by Corollary 3.7, its canonical extension is taut by Lemma 3.3, and $E(S')$ is taut by Lemma 3.3.

This surface S' is constructed as follows. Consider any trivial curve c of ∂S that is neither a component of ∂S_1 nor of ∂S_2 . Thus, ∂S is comprised of arcs of $\partial S_1 - \partial S_2$ and arcs of $\partial S_2 - \partial S_1$ attached to each other. Since c is trivial, it bounds a disc in $\partial E(M, \gamma)$ disjoint from the sutures. By choosing c appropriately, we may assume that if any curve of ∂S lies in the interior of this disc, then it is a component of ∂S_1 or ∂S_2 . This implies that every arc of $\partial S_1 - \partial S_2$ and $\partial S_2 - \partial S_1$ in c is a tubing arc. Pick all such arcs of $\partial S_1 - \partial S_2$ in c , and attach a tube to S_2 along them. The new surface is a generalised sum. Continuing in this way as far as possible, we end with a surface S' which is a generalised sum of surfaces S_1 and S'_2 . These have the property that any trivial curve of $\partial S'$ is a component of ∂S_1 or $\partial S'_2$.

This produces new trace curves, all of which are two-sided, and hence new trace annuli. Some of these may be trivial in $(M_{S'}, \gamma_{S'})$. But the trivial trace annuli in (M_S, γ_S) remain trivial trace annuli in $(M_{S'}, \gamma_{S'})$, by Lemma 4.6. So, there is certainly at least one trivial trace annulus in $(M_{S'}, \gamma_{S'})$. One of its boundary curves bounds a disc in $E(M_{S'}, \gamma_{S'})$. The restriction of this disc to $M_{S'}$ is a planar surface, and the intersection between this and S' is a planar surface \tilde{P} . This is a union of patches. Any boundary component of \tilde{P} other than the one lying in the trace annulus is a trivial curve of $\partial S'$. Hence, it lies entirely in ∂S_1 or $\partial S'_2$. Therefore, there are only trace curves in \tilde{P} , and no trace arcs. Consider one that is innermost in \tilde{P} . This encloses a patch P_1 . This extends to a disc D_1 in $\partial E(M_{S'}, \gamma_{S'})$ disjoint from the sutures. Now consider the trace annulus A incident to P_1 . The other boundary component of A is also a trivial curve in $\partial M_{S'}$. This is because $E(M_{S'}, \gamma_{S'})$ is taut by Lemma 3.3. Hence, this component of ∂A also bounds a disc \tilde{D}_1 in $\partial E(M_{S'}, \gamma_{S'})$ disjoint from the sutures. The intersection between this disc and S' is a planar surface \tilde{P}_1 that is a union of patches. We say that \tilde{P}_1 is *associated* with the patch P_1 .

There are several cases to consider. Suppose first that the patch in \tilde{P}_1 incident to ∂A is not opposing P_1 along the trace curve. If \tilde{D}_1 is disjoint from D_1 , then $\tilde{D}_1 \cup D_1$ forms a 2-sphere, which, after a small isotopy, can be made disjoint from $E(S')$ and which separates $E(S')$. This contradicts the fact $E(M_{S'})$ is irreducible or that $E(S')$ is incompressible or that $E(S')$ has no 2-sphere components. So, suppose now that \tilde{D}_1 and D_1 are not disjoint. Since P_1 is a patch, and \tilde{P}_1 is made from a union of patches, we deduce that $P_1 \subset \tilde{P}_1$. In this case, we remove \tilde{P}_1 from S' and replace it by P_1 . Let S'' denote the new surface. This is obtained from $S_1 \cup S'_2$ by performing an irregular switch along ∂D_1 , giving a surface F' and then removing the punctured torus formed from $\tilde{P}_1 - P_1$. Note that S'' is obtained from S' by slicing under discs of contact and then performing an isotopy. Hence, S'' also extends to an allowable hierarchy by Lemma 4.3. If S was non-separating, then so is S'' . The surface F' can be transversely oriented so that the patches of F' and the patches of S' have the same transverse orientations. So, by Lemma 7.6, there is a corresponding surface F obtained from $S_1 \cup S_2$ by performing an irregular switch. It is isotopic to a standard surface with complexity strictly less than that of S . By Lemma 7.6, this is obtained from F' by possibly performing some boundary compressions along product discs disjoint from γ . The surface S'' is obtained from F' by removing a component. This component corresponds to a component or components of F . Remove these to give a surface S''' . Then S''' is obtained from S'' by possibly performing boundary compressions along product discs disjoint from γ . Hence, by Proposition 4.7, S''' extends to an allowable hierarchy. Since it is a union of components of F , its complexity is at most that of F , and this is (after an isotopy) strictly less than that of S . If S'' was non-separating, then so is at least one component of S''' . Thus, we contradict our assumption that S had minimal complexity.

Now consider the case where \tilde{P}_1 and P_1 are opposing. Then \tilde{P}_1 cannot be a single patch because the summation $S_1 + S'_2$ would then be reducible, contradicting Lemma 7.4. Hence, \tilde{P}_1 is a union of patches.

Again, we may find one that is innermost in \tilde{P}_1 . Call this patch P_2 . Since there are not trace arcs in \tilde{P}_1 , P_2 is a trivial patch. Continuing in this way, we obtain trivial patches P_1, \dots, P_n and associated planar surfaces $\tilde{P}_1, \dots, \tilde{P}_n$. We stop the process when \tilde{P}_n contains one of the previous trivial patches in its interior. By relabelling, we may assume that this trivial patch is P_1 .

We can view the surface $\tilde{P}_i - P_{i+1}$ (where the index i is modulo n) as a punctured annulus. As explained in the proof of Lemma 2.1 of [13], these punctured annuli combine to form a surface T' , which is a punctured torus. This is an alternative summand for S' . Thus $S' = T' + F'$, for some properly embedded surface F' . We can now fill in each of these punctures of T' , by attaching discs of contact, and the result is a torus \overline{T}' . We can fill in the corresponding discs of contact for S' , giving a surface \overline{S}' , which also extends to an allowable hierarchy by Lemma 4.3. It is not product-separating by Proposition 5.7. Then $\overline{S}' = \overline{T}' + F'$. Hence, F' is in fact ambient isotopic to \overline{S}' . Therefore, F' also extends to an allowable hierarchy with the same reduced length and is not product-separating. Since F' is an alternative summand for S' , Lemma 7.5 implies that S also has a corresponding alternative summand F . By Lemma 7.5, this is obtained from F' by performing boundary-compressions along product discs disjoint from γ . So by Proposition 4.7, F also extends to an allowable hierarchy. By Proposition 5.7, F is not product-separating. If S was non-separating, so too is at least one component of F . Now the summation $S_1 + S_2$ was boundary-regulated, and so the alternative summand F is also boundary-regulated. Since it is a boundary-regulated summand, its complexity is strictly less than that of S . But this contradicts the assumption that S had minimal complexity. \square

The surface S' constructed in the above proof will be useful to us later, and so we record some of its properties now.

Addendum 7.7. *Let S, S_1 and S_2 be as in Proposition 7.3. Then there is a surface S' properly embedded in (M, γ) with the following properties:*

- (i) S' is a generalised sum of S_1 and a surface S'_2 ;
- (ii) when S_1 and S_2 have transverse orientations and S is the oriented double-curve sum of S_1 and S_2 , then S'_2 also has a transverse orientation and S' is the oriented double-curve sum of S_1 and S'_2 ;
- (iii) no trace annulus for this summation is trivial in the sutured manifold $(M_{S'}, \gamma_{S'})$ that is obtained by decomposing along S' ;
- (iv) S' is obtained from S by tubing along arcs;
- (v) every trivial curve of $\partial S'$ is a component of ∂S_1 or $\partial S'_2$.

7.14. DECOMPOSITION ALONG FUNDAMENTAL SURFACES

We now come to a central result of this paper.

Theorem 7.8. *Let (M, γ) be a taut decorated sutured manifold with a positive handle structure \mathcal{H} . Suppose that (M, γ) admits an allowable hierarchy. Suppose also that $E(M, \gamma)$ is atoroidal, and that no component of $E(M, \gamma)$ is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Suppose that no component of (M, γ) has boundary a single torus with no sutures. In addition, suppose that (M, γ) is not a product sutured manifold. Then it admits taut allowable sutured manifold decomposition along a fundamental boundary-regulated surface S , which extends to an allowable hierarchy, and such that decomposition along S reduces the complexity of the handle structure.*

Proof. The proof divides into two cases.

Case 1. (M, γ) admits an allowable hierarchy where the first surface S is connected and non-separating.

We may assume that S is incompressible by Lemma 4.2, and so by Lemma 4.9 in [20], it can be placed in standard form. By Theorem 5.5, we may further assume that S is regulated, as long as we

drop the requirement that S is incompressible. We pick S to have minimal complexity among regulated non-separating surfaces that extend to an allowable hierarchy. Let (M_S, γ_S) be obtained from (M, γ) by decomposing along S .

Suppose that S is not fundamental as a boundary-regulated surface. Then it can be written as a sum of non-empty boundary-regulated surfaces S_1 and S_2 . We may perform a normal isotopy to S_1 and S_2 , leaving γ invariant, so that they intersect in a collection of simple closed curves and properly embedded arcs. We may assume that S_1 and S_2 have been chosen so that the number of these curves and arcs is minimal.

Each patch of this summation inherits a transverse orientation from S .

Case 1A. These orientations induce transverse orientations of S_1 and S_2 .

Then S is the oriented double-curve sum of S_1 and S_2 . Since the summation respects the transverse orientations on S_1 and S_2 , we deduce that $[S, \partial S] = [S_1, \partial S_1] + [S_2, \partial S_2] \in H_2(M, \partial M)$. Therefore, at least one of $[S_1, \partial S_1]$ and $[S_2, \partial S_2]$ is non-trivial, say the former. We will show that S_1 extends to an allowable hierarchy. Since S_1 is oriented and boundary-regulated, it is regulated. But we have then found a non-separating, regulated surface S_1 that extends to an allowable hierarchy, but with smaller complexity than that of S . This contradicts our minimality assumption.

Note first that S_1 is connected. For if it is not, then it has a homologically non-trivial component F . We can then instead consider the summation $S = F + (S_2 + (S_1 - F))$, where $S_2 + (S_1 - F)$ is the oriented double-curve sum of S_2 and $S_1 - F$. There are fewer curves of intersection between F and $S_2 + (S_1 - F)$ than between S_1 and S_2 , contradicting our minimality assumption.

By Addendum 7.7, we have a surface S' as in the statement there. This is the oriented double-curve sum of S_1 and a surface S'_2 . We claim that the trivial boundary curves of S' are precisely the trivial curves of ∂S_1 and $\partial S'_2$. Certainly, every trivial curve of $\partial S'$ arises as a trivial curve of ∂S_1 or $\partial S'_2$, by (v) of Addendum 7.7. To prove the claim, we must show that every trivial curve of ∂S_1 and $\partial S'_2$ is preserved as a curve of ∂S . We will show this for curves of ∂S_1 , as the other case is similar. Consider a trivial curve C of ∂S_1 and its trivialising planar surface F . The intersection between F and $\partial S'_2$ is a collection of simple closed and arcs. Consider just the arcs, which all start and end on C , because the remaining components of ∂F are u-sutures of γ . Transversely orient the arcs according to the transverse orientation of S'_2 . We may construct a tree, where each vertex corresponds to a component of the complement of these arcs. This has a sink vertex and a source vertex. One of these corresponds to a trivialising planar surface for a component of $\partial S'$. This contradicts (v), and so the claim is proved.

Let P be the trace discs, annuli and Möbius bands arising from the summation $S' = S_1 + S'_2$. In fact, because the summation respects the transverse orientations of S_1 and S'_2 , P has no Möbius band components, because incident to each component of P , there are parts of S' pointing away from P and parts of S' pointing towards P . Give P some transverse orientation. These then form a collection of product discs and product annuli. By Addendum 7.7, the annuli are all non-trivial in $(M_{S'}, \gamma_{S'})$. We have a commutative diagram of sutured manifold decompositions

$$\begin{array}{ccc} (M, \gamma) & \xrightarrow{S'} & (M_{S'}, \gamma_{S'}) \\ \downarrow S_1 & & \downarrow P \\ (M_1, \gamma_1) & \xrightarrow{S'_2 - \text{int}(N(S_1))} & (M_{12}, \gamma_{12}). \end{array}$$

We claim that any trivial boundary curve C of $S'_2 - \text{int}(N(S_1))$ is disjoint from S_1 . Suppose not and consider the trivialising planar surface P bounded by C . This is divided into subsurfaces $P \cap R_{\pm}(M)$ and $P \cap S_1$. Any component of $P \cap R_{\pm}(M)$ incident to C gives rise to a trivial boundary curve of S' that is not a component of ∂S_1 or $\partial S'_2$. This contradicts (v) in Addendum 7.7. Thus if C has non-empty intersection with S_1 , then it must lie entirely in S_1 . But it then gives rise to a trivial trace annulus in $(M_{S'}, \gamma_{S'})$, which contradicts Addendum 7.7. This proves the claim.

We claim that all the decompositions in the above commutative diagram are allowable. Since S' is obtained from S by tubing along arcs, decomposition along S' is allowable, by Lemma 4.5. The non-trivial annuli in P are all allowable. Moreover, the product discs in P are allowable, because they are disjoint from the u-sutures of $\gamma_{S'}$. This is because each such product disc lies within a regular neighbourhood of a trace arc and so is clearly disjoint from the u-sutures of γ . It is also disjoint from the trivial boundary curves of S' , as each trivial boundary curve of S' is a component of ∂S_1 or $\partial S'_2$. Note also that ∂S_1 is disjoint from the u-sutures of γ . Also, any trivial boundary curve of S_1 is a trivial boundary curve of S , by the above claim. Hence, $S'_2 - \text{int}(N(S_1))$ is disjoint from these u-sutures of γ_1 . Thus, we have shown that all of these surfaces are disjoint from the u-sutures. Finally, any trivial boundary curve of S_1 or $S'_2 - \text{int}(N(S_1))$ is a trivial curve of $\partial S'$, and so its trivialising planar surface is correctly oriented. This proves the claim.

Thus, the above diagram consists of decorated sutured manifolds. We show that the decoration on (M_{12}, γ_{12}) is the same, no matter which way one goes round the digram. If we decompose along S' then P , the u-sutures come from u-sutures of γ and trivial curves of $\partial S'$. Note that decomposition along P does not create any u-sutures because every annular component of P is non-trivial. If we decompose along S_1 then $S'_2 - \text{int}(N(S_1))$, then the u-sutures come from the u-sutures of γ and the trivial boundary curves of ∂S_1 and $\partial S'_2$, which are disjoint. Thus, the u-sutures of γ_{12} are the same in both cases.

By (iv) in Addendum 7.7, S' is obtained from S by tubing along arcs, and so by Lemma 4.5, S' extends to an allowable hierarchy. By Lemma 4.4 and Proposition 4.9, (M_{12}, γ_{12}) admits an allowable hierarchy.

However, this does not yet prove that S_1 extends to an allowable hierarchy, because it is not clear that S_1 and $S'_2 - \text{int}(N(S_1))$ can be part of an allowable hierarchy, since they may have components that are planar, disjoint from the sutures and where all but at most one boundary curve is trivial. In fact, S_1 cannot have such a component. Since S_1 is connected, this component would be all of S_1 . Its canonical extension would have to be a boundary parallel disc in $E(M, \gamma)$ by the tautness of $E(M, \gamma)$. Hence, it would be homologically trivial in M , which is a contradiction.

Now, $S'_2 - \text{int}(N(S_1))$ may possibly have a planar component F disjoint from the sutures and where all but at most one boundary curve is trivial. We now show how to remove such a component F .

Since (M_{12}, γ_{12}) admits an allowable hierarchy, $E(M_{12}, \gamma_{12})$ is taut by Lemma 3.3. We have a decomposition

$$E(M_1, \gamma_1) \xrightarrow{E(S'_2 - \text{int}(N(S_1)))} E(M_{12}, \gamma_{12}),$$

which we will show is taut using Theorem 2.3. No boundary component of $E(S'_2 - \text{int}(N(S_1)))$ bounds a disc in $E(M_1, \gamma_1)$ disjoint from the sutures. Also, no component of $E(M_1, \gamma_1)$ can be a solid torus with no sutures. This is because this would be a component of (M_1, γ_1) , and this would imply that (M, γ) is not taut or that S_1 is homologically trivial. Thus, Theorem 2.3 implies that the above decomposition is taut and that $E(S'_2 - \text{int}(N(S_1)))$ is taut. This implies that $E(F)$ is a boundary-parallel disc in $E(M_1, \gamma_1)$. By choosing F appropriately, we can assume that no component of $S'_2 - \text{int}(N(S_1))$ lies in the product region that $E(F)$ separates off. Hence, $E(F)$ separates off a ball component of $E(M_{12}, \gamma_{12})$. This gives a pre-ball component of (M_{12}, γ_{12}) , by Lemma 3.9. Hence, F is product-separating in (M_1, γ_1) . We remove F , as follows. We first slice under any discs of contact, creating a surface \bar{F} . The effect of this on (M_{12}, γ_{12}) is to remove some u-sutures and attach 2-handles along them. Hence, the resulting sutured manifold still admits an allowable hierarchy. But then \bar{F} is then boundary-parallel, by Lemma 4.10, separating off a pre-ball. Hence, we can remove it, and again the resulting sutured manifold admits an allowable hierarchy.

Thus, S_1 does extend to an allowable hierarchy, as required.

Case 1B. The orientations on the patches do not induce transverse orientations of S_1 and S_2 .

Then, for some trace curve C of $S_1 \cap S_2$, the orientations on S_1 or S_2 near that curve do not agree. By Lemma 7.2, C must be a simple closed curve, rather than an arc. So, there is an annulus or Möbius band A embedded in a regular neighbourhood of this curve such that $A \cap S = \partial S$. The transverse orientations

on S near ∂A either all point towards A or all point away from A . Let S' be the surface that results from performing an annular swap to S , using the annulus or Möbius band A . In other words, we perform an irregular switch along C . Orient S' so that the orientations of S' and S agree on their intersection. If A is an annulus, no component of ∂A bounds a disc in $E(S)$, by Proposition 7.3. If A is a Möbius band, ∂A also cannot bound a disc in $E(S)$, since this would give rise to an $\mathbb{R}P^3$ summand of M . So, by Proposition 4.8, either decomposition along S' extends to an allowable hierarchy or decomposition along $S' - S'_1$ does, where S'_1 is a component of S' that separates off a solid torus. Therefore, since S was non-separating, so too is at least one component of the new decomposing surface. Note that S and S' have the same complexity, but there is an ambient isotopy, leaving γ invariant, that makes S' standard, and reduces its complexity. This contradicts the fact that S had minimal complexity.

Case 2. For every allowable hierarchy for (M, γ) using connected decomposing surfaces, the first decomposing surface is separating.

By Proposition 4.11, we may assume that the first decomposing surface is not product-separating. As in Case 1, we may arrange, using Lemma 4.2 and Theorem 5.5, for S also to be regulated. We pick S to have minimal complexity among connected regulated surfaces that extend to an allowable hierarchy, and that are not product-separating. Suppose that S is not fundamental as a boundary-regulated surface. Then it can be written as a sum of non-empty boundary-regulated surfaces S_1 and S_2 . Again, as in Case 1, we may assume that this summation has no trivial trace annulus. Each patch inherits a transverse orientation from S .

Case 2A. These orientations induce transverse orientations of S_1 and S_2 .

As explained in Case 1A, both S_1 and S_2 extend to allowable hierarchies. We claim that at least one of $[S_1, \partial S_1]$ and $[S_2, \partial S_2]$ is non-trivial in $H_2(M, \partial M)$. This will contradict the hypothesis that there is no connected non-separating surface in M that extends to an allowable hierarchy.

Suppose that, contrary to the claim, both S_1 and S_2 are trivial in $H_2(M, \partial M)$. Pick a point p in M disjoint from $S_1 \cup S_2$. We define an integer associated to each component of $M - (S_1 \cup S_2)$, as follows. Pick an arc that starts at p and ends within some component R of $M - (S_1 \cup S_2)$. Make this arc transverse to S_1 and S_2 , and define the value of R to be the signed intersection number between this arc and $S_1 \cup S_2$. This is well-defined because both S_1 and S_2 are homologically trivial. Since S is homologically trivial, one can also define integers associated with the components of $M - S$. These are just the same as those for $M - (S_1 \cup S_2)$, except near $S_1 \cap S_2$, where components of $M - (S_1 \cup S_2)$ with the same value are glued together to form $M - S$. Now, as one crosses S_1 or S_2 , these integers change by 1. So, near an arc or curve of $S_1 \cap S_2$, these integers take three different values. But this is a contradiction, because, as S is connected, $M - S$ has only two components. Therefore, we deduce that $S_1 \cap S_2$ is empty. Hence, S is the disjoint union of S_1 and S_2 , and so is disconnected. This is contrary to hypothesis.

Case 2B. The orientations on the patches do not induce transverse orientations for S_1 and S_2 .

We would like to apply the same argument as in Case 1B, but there are some added complications. Again, there is some simple closed curve C of $S_1 \cap S_2$, such that the orientations on S_1 or S_2 near that curve do not agree. So, there is an annulus or Möbius band A embedded in a regular neighbourhood of this curve such that $A \cap S = \partial S$. The transverse orientations on S near ∂A either all point towards A or all point away from A . Let S' be the surface that results from performing an annular swap to S , using A . In other words, we perform an irregular switch along C . Let A' be the annulus or Möbius band such that applying an annular swap to S' along A' gives back S . Orient S' so that the orientations of S' and S agree on their intersection. If some component of S' separates off a solid torus with no sutures, call this S'_1 . Otherwise set $S'_1 = \emptyset$. Let $S'' = S' - S'_1$. Then, by Proposition 4.8, decomposition along S'' extends an allowable hierarchy. The difficulty is that each component of the new decomposing surface S'' may be product-separating. If this is not the case, then we have found a decomposing surface with smaller complexity than S , that is not product-separating and that extends to an allowable hierarchy. This contradicts our minimality assumption about S .

So, suppose that S'' is product separating. Our aim is to reach a contradiction.

We slice under any discs of contact for S' . This does not change the fact that S'' (which is a union of components of S') is product-separating. But by Lemma 4.10, each component of S'' is now boundary-parallel. Note that A' still satisfies $A' \cap S' = \partial A'$. When we perform an annular swap along A' , we get a surface that is obtained from S by slicing under discs of contact.

Case 2B(i). $S'' = S'$ and S' is connected.

Let $F \times I$ be the product sutured manifold that S' separates off. Suppose first that the interior of A' is disjoint from $F \times I$. If A' separates the component of $(M_{S'}, \gamma_{S'})$ that contains it, then when we apply the annular swap along A' , the resulting surface S is disconnected, which is contrary to assumption. So, A' is non-separating. Using the product structure on $F \times I$, extend A' into $F \times I$, to form a properly embedded annulus in M , disjoint from γ . This annulus is non-separating in M . But then this annulus extends to an allowable hierarchy, by Proposition 4.9. This contradicts our assumption that no non-separating surface extends to an allowable hierarchy.

Suppose now that A' lies in $F \times I$. Then it is a non-trivial annulus with both boundary components in the same component of $F \times \partial I$. It is therefore boundary parallel in $F \times I$. So, when we perform the annular swap to S' along A' , the resulting surface S is disconnected, which is contrary to assumption.

Case 2B(ii). $S'' = S'$ and S' is disconnected.

We are assuming that S'' is product-separating, and hence each component of S' separates off a product manifold. If these are disjoint from each other, then the interior of A' is also disjoint from them, and the argument of Case 2B(i) applies. So, suppose that the product sutured manifolds overlap. Denote these by $F_1 \times I$ and $F_2 \times I$, which are separated off by S'_1 and S'_2 respectively.

Suppose first that these are nested: say that $F_2 \times I$ is a subset of $F_1 \times I$. Say that S'_1 has transverse orientation pointing out of $F_1 \times I$. Then $R_+(F_1 \times I) - S'_1$ just consists of discs and annuli by Lemma 4.10. Since the transverse orientation on S'_1 points away from A' , so does the transverse orientation on S'_2 . So, it points into $F_2 \times I$. Hence, $R_+(F_2 \times I)$ is a subset of $R_+(F_1 \times I) - S'_1$. It is an essential subsurface of $R_+(F_1 \times I) - S'_1$, since otherwise S' has a disc of contact, but we have removed any of those. Also, $R_+(F_2 \times I)$ cannot be disc, because this implies that a curve of ∂A is trivial. Hence, S'_2 is an annulus, which is parallel to an annulus in $R_+(F_1 \times I)$. So, when we perform an annular swap to S' along A' , we deduce that S is disconnected, which is a contradiction.

Now consider the case where $F_1 \times I$ and $F_2 \times I$ are not nested. Then S'_2 is an incompressible surface in $F_1 \times I$, and hence it is parallel to a subsurface G of $\partial M \cap (F_1 \times I)$. Note that $\partial S'_2$ cannot intersect any of the disc components of $R_+(F_1 \times I) - S'_1$, because an innermost curve of intersection between one of these discs and $\partial S'_2$ would form a disc of contact, and an outermost arc of intersection between one of these discs and $\partial S'_2$ would separate off a disc component of $R_+(M_{S'}, \gamma_{S'})$, which would imply that S'_2 must be a disc. Hence, the entirety of G lies in $R_-(M)$, except possibly some annular collars on ∂G . Now, S'_2 also separates off a product sutured manifold from $F_2 \times I$, and by assumption, this does not lie wholly in $F_1 \times I$. Hence, its interior is disjoint from the product region between S'_2 and G . Therefore, the union of $F_2 \times I$ and this product region is a product sutured manifold, which is all of (M, γ) . This is contrary to assumption.

Case 2B(iii). S' has two components S'_1 and S'_2 , and one of these, S'_1 say, separates off a solid torus with no sutures.

So, $S'' = S'_2$ and this is product separating. Let V be the solid torus bounded by S'_1 .

We first consider the case where V lies in the product manifold $F \times I$ separated off by S'_2 . Remove a regular neighbourhood of $A' \cap \partial V$ from ∂V and attach the two copies of A' . The result is an annulus in the product manifold. Remove the parts of this annulus that intersect $F \times \partial I$. The result is a collection of annuli properly embedded in $F \times I$. Decomposing $F \times I$ along these annuli gives a product manifold. But

the resulting manifold is a component of (M_S, γ_S) . Hence, we deduce that S is product-separating, which is a contradiction.

So we now suppose that V is disjoint from the product manifold $F \times I$. Consider all the trace annuli with interiors lying in V . Suppose that there is at least one of these. Each one separates V , and so there is one that is outermost in V . This separates off a solid torus in V with interior that is disjoint from the trace annuli and from A . So, we may assume that the interior of V contains no trace annuli.

The solid torus V cannot intersect ∂M for the following reason. If it did, we could find an annulus in ∂V that runs between this intersection with ∂M up to A' . Then extend the annulus across the product region that S'_2 separates off. We obtain an annulus properly embedded in M disjoint from γ and also disjoint from S . If this annulus is separating, then because parts of S lie on both sides of it, we deduce that S is disconnected, which is a contradiction. On the other hand, if the annulus is non-separating, then we again have a contradiction to our assumption that no non-separating surface extends to an allowable hierarchy.

We may give V a Seifert fibration so that its intersection with A' is a fibre. Then V must have an exceptional fibre, as otherwise S is isotopic to S'_2 and we deduce that S is product-separating.

Now we have arranged that no trace annulus has interior that intersects V . However, other trace annuli may intersect V , because their interiors may lie outside of V but their boundaries lie on V . Suppose first that there is no such trace annulus. In other words, suppose that the only trace annulus to intersect V is A . The two parts of S in ∂V near A lie in S_1 and S_2 . So, as one goes around ∂V , the surface must switch from S_1 to S_2 . It cannot do so at a trace annulus, because by assumption no others intersect ∂V . So, we deduce that ∂V has non-empty intersection with ∂M , and we have already ruled this case out.

So, we may assume that at least one trace annulus intersects ∂V but has interior outside of V . This annulus therefore has interior disjoint from V and from the product region $F \times I$. Because of the orientations on S' , we deduce that each such trace annulus gives an irregular switch. So, we may concentrate instead on one of these annuli. The only problem case is when the switch gives a decomposing surface that separates off a solid torus V' with no sutures. We give V and V' Seifert fibrations in which the trace annuli are unions of fibres. These fibrations must each have a singular fibre. Since V and V' intersect along at least an essential annulus, we deduce that there is a Seifert fibre space embedded in M with base space a disc and with two singular fibres. We deduce that $E(M, \gamma)$ is either toroidal or has a Seifert fibred component which is neither a solid torus nor a copy of $T^2 \times I$. This is contrary to assumption. \square

The above argument justifies the introduction of boundary-regulated surfaces. It is critical that we control the complexity of our decomposing surfaces, by ensuring that they are fundamental. So, we are forced to consider the situation where a decomposing surface S is a sum $S_1 + S_2$. The problem is that this summation need not respect orientations; indeed S_1 and S_2 need not even be transversely orientable. This sort of issue was faced by Tollefson and Wang [27], who showed that if the summation does not respect orientations, then one should perform an irregular switch along a trace curve and then simplify the resulting surface S' by reducing its weight. The difficulty here is that it is not at all clear that the decomposition along S' extends to an allowable hierarchy or is just taut. If the trace curve is a simple closed curve, then Corollary 7.5 and Proposition 4.8 give this. But this argument breaks down when a trace curve is an arc. Fortunately, it is at this point that Lemma 7.2 intervenes. This asserts that when a surface is a sum, as a boundary-regulated surface, then the trace curves that intersect ∂M must respect the orientations on the surfaces.

8. BOUNDING THE COMPLEXITY OF NORMAL AND BOUNDARY-REGULATED SURFACES

We have been encoding normal surfaces and boundary-regulated surfaces using vectors, which count the number of elementary discs of each type in the surface. In this section, we find upper bounds on the weights of these surfaces, by exploiting algebraic methods. Much of the material here is fairly well known, and goes back to Haken [8] and Hass-Lagarias [9].

We consider the following general set-up. Suppose that A is an $m \times n$ matrix with integer entries. We will examine solutions to the system

$$Ax = 0 \tag{1}$$

where $x = (x_1, \dots, x_n)^T \in \mathbb{R}^n$, subject also to the inequalities

$$x_i \geq 0 \quad \text{for } i = 1, \dots, n. \tag{2}$$

Consider also a system of constraints, each of which is of the form

$$x_i = 0 \quad \text{or} \quad x_j = 0 \tag{3}$$

where $1 \leq i, j \leq n$. We call (3) the *compatibility constraints*.

We say that (1), (2) and (3) form a system Σ . We say that a solution $x \in \mathbb{R}^n$ to this system is *integral* if all its co-ordinates are integers. An integral solution x is *fundamental* if it cannot be written as $x = y + z$ where y and z are also integral solutions to the system, neither of which is zero.

An integral solution x to Σ is a *vertex* solution if

- (i) whenever $kx = y + z$, for some positive real number k and some solutions y and z , then both y and z are multiples of x ; and
- (ii) $x \neq ky$ for some integer $k \geq 2$ and some integral solution y .

One can interpret vertex solutions another way. It is well known that the set of real solutions to the system Σ forms a union P of finitely many convex polytopes. Each polytope is obtained as follows. For each of the compatibility constraints, which asserts that $x_i = 0$ or $x_j = 0$ for some pair of integers i and j between 1 and n , one makes a choice to fix x_i to be zero or to fix x_j to be zero. Once this choice is made, the resulting system is a collection of equations $Ax = 0$ together with inequalities $x_i \geq 0$ for all i , and equalities of the form $x_i = 0$ for certain values of i . The solution set of this restricted system is invariant under dilations based at the origin. Hence it forms a cone, with cone point being the origin. Each line through the origin in P intersects the plane $x_1 + \dots + x_n = 1$ in exactly one point. If we consider the subset of P that satisfies this extra normalising constraint, then we obtain a union of finitely many compact convex polytopes. A vertex solution is exactly a vertex of one of these polytopes, rescaled minimally so that its coefficients are integers.

Theorem 8.1. *Suppose that each row of A has ℓ^2 norm at most k . Then we have the following bounds.*

- (1) *For any vertex solution x to the system Σ , each co-ordinate of x is at most $n^{1/2}k^{n-1}$.*
- (2) *For any fundamental solution x to the system Σ , each co-ordinate of x is at most $n^{3/2}k^{n-1}$.*

Proof. (1) There is some positive real number k and some vector $y = (y_1, \dots, y_n)^T$ satisfying $y = kx$, such that y is the unique solution to the following system of linear equations. The equations include the equations $Ay = 0$. They also include extra equalities of the form $y_i = 0$ for some i between 1 and n . Finally, there is the normalising equation

$$y_1 + \dots + y_n = 1.$$

We can write this system as

$$\tilde{A}y = \tilde{b}$$

where the first m rows of \tilde{A} are A . The vector \tilde{b} has all co-ordinates zero except the last, which is 1. The solution y is unique and hence the kernel of \tilde{A} is zero. This implies that the rank of \tilde{A} is n . Hence, there is some collection of n rows of \tilde{A} with rank n . Let \hat{A} be this $n \times n$ matrix, and let \hat{b} be the corresponding n entries of \tilde{b} . Then y is the unique solution to

$$\hat{A}y = \hat{b}.$$

The $n \times n$ matrix \widehat{A} is invertible, and so

$$y = (\widehat{A})^{-1}\widehat{b}.$$

Now $(\widehat{A})^{-1} = (\det(\widehat{A}))^{-1}\text{adj}(\widehat{A})$, where $\text{adj}(\widehat{A})$ is the adjugate matrix for \widehat{A} . Hence, $\det(\widehat{A})y = \text{adj}(\widehat{A})\widehat{b}$ is an integral vector. Therefore, x is obtained from $\det(\widehat{A})y$ by scaling by a factor at most 1. We may therefore bound the entries of x by bounding the entries of $\det(\widehat{A})y$. Each entry of $\text{adj}(\widehat{A})$ is equal to the determinant of an $(n-1) \times (n-1)$ minor of \widehat{A} . The modulus of this determinant is at most the product of the ℓ^2 norms of its rows, and this is at most $n^{1/2}k^{n-1}$, for the following reason. Each row of \widehat{A} is either a row of A , a row consisting of a single 1 and the remaining entries zero, or consists entirely of 1s, in the case where it is the normalising equation. So, each co-ordinate of $\text{adj}(\widehat{A})\widehat{b}$ has modulus at most $n^{1/2}k^{n-1}$.

(2) As described above, the real solutions to the system Σ and the normalising condition $x_1 + \dots + x_n = 1$ form a union of convex compact polytopes. The vertices of the polytopes are multiples of the vertex solutions. Hence, any solution to Σ can be written as a linear combination $\sum \lambda_i v_i$, where $\lambda_i \geq 0$ for each i , and each v_i is a vertex solution. One may choose this expression so that at most n of the λ_i are non-zero. It is clear that if $x = \sum \lambda_i v_i$ is fundamental, then each $\lambda_i < 1$, since otherwise x has some v_i as a summand. Therefore, we deduce that each co-ordinate of x has modulus at most n times the maximal size of a co-ordinate for a vertex solution. So, by (1), each co-ordinate of x has modulus at most $n^{3/2}k^{n-1}$. \square

Theorem 8.2. *There is a universal computable constant c with the following property. Let (M, γ) be a sutured manifold, and let \mathcal{H} be a handle structure of uniform type. Then, for each fundamental boundary-regulated surface S properly embedded in M , the weight of S is at most c^h , where h is the number of 0-handles of \mathcal{H} .*

Proof. Each boundary-regulated surface can be described as an integral solution to a system Σ as above. The number of variables n is equal to the number of disc types in the 0-handles. Since \mathcal{H} is of uniform type, the number of disc types n is at most $k_1 h$, for some universal computable constant k_1 . Also, there is some universal computable constant k_2 which is an upper bound for the ℓ_2 norm of each row of the matrix A . So, by (2) of Theorem 8.1, each co-ordinate of a fundamental boundary-regulated surface has modulus at most $n^{3/2}k_2^{n-1}$. Furthermore, each elementary disc intersects at most k_3 2-handles, for some universal computable constant k_3 . So, the weight of the surface is at most $k_3 n^{3/2}k_2^{n-1}$. Since $n^{3/2} < 2^n$ for all $n \geq 2$, we get an upper bound of c^h , for some universal computable constant c . \square

The above result is sufficient for most surfaces in our hierarchy. However, the first surface is different, because its homology class is fixed. It is the Thurston norm of this homology class that is being certified. In order to deal with this, we use some work of Tollefson and Wang [27]. They used triangulations rather than handle structures, and so we will do the same here.

Central to the approach of Tollefson and Wang is the notion of a lw-taut surface. They say that a normal surface S properly embedded in a compact triangulated 3-manifold M is *lw-taut* if

- (i) it has smallest Thurston complexity in its class in $H_2(M, \partial M)$;
- (ii) no collection of components of S is homologically trivial; and
- (iii) it has smallest possible weight among normal surfaces satisfying (i) and (ii) in the same homology class of $H_2(M, \partial M)$.

Note that there is an important distinction in the way that Tollefson and Wang use the word ‘taut’ compared with its use by Scharlemann [24] and Gabai [3]. Tollefson and Wang require a ‘taut’ surface to have minimal Thurston complexity in its class in $H_2(M, \partial M)$, whereas Scharlemann and Gabai require Thurston complexity only to be minimised in the class in $H_2(M, N(\partial S))$. Fortunately when ∂M is a collection of tori, these concepts are closely related, using Lemma 2.1.

Our result bounds the weight of the first surface in the hierarchy, as follows.

Theorem 8.3. *Let M be a compact orientable irreducible 3-manifold with boundary a (possibly empty) collection of tori. Let \mathcal{T} be a triangulation of M . Let ϕ be a simplicial 1-cocycle. Then there is a compact oriented lw-taut normal surface S such that $[S, \partial S]$ is Poincaré dual to $[\phi]$, satisfying $w(S) \leq k^t \|\phi\|_1$, where k is a universal computable constant and t is the number of tetrahedra of \mathcal{T} . Here, $\|\phi\|_1$ is the ℓ_1 norm of ϕ .*

Tollefson and Wang consider the non-negative real solutions to the normal surface matching equations and compatibility constraints. These form a union of convex polytopes. The boundary of any such polytope has a natural decomposition into faces. The polytope itself is also called a face. A solution lying in such a face is said to be *carried* by the face, and when the solution corresponds to a normal surface S , we also say that S is *carried* by the face. There is a unique face of minimal dimension that carries a given solution.

The following result of Tollefson and Wang is a central part of their work (see Theorem 3.3 in [27]).

Theorem 8.4. *Let M be a compact orientable triangulated 3-manifold. Let S be an oriented, lw-taut normal surface, and let C_S be the minimal face carrying S . Then every surface carried by C_S is lw-taut. Moreover, every surface carried by C_S inherits a well-defined orientation, and for any two surfaces G and H carried by C_S , the homology class in $H_2(M, \partial M)$ of the normal sum $G + H$ is equal to the sum of the homology classes of G and H .*

Proof of Theorem 8.3. Let S be an oriented, lw-taut surface such that $[S, \partial S]$ is Poincaré dual to $[\phi]$, and let C_S be the minimal face carrying S . This face is the convex hull of the rays through finitely many vertex surfaces S_1, \dots, S_m . By Theorem 8.1 (1), each co-ordinate of each S_i has modulus at most $(c_1)^t$, for some universal computable constant c_1 . Now, any element of C_S can be written as a linear combination $\sum_i \lambda_i S_i$, where each $\lambda_i \geq 0$. We may in fact ensure that there is no linear dependence between the homology classes of the S_i for which λ_i is non-zero, as follows. Say that $\sum \mu_i [S_i, \partial S_i] = 0$, where not all the μ_i are zero. Then we add a small multiple of $\sum \mu_i S_i$ to $\sum_i \lambda_i S_i$. By Theorem 8.4, the effect on the homology class is to add a multiple of $\sum \mu_i [S_i, \partial S_i]$, which is zero. By choosing this multiple of $\sum \mu_i S_i$ correctly, we may arrange that in $\sum_i (\lambda_i + \mu_i) S_i$, all the coefficients $\lambda_i + \mu_i$ are non-negative and at least one more coefficient is zero than is the case for $\sum \lambda_i S_i$. In this way, we may decrease the number of non-zero coefficients. When this is minimal, the S_i with non-zero coefficients are homologically linearly independent. When we apply this procedure to S , we get $S = \sum_i \lambda_i S_i$ with exactly r non-zero terms, say, where $r \leq \dim(H_2(M, \partial M; \mathbb{Q})) \leq 2t$. By relabelling, we may ensure that $\lambda_1, \dots, \lambda_r$ are non-zero.

Pick a maximal tree in the 1-skeleton of \mathcal{T} . Then each edge not in the tree determines a loop in M that starts at some basepoint, travels in the tree to the start of the edge, runs along the edge, and then back through the tree to the basepoint. These loops form a generating set for $\pi_1(M)$ and hence a generating set for $H_1(M)$. So, some subset ℓ_1, \dots, ℓ_d forms a basis for $H_1(M; \mathbb{Q})$. Therefore, an element of $H_2(M, \partial M)$ is determined by its signed intersection numbers with these loops ℓ_1, \dots, ℓ_d . In the subspace of $H_2(M, \partial M)$ spanned by the surfaces in S_1, \dots, S_r , we may determine the class of the surface $\sum_i \lambda_i S_i$ by examining its signed intersection number with just r loops in ℓ_1, \dots, ℓ_d . By re-ordering, we may assume that these loops are ℓ_1, \dots, ℓ_r .

Form an $r \times r$ matrix P , with (i, j) entry equal to the signed intersection number between ℓ_i and S_j . Similarly form an $e \times r$ matrix Q , with (k, j) entry equal to the number of intersection points between an edge e_k and S_j , where e is the number of edges of the triangulation. Hence, for a surface $\sum_{i=1}^r \lambda_i S_i$ carried by C_S , the unsigned intersection numbers with the edges are the entries of $Q\lambda$, where $\lambda = (\lambda_1, \dots, \lambda_r)^T$. Similarly, the signed intersection numbers with the loops ℓ_1, \dots, ℓ_r are $P\lambda$, since by Tollefson and Wang's Theorem 8.4, the homology class of $\sum \lambda_i S_i$ is $\sum \lambda_i [S_i, \partial S_i]$. Now, the signed intersection number of S with the loops ℓ_i is the evaluation $\phi(\ell_i)$, where ϕ is the given simplicial 1-cocycle. Let μ be the vector with entries $\phi(\ell_i)$. Hence, the unsigned intersection numbers of S with the edges of \mathcal{T} are given by $QP^{-1}\mu$.

Each co-ordinate of the normal surface vector (S_j) is, by Theorem 8.1(1), at most $(c_1)^t$. So, each entry of P has modulus at most $20t(c_1)^t$, since each elementary normal disc intersects the 1-skeleton in at most 4 points and there are at most $5t$ types of elementary disc in S_i . Similarly, each entry of Q is at most $20t(c_1)^t$.

Therefore, the ℓ^2 norm of each row of P is at most $r^{1/2}20t(c_1)^t$. Hence, each entry of P^{-1} has modulus at most $(r^{1/2}20t(c_1)^t)^r$. Therefore, each entry of $QP^{-1}\mu$ has modulus at most $20t(c_1)^tr^2(r^{1/2}20t(c_1)^t)^r\|\phi\|_1$. This is at most $k^{t^2}\|\phi\|_1$, for some universal computable constant k . \square

9. DETERMINING THE COMPONENTS OF THE PARALLELITY BUNDLE

We saw in Section 7 that when (M, γ) is a taut sutured manifold with a positive handle structure \mathcal{H} satisfying some natural conditions, then a taut sutured manifold decomposition may be performed along a regulated surface S that is fundamental as a boundary-regulated surface. Hence, by Theorem 8.2, there is an upper bound to the number of elementary discs of S , that is an exponential function of the number of 0-handles of \mathcal{H} . Examples due to Hass, Thurston and Snoeyink [11] demonstrate that one cannot, in general, find a better upper bound than this. This is potentially problematic, because the sutured manifold (M', γ') obtained by decomposing along S inherits a handle structure \mathcal{H}' that may have many more handles than \mathcal{H} . Then, continuing this down the hierarchy, the resulting handle structures may become so complex that they cannot be analysed efficiently and so are not suitable for an NP algorithm.

Fortunately, there is a method to get around this problem. Although S may have a large number of elementary discs, there is a linear upper bound (as a function of the number of handles of \mathcal{H}) on the number of elementary disc types of S . Between two discs of the same type, there is a subset of M' homeomorphic to $D^2 \times I$. These pieces patch together to form an I -bundle embedded in M' known as the *parallelity bundle*. There is a linear upper bound on the number of handles of \mathcal{H}' that are not part of this bundle. Thus, when S has very large weight, almost all the handles of M' lie within this bundle. The solution, therefore, is to remove or modify this bundle. The parts that are I -bundles over discs are replaced by 2-handles. The parts that are I -bundles over more complicated surfaces are typically removed by decomposing along the annuli that separate them from the remainder of M' . In order to achieve this algorithmically, it is important to be able determine the components of the parallelity bundle and to decide efficiently which are I -bundles over discs.

In this section, we explain how to do this. In Section 9.1, we give the precise definition of the parallelity bundle. In Section 9.2, we describe an algorithm, due to Agol, Hass and Thurston [2], that was originally designed to determine the Euler characteristic of the components of a normal surface. In Section 9.3, we give our main application, which is a method for determining the components of the parallelity bundle of the exterior of a normal or regulated surface, and for deciding which of these components are I -bundles over discs. In Sections 9.4 and 9.5, we give further applications of the Agol-Hass-Thurston algorithm. The first of these is concerned with ‘normal’ 1-manifolds in a surface, and it gives an efficient method for determining the components of the complement of such a 1-manifold. In Section 9.5, we give a method for determining whether a boundary-regulated surface can be transversely oriented.

9.1. THE PARALLELITY BUNDLE FOR SUTURED MANIFOLDS

Let \mathcal{H} be a handle structure for the sutured manifold (M, γ) . A handle H of \mathcal{H} is a *parallelity handle* if it satisfies one of the following:

- (i) it is a 2-handle that is disjoint from 3-handles;
- (ii) it is a 1-handle that intersects $\mathcal{H}^2 \cup \gamma$ in precisely two components and that is disjoint from the 3-handles; or
- (iii) it is a 0-handle H such that $H \cap (\mathcal{F} \cup \gamma)$ is connected, every 0-handle of $H \cap \mathcal{F}^0$ has index 0 and H is disjoint from the 3-handles.

A product structure $D^2 \times I$ may be imposed on any parallelity handle H so that

- (i) the intersection between H and any other handle is of the form $\beta \times I$, for a subset β of ∂D^2 , and
- (ii) each component of $H \cap \gamma$ is of the form $\beta \times \{1/2\}$ for some subset β of ∂D^2 .

These product structures may all be chosen so that the I -bundles agree on the intersection between any collection of parallelity handles. Hence, they form an I -bundle \mathcal{B} over a surface F , known as the *base surface*. (See the proof of Lemma 5.3 in [21] for example.) This I -bundle is called the *parallelity bundle*. The *horizontal boundary* $\partial_h \mathcal{B}$ is the ∂I -bundle over F . The *vertical boundary* $\partial_v \mathcal{B}$ is the I -bundle over ∂F . The intersection between \mathcal{B} and the remainder of M is a subset of $\partial_v \mathcal{B}$ consisting of annuli disjoint from γ and product discs.

Note that this notion of the parallelity bundle is related to, but not the same as, the notion of an ‘amalgam’ presented in [20]. An amalgam was also an I -bundle over a surface. However, away from its vertical boundary, it was allowed to contain handles that were not parallelity handles. Another key difference was that its vertical boundary was required to have interior disjoint from ∂M , and hence be annuli properly embedded in M .

9.2. THE PARALLELITY BUNDLE FOR PAIRS

There is also a version of the parallelity bundle in the case where the 3-manifold has a specified subsurface S in its boundary, which is as follows.

Let \mathcal{H} be a handle structure for the compact orientable 3-manifold M . Let S be a compact subsurface of ∂M , such that ∂S is disjoint from the 2-handles and respects the product structure on the 1-handles of \mathcal{H} . In this case we say that \mathcal{H} is a *handle structure* for the pair (M, S) .

A handle H of \mathcal{H} is a parallelity handle if it admits a product structure $D^2 \times I$ such that

- (i) $D^2 \times \partial I = H \cap S$;
- (ii) each component of $\mathcal{F}^0 \cap H$ and $\mathcal{F}^1 \cap H$ is $\beta \times I$, for a subset β of ∂D^2 .

Again, we will view this as an I -bundle over D^2 , and again, these bundle structures may be chosen so that they agree on the intersection between handles. So, their union is an I -bundle over a surface. This is the *parallelity bundle* for the pair (M, S) .

The main case when we consider pairs in this way is the following situation. Suppose that M is a compact orientable 3-manifold with a triangulation \mathcal{T} . Suppose that S is a properly embedded orientable normal surface in M . Then cutting M along S gives a 3-manifold M' . It has a handle structure \mathcal{H}' , which arises by first dualising \mathcal{T} to form a handle structure for M , and then decomposing this along S . Inside $\partial M'$, the two copies of S form a subsurface S' . Then we will frequently consider the parallelity bundle for the pair (M', S') . In this case, parallelity handles arise precisely in the space between two elementary normal discs of S .

9.3. THE EXTENDED ORBIT COUNTING ALGORITHM OF AGOL, HASS AND THURSTON

In this subsection, we recall the algorithm of Agol, Hass and Thurston [2]. The setting for the algorithm is the natural numbers $\{1, 2, \dots, N\}$, which we denote by $[1, N]$. For integers $a \leq b \in [1, N]$, the set of integers lying between a and b , including a and b , is denoted $[a, b]$.

The input to the algorithm is the following data:

- (i) positive integers N, k, d and m ;
- (ii) a collection of linear bijections $g_i: [a_i, b_i] \rightarrow [c_i, d_i]$, where $a_i, b_i, c_i, d_i \in [1, N]$, and where $1 \leq i \leq k$;

(iii) a function $z: [1, N] \rightarrow \mathbb{Z}^d$, with the property that fewer than m integers j satisfy $z(j) \neq z(j+1)$.

The linear bijections are called *pairings*. Note that because each g_i is a linear bijection, it sends the endpoints of the interval $[a_i, b_i]$ to the endpoints of $[c_i, d_i]$. It is *orientation preserving* if $g_i(a_i) = c_i$ and $g_i(b_i) = d_i$. Otherwise it is *orientation reversing*. Obviously, an orientation preserving pairing is just a translation.

Two integers x and y in $[1, N]$ are said to be in the same *orbit* if there is a sequence of these bijections g_{i_1}, \dots, g_{i_s} such that $g_{i_s}^{\pm 1} \dots g_{i_1}^{\pm 1}(x) = y$. Clearly, this forms an equivalence relation and the equivalence classes are the orbits.

The *weight* of an orbit is the sum, over all elements x in the orbit, of $z(x)$.

Let D be the maximal value, over all integers $x \in [1, N]$ of $\|z(x)\|_1$, where the ℓ_1 norm on \mathbb{Z}^d is used.

The following theorem is due to Agol, Hass and Thurston [2].

Theorem 9.1. *There is an algorithm, with running time that is bounded above by a polynomial in $kmd(\log D)(\log N)$, that produces a list of all orbits with their weights.*

It is the $\log N$ that is critical here.

It is perhaps helpful here to consider a simplified situation, where $d = 1$ and the weight function z takes the constant value 1. Then the weight of an orbit is just the number of elements in it. Thus, the algorithm in this situation produces a list of orbits, together with the number of elements in each orbit.

However, it is useful to allow more complicated weight functions. For example, suppose that we wanted to determine the number of elements in the orbit of x , for some given $x \in [1, N]$. Then we could set $d = 2$, and assign every element to have weight $(1, 0)$, except for x which has weight $(1, 1)$. Hence, the weight of an orbit is either $(r, 0)$ or $(r, 1)$, where r is the number of elements in the orbit, and where the second co-ordinate is 1 if and only if the orbit contains x . Hence, by examining the unique orbit of weight $(r, 1)$, we can determine the number elements in the orbit of x .

Agol, Hass and Thurston used their algorithm to determine the Euler characteristic of the components of a normal surface S in a compact triangulated 3-manifold M , as follows. The input to the algorithm is the triangulation \mathcal{T} of M and the normal surface vector (S) . The integer N is $w(S)$, the weight of S . We think of $[1, N]$ as the points of intersection between S and the 1-skeleton of \mathcal{T} . More precisely, each edge e of \mathcal{T} is oriented in some way, and the points $S \cap e$ are identified with some interval $[a, b]$ in $[1, N]$, so that the linear ordering of the points along e agrees with the linear order on $[a, b]$. The bijections g_i arise from the faces of the triangulation. For each face f , the arcs $S \cap f$ come in at most three types. The arcs of each type run from the 1-skeleton to the 1-skeleton, and so specify a linear bijection from a sub-interval of $[1, N]$ to another such sub-interval.

It is fairly clear that there is a one-one correspondence between the components of S and the orbits, for the following reason. Each component of S contains a point of intersection p with the 1-skeleton, and this corresponds to an integer x in $[1, N]$. For any other point of intersection with the 1-skeleton, corresponding to an integer y , this lies in the same component of S as p if and only if there is a path joining it to p that lies in the intersection between S and the 2-skeleton of \mathcal{T} . This is equivalent to the existence of an equation $g_{i_s}^{\pm 1} \dots g_{i_1}^{\pm 1}(x) = y$, for some bijections g_{i_1}, \dots, g_{i_s} .

By choosing the weight function z appropriately, we can count the number of times components of S intersect each edge e of the triangulation. The Euler characteristic of a component S' is a linear function of the integers $|S' \cap e|$, as e runs over all edges of the triangulation. Hence, this Euler characteristic can be calculated in time that is a polynomial function of $t \log w(S)$, where t is the number of tetrahedra of \mathcal{T} .

9.4. DETERMINING THE COMPONENT TYPES OF THE PARALLELITY BUNDLE

We now come to our main applications of Theorem 9.1.

Theorem 9.2. *There is an algorithm that takes, as its input,*

- (i) *a handle structure \mathcal{H} , of uniform type, with h handles, for a sutured manifold (M, γ) ;*
- (ii) *a boundary-regulated vector $(S)_{\partial r}$ for a regulated surface S ;*

and provides as its output, the following data. If (M', γ') is the sutured manifold that results from decomposing along S , and \mathcal{B} is the parallelity bundle for the handle structure that it inherits, then the algorithm produces a handle structure for $\text{cl}(M' - \mathcal{B})$ and, for each component B of \mathcal{B} , it determines:

- (i) *the genus and number of boundary components of its base surface;*
- (ii) *whether B is a product or twisted I -bundle; and*
- (iii) *the location of $\partial_v B$ in $\text{cl}(M' - \mathcal{B})$.*

This algorithm runs in time that is bounded by a polynomial in $h \log(w(S))$.

We will also obtain the following variant, which deals with pairs.

Theorem 9.3. *There is an algorithm that takes, as its input,*

- (i) *a triangulation \mathcal{T} , with t tetrahedra, for a compact orientable manifold M ;*
- (ii) *a vector (S) for an orientable normal surface S ;*

and provides as its output, the following data. If M' is the manifold that results from decomposing along S , and S' is the two copies of S in $\partial M'$, and \mathcal{B} is the parallelity bundle for the pair (M', S') with its induced handle structure, then the algorithm determines the same information as in Theorem 9.2. It runs in time that is bounded by a polynomial in $t \log(w(S))$.

Proof. We will focus on the proof of Theorem 9.2. The proof of Theorem 9.3 is very similar, and requires only very minor modifications.

The first stage in the procedure is to construct the handle structure of $\text{cl}(M' - \mathcal{B})$. Each handle H of \mathcal{H} is divided up into handles of \mathcal{H}' . It is a straightforward matter to determine those handles of $H \cap \mathcal{H}'$ that are not parallelity handles. This is done by examining the entries of $(S)_{\partial r}$ corresponding to the elementary disc types in H and noting which of these are zero and which are non-zero. The number of handles of $H \cap \mathcal{H}'$ that are not parallelity handles is bounded above by a universal computable constant, for the following reason. Since \mathcal{H} is of uniform type and S is boundary-regulated, there is a universal computable upper bound on the number of elementary disc types of S in H . If two adjacent elementary discs of the same type are disjoint from ∂M , then the region between them becomes a parallelity handle of $H \cap \mathcal{H}'$. If two adjacent elementary discs of the same type intersect ∂M and are coherently oriented, then also the region between them becomes a parallelity handle of $H \cap \mathcal{H}'$. Because S is boundary-regulated, at most one pair of adjacent discs of the same type intersecting ∂M can be incoherently oriented.

Consider a handle of H' of \mathcal{H}' that is not a parallelity handle. It is adjacent to a bounded number of other handles of \mathcal{H}' . Some of these are parallelity handles, and so will form pieces of $\partial_v \mathcal{B} \cap H'$. The remainder form handles of $\text{cl}(M' - \mathcal{B})$.

Thus, it is clear that a handle structure for $\text{cl}(M' - \mathcal{B})$ can be constructed in time that is bounded above by a polynomial function of h . Moreover, the components of $\partial_v \mathcal{B}$ can be located.

Since \mathcal{H} is of uniform type, there is a universal computable upper bound c to the number of types of elementary discs of S in each handle of \mathcal{H} . Say that there are $n \leq ch$ elementary disc types of S . For each integer i between 1 and n , let x_i be the number of elementary discs of that type. When a disc

type intersects ∂M , we distinguish between the two possible transverse orientations on it. So, when the elementary disc type lies in a 0-handle of \mathcal{H} , then x_i is a co-ordinate of $(S)_{\partial r}$.

For each i between 1 and n , set N_i to be twice the number of parallelity handles that are incident to the x_i elementary discs. So, N_i is the number of components of intersection between $\partial_h \mathcal{B}$ and these parallelity handles. Thus, N_i equals $x_i - 1$, x_i or $x_i + 1$, depending on whether the handles that intersect only one elementary disc of that type are parallelity handles. Let N be the number of elementary discs that make up $\partial_h \mathcal{B}$. So, N is the sum of $\sum_{i=1}^r 2N_i$ and twice the number of parallelity handles of \mathcal{H} that are disjoint from S .

We now define the bijections g_i . Consider any q -handle H_q of \mathcal{H} that is attached to some p -handle H_p where $p < q$. Consider any elementary disc types in H_p and H_q that intersect. These correspond to intervals $[a_j, b_j]$ and $[c_j, d_j]$ in $[1, N]$. Thus, we get a bijection $[a_j, b_j] \rightarrow [c_j, d_j]$. We consider all such bijections, as we run over all pairs of incident handles H_p and H_q and all elementary disc types in H_p and H_q . We also consider the parallelity handles of \mathcal{H} disjoint from S , which also give rise to similar bijections. The number of these bijections is at most $(ch)^2$.

We now define the weight function $z: [1, N] \rightarrow \mathbb{Z}^d$. The integer d is equal to $2|\partial_v \mathcal{B}| + 3$. The first three co-ordinates are integers used for counting. The remaining co-ordinates correspond to the boundary components of $\partial_v \mathcal{B} - \gamma$. Each integer j in $[1, N]$ corresponds to an elementary disc E of $\partial_h \mathcal{B}$. Then $z(j)$ is a d -tuple. The first co-ordinate is set to be 1 if E lies in a 0-handle of \mathcal{H} . The second is set to be 1 if E is in a 1-handle. The third co-ordinate is 1 if E is in a 2-handle. The remaining co-ordinates are set to be zero, unless the handle is incident to $\partial_v \mathcal{B}$. In this case, the relevant co-ordinate of $z(j)$ is set to 1 if E runs over the relevant boundary component of $\partial_v \mathcal{B} - \gamma$; otherwise, it is zero.

It is clear that the orbits correspond to the components of $\partial_h \mathcal{B}$. Moreover, the first three components of the weight function count the number of handles of each index of $\partial_h \mathcal{B}$. Therefore, they can be used to compute the Euler characteristic of this component. The remaining co-ordinates can be used to determine whether a component of $\partial_h \mathcal{B}$ contains one of the boundary components of $\partial_v \mathcal{B} - \gamma$. Theorem 9.1 therefore produces a list of components of $\partial_h \mathcal{B}$, and for each such component F , the location of ∂F in $\partial_v \mathcal{B}$ and the Euler characteristic $\chi(F)$. From this, one readily determines the data required by the theorem. The time taken by the algorithm of Theorem 9.1 is at most a polynomial function of $kmd(\log D)(\log N)$. Here, k is the number of pairings, which is at most $(ch)^2$. The integer D is at most $2|\partial_v \mathcal{B}| + 1$, which is bounded above by a linear function of h . The integer m is the number of values of j such that $z(j) \neq z(j+1)$. These values of j can occur at the outermost elementary discs of the same type, and when an elementary normal disc or one of its neighbours intersects $\partial_v \mathcal{B}$. So, again, there is an upper bound on the number of such values of j , which is a polynomial function of h . Finally, N is at most a linear function of $w(S)$ and h . So, the running time is at most a polynomial function of $h \log(w(S))$, as required. \square

In the remainder of this section, we give two further applications of the Agol-Hass-Thurston algorithm.

9.4. CUTTING A SURFACE ALONG A NORMAL 1-MANIFOLD

In this section, we consider a handle structure \mathcal{H} for a compact surface F . As usual, we let \mathcal{H}^i denote the union of the i -handles, for $i = 0, 1$ and 2 .

We say that a 1-manifold C properly embedded in F is *normal* if

- (i) it misses the 2-handles of \mathcal{H} ;
- (ii) it respects the product structure on each 1-handle;
- (iii) its intersection with each 0-handle is a collection of properly embedded arcs;
- (iv) for each 0-handle H_0 , no arc of $C \cap H_0$ has both endpoints on the same component of $H_0 \cap \mathcal{H}^1$ or on the same component of $\partial H_0 - \mathcal{H}^1$.

It is clear that normal 1-manifolds may be encoded algebraically, as follows. The intersection between a normal 1-manifold C and a 0-handle H_0 is a collection of arcs, which come in finitely many types. One can form a vector (C) which counts the number of arcs of each type. These satisfy a collection of *matching equations*, one for each 1-handle. There are also *compatibility constraints* which prevent distinct arc types within a 0-handle from coexisting if they inevitably intersect each other.

The *weight* of a normal curve C is $|C \cap \mathcal{H}^1|$.

We say that the handle structure \mathcal{H} has *uniform type* if, for each 0-handle H_0 , the number of components of $H_0 \cap \mathcal{H}^1$ is bounded above by some universal computable constant.

Theorem 9.4. *There is an algorithm that takes, as its input,*

- (i) *a handle structure \mathcal{H} , of uniform type, with h handles, for a compact surface F ;*
- (ii) *a vector (C) for a normal 1-manifold C ,*

and provides as its output, the following data for each component F' of $F - \text{int}(N(C))$:

- (1) *the genus of F' ;*
- (2) *the number of boundary components of F' ;*
- (3) *the location of the components of $F' \cap \partial F$;*
- (4) *the vector (C') for each component C' of $\text{cl}(\partial F' - \partial F)$.*

The algorithm runs in time that is bounded by a polynomial in $h \log(w(C))$. This polynomial depends on the implied constant in the definition of uniform.

The proof is an easier version of Theorem 9.2, and so we omit it.

Corollary 9.5. *There is an algorithm that takes, as its input,*

- (i) *a handle structure \mathcal{H} , of uniform type, with h handles, for a decorated sutured manifold (M, γ) , such that for each 0-handle H_0 of \mathcal{H} , $H_0 \cap (\mathcal{F} \cup \gamma)$ is connected;*
- (ii) *a vector $(S)_{\partial r}$ for a boundary-regulated surface S ,*

and determines whether any component of ∂S bounds a disc disjoint from γ . It also determines those components of ∂S disjoint from γ that are trivial, and for each such component, the algorithm determines the side of ∂S on which the trivialising planar surface lies. The algorithm runs in time that is bounded by a polynomial in $h \log(w(S))$.

Proof. The handle structure \mathcal{H} induces a handle structure on $R_{\pm}(M)$. Since \mathcal{H} is of uniform type, there is a universal computable upper bound to the number of times that a 0-handle of $R_{\pm}(M)$ can intersect the 1-handles. Therefore, we may declare that the handle structure on $R_{\pm}(M)$ is of uniform type.

The intersection between ∂S and $R_{\pm}(M)$ is a normal 1-manifold C in this handle structure, possibly after removing components of $\partial S \cap R_{\pm}(M)$ that lie within a single 0-handle H_0 of \mathcal{H} and have endpoints on the same component of $H_0 \cap \gamma$. The vector (C) is a simple linear function of $(S)_{\partial r}$. Using Theorem 9.4, we can determine whether some component F of $R_{\pm}(M) - \text{int}(N(C))$ is a disc with boundary disjoint from $N(\gamma)$. Hence, we can decide whether any component of ∂S bounds a disc disjoint from γ .

We can determine the trivial curves of C as follows. Applying Theorem 9.4, we build a graph. It has a vertex for each component of $R_{\pm}(M) - \text{int}(N(C))$ that does not lie between two normally parallel curves of C . For a vertex corresponding to a component F' of $R_{\pm}(M) - \text{int}(N(C))$, the edges emanating from this vertex correspond to the components of $\text{cl}(\partial F' - \partial F)$. When an edge starts at one vertex and ends at another, the corresponding normal arcs of $\partial S \cap R_{\pm}(M)$ are normally parallel. Each edge of the graph therefore corresponds to a maximal normally parallel collection of arcs or simple closed curves. We call

these *arc-type* and *curve-type* edges respectively.

We now can characterise the trivial curves of C . Each corresponds to a curve-type edge of this graph that separates off a subgraph that is a tree and that corresponds to a subsurface F' in which all the vertices have genus zero, all the edges in this subgraph are of curve-type and all the components of $F' \cap \partial F$ are u -sutures. This can readily be determined using Theorem 9.4. \square

9.5. DETERMINING WHETHER A SURFACE IS TRANSVERSELY ORIENTED

Although boundary-regulated surfaces are convenient to describe by means of their vectors, it is regulated surfaces that we actually need to decompose along. The following result can be used to determine whether a boundary-regulated surface is regulated.

Theorem 9.6. *There is an algorithm that takes, as its input*

- (i) a handle structure \mathcal{H} , of uniform type, with h handles, for a sutured manifold (M, γ) ;
- (ii) a vector $(S)_{\partial r}$ for a boundary-regulated surface S ;
- (iii) a transverse orientation on a collection of at most r elementary discs of S ;
- (iv) a specified 2-handle of \mathcal{H} with a transverse orientation on its cocore e ;

and determines whether there is a transverse orientation on S that is compatible with the transverse orientation on the discs in (iii) and the given transverse orientation on the discs that intersect ∂M . It also computes the signed intersection number with between e and S with this transverse orientation. The algorithm runs in time bounded above by a polynomial in $hr \log(w(S))$.

Using essentially the same proof, we can also determine whether a normal surface is orientable.

Theorem 9.7. *There is an algorithm that takes, as its input*

- (i) a triangulation \mathcal{T} with t tetrahedra for a compact orientable 3-manifold M ;
- (ii) a vector (S) for a normal surface S ;
- (iii) a transverse orientation on a collection of at most r elementary discs of S ;
- (iv) a specified edge e of \mathcal{T} ;

and determines whether there is a transverse orientation on S that is compatible with the transverse orientation on the discs in (iii). It also computes the signed intersection number with between e and S with this transverse orientation. The algorithm runs in time bounded above by a polynomial in $tr \log(w(S))$.

We will focus on the proof of Theorem 9.6, as the proof of Theorem 9.7 is essentially the same.

Proof. From the vector $(S)_{\partial r}$, one may produce a list of components of S . Indeed, this was one of the original applications of the Agol-Hass-Thurston algorithm. These are produced as a collection of vectors. It suffices to check, for each such component, that it admits a transverse orientation that is compatible with the given orientations and to compute its signed intersection number with e . Thus, we may assume that S is connected.

We first determine whether the connected surface S is transversely orientable, but with no regard as to whether this is compatible with the pre-assigned transverse orientations on various elementary discs. To do this, one forms the orientable double cover, and one checks whether it has two or one component, as follows.

From the vector $(S)_{\partial r}$, one may readily compute the number of elementary discs of S of each type. For elementary discs in a 0-handle of \mathcal{H} , this is just a co-ordinate of $(S)_{\partial r}$, or a sum of two co-ordinates. For elementary discs in 1-handles and 2-handles, this is a simple linear function of the vector $(S)_{\partial r}$. Let N

be twice the number of elementary discs of \mathcal{H} . We view $[1, N]$ as divided into two halves. Integers in the first half $[1, N/2]$ correspond to the elementary discs of S with some transverse orientation, and integers in the second half correspond to the same discs but with the opposite orientation. Elementary discs of the same type correspond to sub-intervals of $[1, N/2]$ or $[N/2 + 1, N]$. The number of such subintervals is at most $k_1 h$, for some universal computable k_1 . Whenever a p -handle and q -handle of \mathcal{H} are incident, the elementary normal discs in these handles patch together. Hence, we obtain bijections between sub-intervals of $[1, N]$. These are chosen so that they respect the transverse orientations on the discs. The number k of such pairings is at most $(k_2 h)^2$, for some universal computable k_2 .

Clearly, S induces either one or two orbits, depending on whether S is unorientable or orientable. However, we also need to check whether it has a transverse orientation that is compatible with the given ones on various elementary discs. To do this, we introduce a weight function $z: [1, N] \rightarrow \mathbb{Z}^d$, where d is r plus twice the number of elementary disc types that intersect ∂M . For an integer $j \in [1, N]$ corresponding to an elementary disc, and integer i satisfying $1 < i \leq d$, the i th co-ordinate of $z(j)$ is an ‘error’ co-ordinate. It is set to 1 if the orientation on the disc corresponding to j is contrary to the orientation specified by the fact that it is boundary-regulated, or by the transverse orientation given in (iii). Otherwise, this co-ordinate is set to 0. The first co-ordinate counts the signed intersection number between the elementary disc and the arc e . Thus, it is zero unless the disc lies in the same 2-handle as e , in which case it is ± 1 . So, the number of values of j such that $z(j) \neq z(j+1)$ is at most $2r$ plus twice the number of disc types.

Theorem 9.1 produces a list of orbits and their weights. There is a transverse orientation on the connected surface S that is compatible with the given ones if and only if there are two orbits and at least one of these orbits has weight $(q, 0, 0, \dots, 0)$. In this case, q is the signed intersection number between S and e . \square

10. AN ALGORITHM TO SIMPLIFY HANDLE STRUCTURES

In Theorem 9.2, we showed that, when a taut sutured manifold decomposition is performed along a regulated surface, the components and structure of the parallelity bundle of the resulting sutured manifold (M', γ') may be efficiently determined. In this section, we utilise this information to produce a simplified handle structure for a sutured manifold obtained from (M', γ') .

Theorem 10.1. *Let \mathcal{H} be a positive handle structure of a connected decorated sutured manifold (M, γ) , of uniform type, with h handles. Let S be a regulated surface properly embedded in M . Let (M', γ') be obtained by performing an allowable decomposition along S . Suppose that no component of M' has boundary a single torus with no sutures. Then, there is an algorithm that takes as its input \mathcal{H} and $(S)_{\partial r}$ and either correctly asserts that $E(M', \gamma')$ is not taut, or supplies a handle structure \mathcal{H}'' for a decorated sutured manifold (M'', γ'') with the following properties:*

(i) *there is a decomposition*

$$E(M', \gamma') \xrightarrow{A} E(M'', \gamma'') \# Y$$

for some closed 3-manifold Y , and where A is a collection of oriented non-trivial annuli in M' disjoint from γ' ; this decomposition is taut if and only if $E(M', \gamma')$ is taut;

(ii) *if $E(M', \gamma')$ is taut, then it can be certified in polynomial time as a function of h that Y is homeomorphic to the 3-sphere;*

(iii) *if $E(M', \gamma')$ is atoroidal and no component of $E(M', \gamma')$ is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$, and (M', γ') admits an allowable hierarchy, then so does (M'', γ'') ;*

(iv) *\mathcal{H}'' is of uniform type;*

(v) *each 0-handle of \mathcal{H}'' lies inside a 0-handle of \mathcal{H} ;*

(vi) *for each 0-handle H of \mathcal{H} , the complexity of $H \cap \mathcal{H}''$ is at most that of H , and in the case of equality,*

$H \cap \mathcal{H}''$ contains a single 0-handle isotopic to H ;

(vii) \mathcal{H}'' has no parallelity handles other than 2-handles;

(viii) \mathcal{H}'' is positive.

The algorithm runs in time that is bounded by a polynomial in $h \log(w(S))$.

The idea behind this theorem is simply that if one were to decompose \mathcal{H} by cutting along S , then the resulting handle structure \mathcal{H}' might have many parallelity handles. Our goal is to remove these parallelity handles, primarily by decomposing along the vertical boundary components of the parallelity bundle \mathcal{B} for \mathcal{H}' , or by replacing them by 2-handles. The exact procedure is slightly delicate, and requires modifications to be made in the right order. We follow the method given in Section 8 of [20].

10.1. DECOMPOSITION ALONG PRODUCT DISCS

Some components B of the parallelity bundle \mathcal{B} may have vertical boundary that intersects ∂M . In that case, $\text{cl}(\partial_v B - \partial M)$ consists of at least one product disc, plus possibly some product annuli.

If any of these product discs intersects a u-suture, we attach a 2-handle along the u-suture, and enlarge \mathcal{B} by including this 2-handle in the parallelity bundle. This may increase the complexity of the handle structure, but this is only a temporary problem, because all components of \mathcal{B} will be later removed from the manifold.

So suppose that a component B of \mathcal{B} has vertical boundary that intersects ∂M but that is disjoint from all u-sutures. In this case, we may remove B . The product discs and annuli in $\text{cl}(\partial_v B - \partial M)$ become discs and annuli in the boundary of the new manifold, each of which contains a new arc of a suture. The reason that we may remove B in this way is as follows.

We can decompose along any product discs in $\text{cl}(\partial_v B - \partial M)$, since these are allowable. If B becomes detached from the remainder of M , we stop. If it is still attached, then we may also decompose along vertical product discs in B , until B becomes a collection of collar neighbourhoods of the annuli in $\text{cl}(\partial_v B - \partial M)$. We can then remove these collar neighbourhoods, without changing the homeomorphism type of the manifold.

Thus, in this way, we may assume that every component of $\partial_v \mathcal{B}$ is a properly embedded annulus in M . All future modifications that we make will preserve this property.

10.2. THE BOUNDARY GRAPH

In order to be able to decide which vertical boundary components of the parallelity bundle to cut along, we introduce two graphs. The first of these, the boundary graph, encodes the way that $\partial M'$ is cut into pieces by the parallelity bundle. The second graph, the connectivity graph, stores information about the way that the pieces of $M' - \mathcal{B}$ and \mathcal{B} are connected together.

Let \mathcal{H}' be the handle structure obtained by decomposing \mathcal{H} along S . Let \mathcal{B}' be the union of the components of its parallelity bundle that are not 2-handles.

We define the *boundary graph* X_∂ which encodes information about the boundary of M' . It has two types of vertices: a vertex for each component of $\partial E(M', \gamma') - \mathcal{B}'$ (known as G -vertices, where ' G ' stands for 'guts'), and a vertex for each component of $\partial_h \mathcal{B}'$ (known as B -vertices, where ' B ' stands for 'bundle'). It has an edge for each boundary component of $\partial_v \mathcal{B}'$. Note that every component of $\partial_v \mathcal{B}'$ is an annulus, and so the edges of X_∂ come in pairs, where edges in a pair correspond to boundary components of the same annulus. Since $\partial_v \mathcal{B}'$ separates \mathcal{B}' from the remainder of $E(M', \gamma')$, each edge of X_∂ joins a B -vertex to a G -vertex. We orient the edge, so that it points from the B -vertex to the G -vertex.

We give each vertex of X_∂ an integer value, which we call its χ -value. This is the Euler characteristic of the corresponding component of $\text{cl}(\partial E(M', \gamma') - \mathcal{B}')$ or $\partial_h \mathcal{B}'$.

Each B -vertex is assigned the label R_- or R_+ , according to whether the corresponding component of $\partial_h \mathcal{B}$ lies in $R_-(M')$ or $R_+(M')$. Each edge is also assigned the label R_- or R_+ , according to whether the corresponding curve lies in $R_-(M')$ or $R_+(M')$. This label is the same as that of the B -vertex that it emanates from.

Clearly, X_∂ and its χ -values are constructible in time that is bounded above by a polynomial function of $h \log(w(S))$, for the following reasons. The G -vertices and edges can be readily constructed from the handle structure on $\text{cl}(M' - \mathcal{B}')$ given by Theorem 9.2. The B -vertices, the way that they are incident to the edges of X_∂ and their χ -values are also provided by the algorithm in Theorem 9.2.

It is perhaps convenient to view X_∂ as a graph embedded in $\partial M'$.

10.3. THE CONNECTIVITY GRAPH

The *connectivity graph* X_c has an edge for each component of $\partial_v \mathcal{B}'$. It has a vertex for each component of $M' - \partial_v \mathcal{B}'$. These are clearly of two types: components of \mathcal{B}' and components of $\text{cl}(M' - \mathcal{B}')$. We orient the edges so that they point from the former vertices to the latter ones.

Again, the graph X_c may be readily constructed using the data provided by Theorem 9.2.

Note that there is a morphism of graphs $X_\partial \rightarrow X_c$. This sends edges to edges, by sending each boundary component of $\partial_v \mathcal{B}'$ to the component of $\partial_v \mathcal{B}'$ that contains it. It sends vertices to vertices in a similar way.

10.4. THE ALGORITHM

We now explain how to modify M' and its handle structure. As explained in Section 10.1, the first stage is to attach a 2-handle along any u-suture that intersects \mathcal{B} , and enlarge \mathcal{B} to include these 2-handles. We then remove any component B of \mathcal{B} such that $\text{int}(\partial_v B)$ intersects ∂M . We then construct the boundary graph X_∂ and the connectivity graph X_c .

We now modify the manifold M' , and at the same time modify the boundary graph and the connectivity graph. We follow the same procedure used in Section 8 of [20], and so, for ease of reference, we use the same numbering of cases.

The algorithm starts by checking whether the hypothesis of the first case holds. If it does, then the procedure given in that case is followed. This changes the handle structure and the graphs X_c and X_∂ . The algorithm then goes back to the beginning, and checks whether Case 1 holds in this new structure. On the other hand, if Case 1 does not apply, then the algorithm proceeds to Case 2, and so on. In this way, a case is considered only if all the earlier cases do not apply.

Case 1. When there a B -vertex of X_∂ with χ -value 1.

This corresponds to component of $\partial_h \mathcal{B}'$ that is a disc. This therefore lies in a component of \mathcal{B}' that is an I -bundle over a disc.

At this stage, we remove this component of \mathcal{B}' and replace it by a 2-handle. We therefore remove the two corresponding vertices of X_∂ and the two edges that are incident to them. For each of these edges, we add 1 to the χ -value of the G -vertex that is incident to it. (If both edges are incident to the same G -vertex, we add to 2 to its χ -value.) We also remove the corresponding vertex of X_c and the edge that is incident to it.

We therefore now assume that every B -vertex has χ -value other than 1.

Case 2. When there is a B -vertex with no incident edges.

This corresponds to a component of \mathcal{B}' that is an I -bundle over a closed surface. If this component is incident to both $R_-(M')$ and $R_+(M')$ (which is information provided by the algorithm in Theorem 9.2), then there is a sequence of decompositions along product annuli and then product discs which takes it to a taut ball. So, we replace this component of \mathcal{B}' with a single 0-handle containing a single suture.

On the other hand, if this component of \mathcal{B}' is incident only to $R_-(M')$ or to $R_+(M')$, then our procedure checks the χ -value of this vertex. If it is zero, then this component of M' is an I -bundle over a torus or Klein bottle, and again there is a sequence of taut decompositions along annuli and discs taking it to a taut ball. If the χ -value of the vertex is non-zero, then this component of $E(M', \gamma')$ is not taut, and the algorithm terminates with a declaration to this effect.

We therefore now assume that every B -vertex is incident to some edge.

Case 3. When every vertex of X_∂ has non-positive χ -value, and for every pair of edges in X_∂ , one lies in R_- and one lies in R_+ .

This is exactly the situation where every component of $\partial_v \mathcal{B}'$ is a non-trivial product annulus, for the following reason. If one of these annuli was trivial, then one of its boundary components would, by definition, bound a disc in $R_-(E(M', \gamma'))$ or $R_+(E(M', \gamma'))$. The intersection between this disc and $\partial_v \mathcal{B}'$ is a collection of simple closed curves, and an innermost one bounds a disc which corresponds to a vertex of X_∂ with χ -value 1. This is contrary to our hypothesis.

Thus, we decompose (M', γ') along the annuli $\partial_v \mathcal{B}'$. This has the effect of removing \mathcal{B}' from the remainder of M' . Then \mathcal{B}' becomes a product sutured manifold, which we may further decompose along product discs, resulting in a collection of taut 3-balls. The manifold $\text{cl}(M' - \mathcal{B}')$ inherits a handle structure. Each component of $\partial_v \mathcal{B}'$ becomes a suture in this manifold that is not a u-suture. It is shown in [20] that this new handle structure has smaller complexity than that of M' . (Recall that a notion of complexity for handle structures was defined in [20] and summarised in Section 5.1.)

We may therefore assume that some annulus of $\partial_v \mathcal{B}'$ is not a non-trivial product annulus.

Case 4. Some G -vertex v has positive χ -value.

This corresponds to component of $\text{cl}(\partial E(M', \gamma') - \mathcal{B}')$ that is a disc D'_1 . This vertex v is therefore incident to a single edge e of X_∂ . This edge is sent to an edge of the connectivity graph by the map $X_\partial \rightarrow X_c$. If this edge of X_c is non-separating, then the algorithm terminates with the correct declaration that $E(M', \gamma')$ is not taut, for the following reason. Let A be the annulus of $\partial_v \mathcal{B}$ corresponding to this edge of X_c . If the edge is non-separating, then the annulus A is also non-separating, and hence so is the disc $A \cup D'_1$. Hence, $E(M', \gamma')$ contains an essential disc with boundary disjoint from γ' , and therefore $E(M', \gamma')$ is not taut.

Therefore, suppose that this edge of X_c is separating. Let e' be the other edge in this pair in X_∂ . Then e' is separating. The curve in $R_\pm(M)$ corresponding to e' bounds the disc $A \cup D'_1$ in M' . So if $E(M', \gamma')$ is taut, then this curve bounds a disc in $\partial E(M', \gamma')$ disjoint from the sutures. So, the algorithm computes the total χ -values of the vertices in each of the components of $X_\partial - e'$. It determines whether one of these components C has χ -value 1 and is disjoint from the sutures of $E(M', \gamma')$. If this applies to neither component of $X_\partial - e'$, then the algorithm terminates with the correct declaration that $E(M', \gamma')$ is not taut. If it applies to both components of $X_\partial - e'$, then this component of $\partial E(M', \gamma')$ is a sphere with no sutures. This is impossible unless $E(S)$ has a 2-sphere component, in which case $E(M', \gamma')$ is not taut. So, again the algorithm terminates. So, suppose that exactly one component C of $X_\partial - e'$ has χ -value 1 and is disjoint from the sutures of $E(M', \gamma')$. This corresponds to a disc D'_2 in $R_\pm(E(M', \gamma'))$. The argument now divides into two cases.

Case 4A. v does not lie in C .

Then D'_1 and D'_2 are disjoint. Then the union of D'_1 , D'_2 and the annulus A of $\partial_v \mathcal{B}$ incident to them together forms a 2-sphere. Since A is separating, so too is this 2-sphere. Using the Agol-Hass-Thurston algorithm, we can determine whether this 2-sphere separates components of $\partial M'$. If it does, then the algorithm terminates with the correct declaration that $E(M', \gamma')$ is not taut. On other hand, if the entirety of $\partial M'$ lies on one side of the 2-sphere, then we obtain an expression of $E(M', \gamma')$ as a connected sum, where one of the summands Y is closed. If $E(M', \gamma')$ is taut, then Y is the 3-sphere. Using the result of Schleimer [25] and Ivanov [12] that 3-sphere recognition is in NP, then we can certify that Y is the 3-sphere, once we have a handle structure for Y of uniform type and where the number of handles is bounded above by a linear function of h . We can construct this handle structure as follows. Let Z be the component of $\text{cl}(M' - A)$ containing D'_1 and D'_2 . Then Y is obtained from Z by attaching a 3-ball. The parts of Z lying in $M' - \mathcal{B}'$ already have a handle structure. The parts lying in \mathcal{B}' are all I -bundles over planar surfaces, since their horizontal boundaries lie in $D'_1 \cup D'_2$. So we can give them a handle structure using just 1-handles and 2-handles, knowing just the number of vertical boundary components of each component of $\mathcal{B}' \cap Z$ and whether or not it is product I -bundle. This is all information given by Theorem 9.2. Thus, we can construct a handle structure for Y of uniform type, and where the number of handles is bounded by a linear function of h . If Y is a 3-sphere, this can be certified in polynomial time. If so, then Z is a 3-ball, and we can therefore enlarge \mathcal{B} by adding the ball to it. Note that this may change (M', γ') since the ball was a subset of $E(M', \gamma')$ and not necessarily M' . So, the ball may contain some attached 2-handles. But clearly $E(M', \gamma')$ remains unchanged. Moreover, when we attach these 2-handles to the u -sutures, the resulting manifold still has an allowable hierarchy, which is just the original hierarchy.

The way that the algorithm processes this is as follows. The annulus A corresponds to an edge of the connectivity graph. Since A is separating in M' , removing this edge from the connectivity graph creates two components. The edge e points into one of these components, and this corresponds to a component Z of $M' - A$. Once we have verified the certificate that Z is a 3-ball, we discard all handles of $M' - \mathcal{B}$ that lie in Z , together with all components of \mathcal{B} lying in Z . We also discard the corresponding edges and vertices of X_∂ and X_c . The two edges of X_∂ corresponding to the two components of A are also removed. The adjacent B -vertices simply have 1 added to their χ -values.

Case 4B. v does lie in C .

Then the discs D'_1 and D'_2 are nested: D'_1 lies in D'_2 . The union of D'_2 , A and D'_1 pushed a little into the interior of M' is a separating 2-sphere. Again, we determine whether it separates components of $\partial M'$. If it does not, it forms an expression of $E(M', \gamma')$ as a connected sum with a closed summand Y . Again, if Y is the 3-sphere, this can be certified in polynomial time. In this case, we remove the component of $M' - A$ that contains the component of $\text{int}(\mathcal{B})$ incident to A . The handles of $M' - \mathcal{B}$ that lie in this component are removed, as are any components of \mathcal{B} . The connectivity graph is modified accordingly, by removing vertices and edges. Similarly, vertices and edges of X_∂ are removed. The vertex v is removed. But the G -vertex incident to e' is retained, and its χ -value is increased by 1. As shown in Lemma 7.3 in [20], this decreases the complexity of the handle structure.

Case 5. Every vertex of X_∂ has non-positive χ -value, and for some pair of edges of X_∂ , both lie in $R_-(M')$ or both lie in $R_+(M')$.

Let e_1 and e_2 be this pair of edges. Let A be the annulus of $\partial_v \mathcal{B}$ that contains the corresponding curves. Let A_1 and A_2 be parallel copies of this annulus, incoherently oriented in such a way that the region between them inherits four sutures. Isotope A_1 and A_2 a little so that they become standard surfaces.

Case 5A. e_1 and e_2 point into the same component of $X_\partial - \{e_1, e_2\}$ and this has χ -value 0.

Then (as shown in Case 5B(i) of Section 8 in [20]), we decompose along the component of $A_1 \cup A_2$ that is closest to \mathcal{B} .

Case 5B. e_1 and e_2 point out of the same component of $X_\partial - \{e_1, e_2\}$ and this has χ -value 0.

Then we decompose along the component of $A_1 \cup A_2$ that is furthest from \mathcal{B} .

Note that it is not possible for both Case 5A and Case 5B to hold because of our hypothesis that no component of M' has boundary a single torus with no sutures.

Case 5C. Neither Case 5A nor Case 5B applies.

Then we decompose along $A_1 \cup A_2$.

In each of these cases, the decomposition of $E(M', \gamma')$ along these annuli is taut if and only if $E(M', \gamma')$ is taut, for the following reason. If the decomposition is taut, then, by definition $E(M', \gamma')$ is taut. So, suppose that $E(M', \gamma')$ is taut. One cannot apply Proposition 2.5, because the decomposition is not along product annuli. However, the argument to establish tautness is given in Case 5 of Section 8 in [20]. Briefly, there it is shown that the only way that the decomposition can fail to be taut is if the resulting sutured manifold has a solid toral component with no sutures. Similarly, Proposition 4.9 implies that if (M', γ') admits an allowable hierarchy, then so does the manifold obtained by decomposing along non-trivial annuli disjoint from the sutures, unless they separate off a solid torus V with no sutures. The boundary of V intersects $\partial M'$ in a collection of annuli. Each annulus gives a component of $X_\partial - \{e_1, e_2\}$ and this has χ -value 0. The various arrangements for such a solid torus are considered in Case 5 of Section 8 in [20]. In each case, they are avoided by the suitable choice of decomposing surface, either A_1 , or A_2 , or $A_1 \cup A_2$.

It is also shown in Case 5 of Section 8 in [20] that these decompositions do not increase the complexity of any 0-handle. Moreover, if the complexity of a 0-handle is unchanged, then it is only modified by performing a trivial modification.

At the end of this process, we have created a sutured manifold (M'', γ'') with a handle structure \mathcal{H}'' that satisfies all the conditions of Theorem 10.1, except possibly (vi), which requires that \mathcal{H}'' be positive. We now explain how to guarantee this extra condition.

10.5. ARRANGING POSITIVITY

The procedure that we will use will decrease the number of handles in the handle structure. Hence, we must initially find an upper bound on the number of handles in our handle structure for (M'', γ'') . Note the number of handles of \mathcal{H}' that are not parallelity handles is at most ch , for some universal computable constant c . Each of the parallelity handles, other than isolated 2-handles incident to no other parallelity handle, is removed in the above process. Some are replaced with 2-handles. Hence, the number of handles in the resulting handle structure \mathcal{H}'' for (M'', γ'') is at most ch plus the number of components of $\partial_v \mathcal{B}$. This is at most $c'h$ for some universal computable constant c' .

The first stage in the procedure is to arrange that \mathcal{H}'' has no 3-handles. This is achieved using Procedure 6 of Section 6.1. Since no component of M'' is closed, then if \mathcal{H}'' contains a 3-handle, then it contains a 2-handle that intersects \mathcal{H}^3 in a single disc. Hence, Procedure 6 may be applied until there are no more 3-handles. None of the later modifications will introduce 3-handles, and so we may henceforth assume that none of the handle structures that we consider contains them.

Suppose that the handle structure is not positive. Then, as explained in Lemma 7.6 of [20], one of Procedures 1-6 of Section 6.1 can be applied, or the handle structure contains a parallelity handle that is not an isolated 2-handle. However, there are two slight discrepancies between the procedures referred to in Lemma 7.6 of [20] and Procedures 1-6 of Section 6.1. Firstly, a slight variant of Procedure 1 is used in [20], but the version that we present here works just as well; it has the advantage that one does not need to recognise 3-balls algorithmically. Secondly, Procedure 4 may be applied here only when the relevant 1-handle H_1 of \mathcal{H}'' is disjoint from the u-sutures, whereas in [20] there was no such constraint. If just one arc of $H_1 \cap \gamma$ lies in a u-suture or the arcs of $H_1 \cap \gamma$ lie in distinct u-sutures, then we attach a 2-handle along one of these u-sutures. The resulting manifold still has the same canonical extension and it still

admits an allowable hierarchy. However, we have introduced a handle and so we may have increased the complexity of the handle structure. But we can immediately apply Procedure 3, which cancels H_1 and the new 2-handle. The resulting handle structure has no greater complexity than the original one, and it has fewer handles. If both arcs of $H_1 \cap \gamma$ lie in the same u-suture, then we still attach a 2-handle along it. Then H_1 becomes a parallelity handle disjoint from γ . So we apply the procedures given in Sections 10.1, 10.2 and 10.3 again, except that we do not use Theorem 9.2. We again define \mathcal{B} to be the union of the parallelity handles that are not 2-handles. We calculate the boundary graph X_∂ and connectivity graph X_c directly, without using Theorem 9.2. This is possible because the handle structure has at most $c'h$ handles. Then we apply the algorithm given in Section 10.3.

So, we proceed initially, by applying Procedures 1-6 of Section 6.1 as many times as possible. If we come across a 1-handle intersecting γ in two arcs and disjoint from the 2-handles, but which intersects at least one u-suture, then we proceed as above.

This process may create new parallelity handles. If so, we apply the procedures given in Sections 10.1, 10.2 and 10.3 again, except that we do not use Theorem 9.2. Iterating in this way, we eventually reach the desired handle structure \mathcal{H}'' .

10.6. DECOMPOSITION ALONG NORMAL SURFACES

In this section, we give a variation of Theorem 10.1. Instead of decomposing a handle structure along a regulated surface, we decompose a triangulation along a normal surface.

Theorem 10.2. *Let \mathcal{T} be a triangulation of a connected compact orientable 3-manifold M with t tetrahedra. Suppose that every boundary component of M is a torus that is triangulated using just two triangles. Give M a sutured manifold structure (M, γ) where $R_-(M) = \partial M$ and $R_+(M) = \emptyset$. Let S be a taut normal surface properly embedded in M , such that the intersection between S and each component of ∂M is a (possibly empty) collection of coherently oriented essential curves. Let (M', γ') be obtained by decomposing along S . Then, there is an algorithm that takes as its input \mathcal{T} and (S) and either correctly asserts that $E(M', \gamma')$ is not taut, or supplies a handle structure \mathcal{H}'' for a sutured manifold (M'', γ'') with properties (i) - (viii) of Theorem 10.1. The algorithm runs in time that is bounded by a polynomial in $t \log(w(S))$.*

Proof. Each component of ∂M is a torus triangulated using two triangles. This hypothesis is made to ensure that ∂S is very controlled. It is shown in Lemma 3.5 in [14] that two normal curves in such a torus are isotopic if and only if they are normally isotopic. Furthermore, in Theorem 3.6 of [14], a parametrisation is given for all possible normal curves in the torus, and a consequence of this is that the intersection between any essential normal simple closed curve and any edge in the torus is a collection of points all with the same sign. We are assuming that the intersection between ∂S and any component of ∂M is a collection of coherently oriented essential curves. Thus, we deduce that the points of intersection between ∂S and any edge in ∂M all have the same sign.

We now possibly modify \mathcal{T} so that each tetrahedron Δ of \mathcal{T} contains at most one 2-simplex lying in ∂M . The only other alternative is that Δ contains two 2-simplices in ∂M , which comprise the entirety of a component of ∂M . We may then remove this tetrahedron from \mathcal{T} , and the result is a new triangulation \mathcal{T}' for a 3-manifold M' which is a copy of M . This new triangulation again satisfies the condition that every boundary component of M is a torus that is triangulated using just two triangles. The surface $S \cap M'$ is normal in \mathcal{T}' . We claim that there is a homeomorphism from M to M' which takes S to $M' \cap S$. This will imply that $M' \cap S$ satisfies the hypotheses of the theorem.

We can think of Δ as forming a collar on $\partial M'$. In particular, there is a natural product structure on Δ . We will show that $S \cap \Delta$ respects this product structure, which will establish the claim. There are four possible triangle types in Δ , and each of these respects its product structure. There are three possible square types, two of which respect the product structure. However, the third square type gives rise to a boundary-compression disc for S , which contradicts the hypothesis that S is taut.

So, we may assume that each tetrahedron Δ of \mathcal{T} contains at most one 2-simplex lying in ∂M . We now dualise \mathcal{T} to form a handle structure \mathcal{H} . Recall that each i -handle of \mathcal{H} arises from a $(3-i)$ -simplex of \mathcal{T} that does not lie entirely in ∂M . Hence, every 0-handle of \mathcal{H} is subtetrahedral. In particular, a tetrahedron that intersects ∂M in a triangle is dual to a 0-handle H_0 such that $H_0 \cap \mathcal{F}$ is as shown in Figure 23.

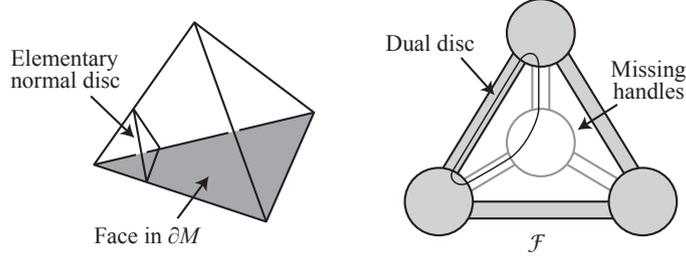


Figure 23: Dual handle structure near ∂M

The normal surface S determines a standard surface in \mathcal{H} . We will call this new surface S' , although it is just a copy of S . Note that S' need be neither normal nor regulated with respect to \mathcal{H} . For example, in Figure 23, an elementary disc in a 0-handle H_0 is shown which is dual to a normal triangle of S . This disc fails to satisfy Condition 3' in the definition of regulated and it fails to satisfy (i) in the definition of normality.

Let M' be the manifold obtained by decomposing along S' , and let \mathcal{H}' be the handle structure that it inherits from \mathcal{H} . Consider the parallelity handles of \mathcal{H}' for the pair $(M', N(S') \cap M')$. As explained in Section 9.2, these arise precisely in the space between two parallel normal discs of S . Hence, in each 0-handle H_0 of \mathcal{H} , at most 6 handles of $H_0 \cap \mathcal{H}$ are not parallelity handles in this sense. If we let \mathcal{B} denote the union of these parallelity handles, we can use Theorem 9.3 to compute a handle structure for $\text{cl}(M' - \mathcal{B})$ and, for each component B of \mathcal{B} , the genus and number of boundary components of the base surface for B , whether it is a product or twisted I -bundle and the location of $\partial_v B$.

However, there is another notion of parallelity handle, given in Section 9.1, for the sutured manifold (M', γ') . We claim that every parallelity handle H for the pair $(M', N(S') \cap M')$ is also a parallelity handle for (M', γ') . In the case where H is disjoint from ∂M , this is automatic. In the case where H intersects ∂M , it lies between two elementary normal discs of S that intersect ∂M , and we have arranged that these two discs are therefore coherently oriented. This is because the points of intersection between S and any edge in ∂M all have the same sign. Hence, H does indeed satisfy (ii) or (iii) in the definition in Section 9.1. Let \mathcal{B}' be the union of the parallelity handles for (M', γ') . Then, we have shown that $\mathcal{B} \subseteq \mathcal{B}'$.

We can readily compute the handles of $\text{cl}(M' - \mathcal{B})$ that lie in \mathcal{B}' . Therefore, we can determine a handle structure for $\text{cl}(M' - \mathcal{B}')$ and, for each component B of \mathcal{B}' , the genus and number of boundary components of the base surface for B , whether it is a product or twisted I -bundle and the location of $\partial_v B$.

This was exactly the information that we required to be able to proceed with the proof of Theorem 10.1. The remainder of the proof proceeds exactly as in that case. \square

10.7. MODIFYING A TRIANGULATION TO SIMPLIFY THE BOUNDARY TORI

A key hypothesis in Theorem 10.2 was that the triangulation of the 3-manifold restricts to a triangulation of each boundary torus with only two triangles. In the following result, we show how to arrange that we have a triangulation of this form.

Proposition 10.3. *Let M be a compact orientable 3-manifold with boundary a union of tori. Let \mathcal{T} be a triangulation of M with t tetrahedra. Then there is an algorithm to construct a triangulation \mathcal{T}' for M with at most $5t$ tetrahedra, and where every boundary torus is triangulated using two triangles. Moreover,*

if ϕ is a simplicial 1-cocycle on \mathcal{T} , then the algorithm computes a simplicial 1-cocycle ϕ on \mathcal{T}' in the same cohomology class as ϕ and satisfying $\|\phi'\|_1 \leq 3^{4t}\|\phi\|_1$. This algorithm runs in time that is bounded by a polynomial function of $t \log \|\phi\|_1$.

Proof. We will modify \mathcal{T} by attaching tetrahedra to ∂M . This will have the effect of changing the boundary triangulation by a Pachner move. Each time we do this, there is a simple way to extend ϕ to a cocycle ϕ'' on the larger triangulation. For most ways of attaching a tetrahedron, ϕ'' is determined by the fact that has to be a cocycle. But when we attach a tetrahedron along a single triangle, there is some choice, and so we simply declare that one of the new edges has evaluation zero under ϕ'' . It is easy to check the sum of the absolute values of the evaluations of the new edges is at most the sum of the values of ϕ on two edges in ∂M . Hence, $\|\phi''\|_1 \leq 3\|\phi\|_1$. We will show that we have to add at most $4t$ tetrahedra, and hence the final cocycle ϕ' satisfies $\|\phi'\|_1 \leq 3^{4t}\|\phi\|_1$. This will also show that the final triangulation has at most $5t$ tetrahedra.

It is well known that any two triangulations of a surface differ by a sequence of Pachner moves. However, we wish to do this algorithmically, and with control over the number of moves. Suppose that a triangulation of a component of ∂M has more than one vertex. Our strategy will be to apply at most 4 Pachner moves to this triangulation to create a triangulation with fewer vertices. Hence, the total number of Pachner moves that we will apply is at most 4 times the number of vertices in the boundary. For Euler characteristic reasons, this is half the number of triangles in the boundary. But each tetrahedron of M contributes at most 2 faces to the boundary. So, the total number of Pachner moves will be at most $4t$.

Suppose first that there is an edge in ∂M with the same triangle on both sides. The two edges of the triangle are then identified. Because ∂M is orientable, the edges are incompatibly oriented as one encircles the triangle. Hence, their common endpoint ends up as a valence one vertex v in ∂M . The remaining edge e of the triangle has a distinct triangle on the other side of it. Thus we may perform a 2-2 Pachner move centred at e . This has the effect of removing e and increasing the valence of v to 2. Incident to v there are now two distinct triangles. For each of these triangles, one of its edges is not incident to v . These edges cannot be the same edge of the triangulation, because this would form a triangulation of the 2-sphere with two triangles. Hence, we may perform a 2-2 Pachner move along one of these edges. This increases the valence of v to 3. We can then perform a 3-1 Pachner move at v , and this reduces the number of vertices in ∂M by 1, as required. So, we may assume that every edge in ∂M is adjacent to two distinct triangles, and hence one may perform a 2-2 Pachner move along it.

Using the fact that the torus has zero Euler characteristic, we can find a vertex v of its triangulation with valence at most 6. Our aim is reduce the valence of this vertex to 3 and then apply a 3-1 Pachner move to remove it.

Suppose first that there is an edge e that starts and ends at v . If we perform a 2-2 Pachner move along e , then the two endpoints of e no longer contribute to the valence. Thus, we reduce the valence of v unless the new edge also starts and ends at v . In this case, the two triangles adjacent to e form a square, and the four sides of the square must be identified in pairs, in order that the valence of v is at most 6. Thus, in this case, we deduce that the triangulation of the torus already has a single vertex, which is contrary to our assumption.

On the other hand, if no edge starts and ends at v , then the star of v is an open disc. We may then apply at most three 2-2 Pachner moves supported in this disc to reduce the valence of v to 3. We then can apply a 3-1 Pachner move to remove v .

Thus in all cases, after at most 4 Pachner moves, we can reduce the number of vertices by 1. \square

11. REDUCTION TO THE ATOROIDAL AND SEIFERT FIBRED CASES

At certain points in the previous sections, we made the assumption that the manifold is atoroidal. In this section, we show why it suffices to consider this case and also the case where the manifold is Seifert fibred.

11.1. THE CANONICAL TORI

Recall that a surface S properly embedded in a compact orientable 3-manifold M is *essential* if it is incompressible, boundary-incompressible and no component is parallel to a subsurface of ∂M .

A properly embedded torus in M is *canonical* if it is essential and furthermore it can be ambient isotoped off any other essential torus.

Canonical tori are also called JSJ tori, since Jaco, Shalen and Johannson developed their theory. In particular, they proved the following result [15, 17].

Theorem 11.1. *Let M be a compact orientable irreducible 3-manifold with incompressible boundary. Then up to ambient isotopy, there exists finitely many canonical tori in M . The union T of these tori is properly embedded. Each component of $M - \text{int}(N(T))$ is either atoroidal or Seifert fibred.*

11.2. THE BEHAVIOUR OF THURSTON NORM UPON CUTTING ALONG ESSENTIAL TORI

Proposition 11.2. *Let M be a compact orientable irreducible 3-manifold with incompressible boundary. Let T be a union of finitely many disjoint incompressible tori properly embedded in M . Let c be a class in $H^1(M)$, and let $i^*(c)$ be its image under the homomorphism $i^*: H^1(M) \rightarrow H^1(M - \text{int}(N(T)))$ induced by inclusion. Then the Poincaré duals of c and $i^*(c)$ have equal Thurston norm.*

Proof. Let S be a representative for the Poincaré dual of c , with minimal Thurston complexity. Since M is irreducible and has incompressible boundary, we may assume that S has no sphere or disc components. If $S \cap T$ contains any simple closed curves that bound discs in T , then these may be removed without increasing the Thurston complexity of S . So, we may assume that $S \cap T$ is essential in T . Hence, by the incompressibility of T , none of these curves bounds a disc in M . In particular, no component of $S - \text{int}(N(T))$ is a disc. Therefore

$$\chi_-(S) = -\chi(S) = -\chi(S - \text{int}(N(T))) = \chi_-(S - \text{int}(N(T))).$$

Hence, the Thurston norm of $[S - \text{int}(N(T))]$ is at most that of $[S]$. Note that $S - \text{int}(N(T))$ is a representative of the Poincaré dual to $i^*(c)$.

Now consider a surface S' representing the Poincaré dual to $i^*(c)$ with minimal Thurston complexity. We may assume that $S' \cap \partial N(T)$ is a collection of essential curves. Moreover, we may arrange, by attaching annuli if necessary, that S' intersects each component of $\partial N(T)$ in a collection of coherently oriented curves. We wish to use S' to build a representative for the Poincaré dual to S . We first note that if T_1 and T_2 are the components of $\partial N(T)$ coming from the same component of T , then $S' \cap T_1$ and $S' \cap T_2$ patch together under the gluing map $T_1 \rightarrow T_2$. This is because $S' \cap T_1$ and $S' \cap T_2$ are determined up to isotopy by the restriction of c to $H^1(T)$. Hence, when we patch together the tori in $\partial N(T)$ to form M , we can get a properly embedded surface S . However, it is not immediately clear that S is Poincaré dual to c . Indeed, all we can say is that c and the Poincaré dual to $[S]$ have the same image under i^* . We have an exact sequence

$$H^1(M, M - \text{int}(N(T))) \rightarrow H^1(M) \xrightarrow{i^*} H^1(M - \text{int}(N(T))).$$

Now, $H^1(M, M - \text{int}(N(T)))$ is isomorphic to $H^1(N(T), \partial N(T))$ by excision, and this is isomorphic to $H_2(N(T))$ by Poincaré duality. Hence, the difference between $[S]$ and the Poincaré dual to c is given by some copies of components of T . But we can introduce these components, and form their double-curve

sum with S , forming a surface S'' , say. Then $\chi_-(S'') = \chi_-(S) = \chi_-(S')$. Hence, the Thurston norm of the Poincaré dual to c is at most that of the Poincaré dual to $i^*(c)$. \square

11.3. THE CANONICAL TORI IN NORMAL FORM

The canonical tori may be placed in normal form with respect to any triangulation. In fact, it has long been known that one may control the weight of this normal surface. For example, an algorithm was given by Jaco and Tollefson for constructing the canonical tori [16]. An explicit upper bound on the weight of the canonical tori in normal form was given by Mijatović (Proposition 2.4 in [23]), as follows.

Theorem 11.3. *Let M be a compact orientable irreducible 3-manifold with incompressible boundary. Let \mathcal{T} be a triangulation of M with t tetrahedra. Then the union of the canonical tori may be placed in normal form with respect to \mathcal{T} , so that it contains at most 2^{80t^2} elementary normal discs.*

11.4. A TRIANGULATION FOR THE EXTERIOR OF NORMAL TORI

Theorem 11.4. *Let M be a compact orientable irreducible 3-manifold with incompressible boundary. Let \mathcal{T} be a triangulation of M with t tetrahedra. Let T be a union of disjoint properly embedded normal tori. Then there is an algorithm to build a triangulation \mathcal{T}' for $M' = M - \text{int}(N(T))$, which runs in time at most a polynomial function of $t \log w(T)$, and where this triangulation has at most $200t$ tetrahedra.*

Proof. Let \mathcal{H} be the handle structure of M that is dual to \mathcal{T} . The normal surface T in \mathcal{T} determines a standard surface, also called T , in \mathcal{H} . So $M - \text{int}(N(T))$ inherits a handle structure \mathcal{H}' . Let \mathcal{B} be the parallelity bundle of \mathcal{H}' for the pair $(M', \partial N(T))$. Let $\partial_v \mathcal{B}$ be its vertical boundary.

For each 0-handle H_0 of \mathcal{H} , all but at most 6 handles of $H_0 \cap \mathcal{H}'$ are parallelity handles. Similarly, for each 1-handle H_1 of \mathcal{H} , all but at most 4 handles of $H_1 \cap \mathcal{H}'$ are parallelity handles. We may triangulate the union of these handles using at most 60 (respectively, 40) tetrahedra, so that the triangulations match on the intersection between handles and so that the intersection with $\partial_v \mathcal{B}$ is simplicial.

It is easy to check that, for each 0-handle H_0 of \mathcal{H} , the number of components of $H_0 \cap \partial_v \mathcal{B}$ is at most 10. Hence, $\partial_v \mathcal{B}$ has at most $10t$ components.

We now apply Theorem 9.3, which provides an algorithm that determines, for each component B of \mathcal{B} , the genus and number of boundary components of its base surface, whether B is a product or twisted I-bundle, and the location of $\partial_v B$ in $\text{cl}(M' - \mathcal{B})$. Using this information, we may build an abstract triangulation for each component of \mathcal{B} that agrees with the one that we have built for $\partial_v \mathcal{B}$. Since the genus of the base surface of each component of \mathcal{B} is at most one, and the number of boundary components is at most $10t$, we may arrange that the number of tetrahedra in this triangulation of \mathcal{B} is at most $100t$. Attaching this to the triangulation of $\text{cl}(M' - \mathcal{B})$ already constructed, we obtain the required triangulation of M' .

As a result of Theorem 9.3, each of these steps takes time that is bounded above by a polynomial function of $t \log w(T)$. \square

Theorem 11.4 gives an algorithm to construct a triangulation for the exterior of the canonical tori, as long as that these tori are provided non-deterministically as a normal surface.

Note however, we are not at this stage claiming to be able verify that a given collection of normal tori T is the canonical collection. In particular, we have not checked that T is incompressible, nor have we checked that each component of the exterior of T is atoroidal or Seifert fibred.

We will use Theorem 11.4 to reduce the main theorem to the case where the 3-manifold is atoroidal or Seifert fibred. But to be able to do this, we need to be able to restrict the given 1-cocycle on M to an explicit 1-cocycle on M' .

Addendum 11.5. Let $M, \mathcal{T}, t, \mathcal{T}'$ and M' be as in Theorem 11.4. Let ϕ be a simplicial 1-cocycle on \mathcal{T} . Let $i: M' \rightarrow M$ be inclusion. Then we may construct a simplicial 1-cocycle on \mathcal{T}' with $\|\phi'\|_1 \leq 1200t\|\phi\|_1$, where $\|\cdot\|_1$ denotes the ℓ_1 norm of a cocycle and such that the Poincaré duals of $i^*([\phi])$ and $[\phi']$ have the same Thurston norm. This can be achieved in time at most a polynomial function of $t \log w(T) \log \|\phi\|_1$.

Proof. A cocycle representing $i^*([\phi])$ may be formed by first dualising ϕ to form a surface S in M , then intersecting S with M' to form a surface S' , then dualising this to give $i^*([\phi])$. In order to complete this successfully, we first control the position of S in M . From the cocycle ϕ , there is a natural way of building S , as follows. For each edge e of the triangulation, one orients e so that $\phi(e)$ is non-negative and one places $|\phi(e)|$ points in the interior of e . For each face of the triangulation, one inserts arcs joining these points, in such a way that arcs start and end on coherently oriented edges. This is possible because ϕ is a cocycle, and so the total evaluation of the edges encircling the face is zero. For each tetrahedron, we have now specified a collection of simple closed curves in its boundary. Each bounds an elementary normal disc. The union of these discs we take to be S .

We perform a normal isotopy of S so that its intersection with T is in general position. More specifically, we arrange that the intersection between any normal disc of S and any normal disc of T is at most one arc. We may also arrange that $S \cap \mathcal{B}$ is vertical in \mathcal{B} . Then $S - \text{int}(N(T))$ is dual to $i^*([\phi])$.

We now need to specify the triangulation \mathcal{T}' a little more precisely than in the proof of Theorem 11.4. The parts of $\text{cl}(M' - \mathcal{B})$ were triangulated in a relatively simple way. It is clear that we may choose this triangulation so that, for each edge in this part of the triangulation, the number of times it intersects $S - \text{int}(N(T))$ is at most the number of elementary discs of S in the tetrahedron containing it. We take this algebraic intersection number to be ϕ' on that edge. So, the evaluation of each such edge under ϕ' is at most $\|\phi\|_1$.

The remainder of \mathcal{T}' was formed by triangulating \mathcal{B} , which is an I -bundle over a surface F . Using Theorem 9.3, we know the topological type for each component of F . We pick a triangulation for F with a certain number of vertices in its boundary and none in its interior. Then, for each simplex σ of this triangulation, its inverse image in \mathcal{B} is of the form $\sigma \times I$. We can triangulate this in a simple way. When σ is a 1-simplex, we triangulate $\sigma \times I$ using two 2-simplices. When σ is a 2-simplex, we triangulate $\sigma \times I$ using three 3-simplices. Note that $\partial_v \mathcal{B}$ then inherits a triangulation, which agrees with the one arising from the triangulation on $\text{cl}(M' - \mathcal{B})$, as long the number of vertices in each component of ∂F is chosen correctly.

Thus, we have triangulated M' and have defined ϕ' on $\text{cl}(M' - \mathcal{B})$. We still need to define ϕ' on \mathcal{B} . Now $S - \text{int}(N(T))$ is vertical in \mathcal{B} . But to determine this precise location of this surface might take too long. So, instead, we pick any vertical transversely oriented surface in \mathcal{B} that agrees with S on $\partial_v \mathcal{B}$. This can be achieved by choosing arcs in F transverse to the 1-skeleton of F and that join the points of $S \cap F$ with the correct orientations. These arcs can be chosen so that the number of times they intersect each edge of F is at most $|S \cap \partial_v \mathcal{B}|$. They then specify a collection of transversely oriented vertical discs in \mathcal{B} . We define ϕ' on the edges in \mathcal{B} to be the intersection numbers of these edges with these discs. Hence, $\|\phi'\|_1 \leq 1200t\|\phi\|_1$, because there are at most $1200t$ edges of \mathcal{T}' and each edge has evaluation at most $\|\phi\|_1$ under ϕ' .

Now, $[\phi']$ and $i^*([\phi])$ need not be equal cohomology classes, because their dual surfaces may differ in \mathcal{B} . But their difference is represented by a union of vertical annuli in \mathcal{B} . Hence, the duals have the same Thurston norm. \square

12.1. PRODUCTS

Throughout this paper, we have used allowable hierarchies, which terminate in product sutured manifolds. It will be important that we can certify that a given sutured manifold is indeed a product. The existence of such a certificate is presented in this subsection.

Theorem 12.1. *Let (M, γ) be a product sutured manifold with handle structure \mathcal{H} of uniform type. Let h be the number of 0-handles of \mathcal{H} . Suppose that (M, γ) is a product sutured manifold. Then there is a certificate that proves that (M, γ) is indeed a product. This can be verified in time that is bounded above by a polynomial function of h .*

A complete collection of product discs for a product sutured manifold (M, γ) is a collection of disjoint product discs D , such that decomposing along D gives a collection of taut 3-balls.

Theorem 12.2. *Let (M, γ) be a sutured manifold with handle structure \mathcal{H} of uniform type. Then (M, γ) contains a complete collection D of product discs in normal form with weight satisfying $w(D) \leq c^{h^2}$, where h is the number of 0-handles of \mathcal{H} and c is a universal computable constant.*

We will prove this result in the next subsection. But first, we show how it can be used to prove Theorem 12.1.

We are given a handle structure \mathcal{H} for the product sutured manifold (M, γ) . We need to certify that it is indeed a product. The certificate that we use will be of the following form:

- (i) a complete collection D of product discs as in Theorem 12.2;
- (ii) a handle structure \mathcal{H}' for a sutured (M', γ') with at most ch handles; in fact, this will be the sutured manifold obtained by decomposing (M, γ) along D ;
- (iii) a certificate that M' is a collection of 3-balls, as provided by Schleimer [25] or Ivanov [12].

The algorithm to verify this certificate is as follows:

- (1) Verify that D is a collection of product discs, using the algorithm of Agol-Hass-Thurston to determine the topological types of the components of D , and to verify that each component intersects γ exactly twice.
- (2) Denote the manifold obtained by decomposing (M, γ) along D by (M_D, γ_D) . This inherits a handle structure \mathcal{H}_D . Denote the parallelity bundle of the pair $(M_D, M_D \cap N(D))$ by \mathcal{B} . Apply the algorithm given in Theorem 9.3 (adapted in the obvious way to handle structures of uniform type) to determine a handle structure for $\text{cl}(M_D - \mathcal{B})$ and to determine, for each component B of the parallelity bundle \mathcal{B} , the number of boundary components of its base surface, whether B is a product or twisted I -bundle, and the location of $\partial_v B$ in $\text{cl}(M_D - \mathcal{B})$. The total number of 0-handles in the handle structure of $\text{cl}(M_D - \mathcal{B})$ is at most ch for some universal computable constant c .
- (3) Using the information provided in (2), express \mathcal{B} as a union of parallelity 1-handles and 2-handles, and thereby form a handle structure for (M_D, γ_D) . Verify that this is equal to the given handle structure \mathcal{H}' .
- (4) A verification of the certificate that M' is a collection of 3-balls, using the algorithm of Schleimer [25] or Ivanov [12], together with a verification that each component of $R_{\pm}(M')$ is a disc.

The time taken to complete this verification is at most a polynomial function of $h \log(w(D))$, which by Theorem 12.2, at most a polynomial function of h .

12.2. PRODUCT DISCS IN NORMAL FORM

In this subsection, we prove Theorem 12.2. Our method is rather similar to that used by Mijatović in the proof of Theorem 11.3, but somewhat more straightforward.

Let (M, γ) be a connected product sutured manifold with a handle structure \mathcal{H} of uniform type. We may suppose that it is not a 3-ball, as in this case, we may take D to be empty.

The first step in the proof is to show that there is some essential product disc P_1 that is in normal form with respect to \mathcal{H} and that is fundamental. This is well known. A proof, in the related case where M is triangulated, is given in Lemma 4.1.18 of [22]. Hence, by Theorem 8.1, the weight $w(P_1)$ is at most $(c_1)^h$, where h is the number of 0-handles of \mathcal{H} and c_1 is a universal computable constant.

We now decompose (M, γ) along P_1 , giving a manifold (M_2, γ_2) . This inherits a handle structure \mathcal{H}_2 . We denote by \mathcal{B} the parallelity bundle for the pair $(M_2, M_2 \cap N(P_1))$ with handle structure \mathcal{H}_2 .

Our next goal is to find an essential product disc P_2 in (M_2, γ_2) that is disjoint from the copies of P_1 in ∂M_2 . It therefore is an essential product disc in (M, γ) also. Our aim to ensure that P_2 is in fact normal in \mathcal{H} , and that there is a good upper bound on its weight. Simple normal surface theory applied to the handle structure \mathcal{H}_2 will not suffice. This is because the bound on the weight of a fundamental normal surface in \mathcal{H}_2 given by Theorem 8.1 is of the form $(c_1)^{h_2}$, where h_2 is the number of 0-handles of \mathcal{H}_2 , but h_2 may not be bounded by a linear function of h .

So, we develop a modified normal surface theory in \mathcal{H}_2 . We consider normal surfaces S properly embedded in \mathcal{H} that are disjoint from P_1 . For any such surface S , its intersection with \mathcal{B} is horizontal and so, for each component B of \mathcal{B} , $B \cap S$ intersects every fibre of B the same number of times. We call this the number of *sheets* of $B \cap S$.

One can encode such a surface by means of a vector, which counts the number of elementary discs of each type in each 0-handle in $\text{cl}(M_2 - \mathcal{B})$ and the number of sheets of $B \cap S$ for each component B of \mathcal{B} . The equations come in two types. Equations of the first type arise from 1-handles in $\text{cl}(M_2 - \mathcal{B})$, and they simply specify that the elementary discs in the adjacent 0-handles can be attached to suitable elementary discs in the 1-handles. Equations of the second type are associated with components B of \mathcal{B} . They assert that, for each 0-handle H_0 in $\text{cl}(M_2 - \mathcal{B})$ and each component of $\partial_v B \cap H_0$, the number of arcs of intersection between S and this component is equal to the number of sheets of $B \cap S$.

Because of this algebraic structure, one may speak of such a surface being fundamental. We may find an essential product disc P_2 which is normal in \mathcal{H} and fundamental in the above sense, using the proof of Lemma 4.1.18 of [22]. Then Theorem 8.1 gives an upper bound $(c_1)^h$ on the size of each co-ordinate in the vector for P_2 . However, care must be taken when translating this to a bound on the weight of P_2 as a normal surface in \mathcal{H} , because the part of P_2 running through \mathcal{B} also contributes to this weight. For each component B of \mathcal{B} , the number of sheets of $B \cap P_2$ is bounded above by $(c_1)^h$. Each sheet of $B \cap P_2$ gives rise to a number of elementary discs of P_2 , which is at most the sum of the co-ordinates of P_1 . So, $w(P_2) \leq (c_1)^h (c_1)^h$.

We now repeat by cutting M_2 along P_2 to give another 3-manifold M_3 . We can view this as the result of cutting M along $P_1 \cup P_2$. The above argument gives another essential product disc P_3 , and the number of elementary discs of P_3 is at most $(c_1)^{3h}$.

We repeat this process until we have no more essential product discs, in other words, when we have decomposed (M, γ) to a collection of balls. The number of times that we iterated this procedure was at most $c_2 h$, for some universal computable constant c_2 , since this is an upper bound for the number of disjoint non-parallel properly embedded normal surfaces in \mathcal{H} . So, we end with a complete collection D of essential normal product discs in (M, γ) with total weight at most $((c_1)^h)^{c_2 h + 1}$, as required. \square

12.3. SEIFERT FIBRE SPACES

As explained in the previous section, we will cut our given compact orientable 3-manifold along its canonical tori. Each component of the resulting 3-manifold M' will be atoroidal or Seifert fibred. The reason for doing this is that many of the arguments given in Section 3 and 4 required the manifold to be atoroidal. However, that still leaves the case of Seifert fibred manifolds. In this subsection, we will explain how our main theorem is proved in this case.

Theorem 12.3. *Let M be a connected Seifert fibre space other than a solid torus, with a triangulation \mathcal{T} having t tetrahedra. Let ϕ be a simplicial 1-cocycle on M , and let m be the Thurston norm of the Poincaré dual of $[\phi]$. Then there is a certificate that certifies that the Thurston norm of this class is indeed m and that M is irreducible and has incompressible boundary. The algorithm to verify this certificate can be completed in time that is bounded above by a polynomial function of $t \log \|\phi\|_1 \log(m + 2)$.*

Note that we do not require the Seifert fibre space structure to be provided to us in any way. All we need for the existence of the certificate is that there is some Seifert fibration on M .

The first stage in the algorithm to apply Proposition 10.3 to obtain a triangulation (which we will also call \mathcal{T}) for M in which every boundary torus has precisely two triangles.

Any class in $H_2(M, \partial M)$ is represented by a compact oriented incompressible surface. There is a well known classification of compact orientable incompressible surfaces in a Seifert fibre space M . Any such surface is isotopic to a surface that is *horizontal* or *vertical*. By definition, a surface is *horizontal* if it is everywhere transverse to the fibres. A surface is *vertical* if it is a union of fibres. The proof divides into these two cases.

Case 1. The Poincaré dual of ϕ is represented by a horizontal surface S .

Then $M - \text{int}(N(S))$ is a product I -bundle over S . This implies that S is in fact a fibre in a fibration of M over the circle. Hence, it has minimal Thurston complexity in its homology class.

Using Theorem 8.3, there is a taut normal surface S representing this class such that $w(S) \leq k^{t^2} \|\phi\|_1$, where k is a universal computable constant. The vector (S) will form part of our certificate. Using Theorem 9.4, we can verify that ∂S contains no component that bounds a disc in ∂M , in time at most a polynomial function of $t \log w(S)$. Let (M', γ') be the product sutured manifold that is obtained by decomposing along S . Theorem 10.2 is applied, and it supplies a positive handle structure \mathcal{H}'' for a sutured manifold (M'', γ'') that is obtained from (M', γ') by decomposing along some non-trivial annuli disjoint from γ' . According to Theorem 10.2, $(M'', \gamma'') = E(M'', \gamma'')$ is taut. The time taken to complete the algorithm in Theorem 10.2 is at most a polynomial function of $t \log w(S)$. Since each annulus is vertical in the product structure or boundary-parallel, (M'', γ'') is again a product sutured manifold. Theorem 10.2 provides a handle structure for (M'', γ'') of uniform type, with at most ct 0-handles, for a universal computable constant c .

Using Theorem 12.1, there is a certificate which establishes that (M'', γ'') is a product sutured manifold. The time that it takes to verify this certificate is bounded above by a polynomial function of the number of 0-handles, and so by a polynomial function of t . Assuming that this certificate is correctly verified, then by Theorem 10.2, (M', γ') is also taut. Therefore, S is taut and (M, γ) is taut. So, by Lemma 2.2, we know that S has minimal Thurston complexity in its class in $H_2(M, \partial M)$. So, we have established that the Thurston norm of this class is indeed m .

Case 2. The Poincaré dual of ϕ is represented by a vertical surface.

This is a union of annuli and tori. In particular, the Thurston norm of this class is zero. By Theorem 8.3, there is a taut normal surface S representing this class such that $w(S) \leq k^{t^2} \|\phi\|_1$. This surface also has zero Thurston complexity. Since the Thurston complexity of S may be computed, using the Agol-Hass-Thurston algorithm, in time bounded above by $t \log(w(S))$, we may readily verify that it is zero. Thus, the part of Theorem 12.3 concerned with Thurston norm is completed.

However, we will still need to verify that M is irreducible and has incompressible boundary, or, in other words, that (M, γ) is taut, where $R_-(M) = \partial M$ and $R_+(M) = \emptyset$. In fact, we will verify that the decomposition $(M, \gamma) \xrightarrow{S} (M', \gamma')$ is taut. We apply Theorem 10.2 to find a handle structure \mathcal{H}'' for a sutured manifold (M'', γ'') that is obtained from (M', γ') by decomposing along some annuli disjoint from γ' . This handle structure is of uniform type and has at most ct 0-handles. According to Theorem 10.2, this is taut if and only if (M', γ') is. Note that M'' is Seifert fibred. Since it has non-empty boundary, there is some properly embedded, orientable surface that is horizontal in M'' . We may find such a surface S_2 that is fundamental. In particular, $w(S_2)$ is at most k^{ct} , by Theorem 8.2. We may now revert to Case 1, which provides a certificate to establish that (M'', γ'') is taut. This is verifiable in time at most a polynomial function of $ch \log w(S_2)$, and this is at most a polynomial function of t . \square

13. THE CERTIFICATE AND ITS VERIFICATION

In this section, we complete the proof of Theorem 1.5, and of Theorems 1.1 and 1.3, by describing the NP algorithm for determining the Thurston norm of a homology class. We also give the proof of Theorem 1.6 by showing how to certify that a compact orientable irreducible 3-manifold with toroidal boundary and positive first Betti number is indeed irreducible.

13.1. A DESCRIPTION OF THE CERTIFICATE

In Theorem 1.5, we are given, as an input, a triangulation \mathcal{T} for the 3-manifold M , a simplicial 1-cocycle ϕ and an integer m . Our aim is determine whether the Thurston norm of the Poincaré dual of $[\phi]$ is m . In Theorem 1.6, we are given a triangulation \mathcal{T} for a 3-manifold M , and our aim is to certify that M is irreducible. In both cases, let t be the number of tetrahedra in \mathcal{T} .

In the situation of Theorem 1.6, we are assuming that $b_1(M) > 0$, and so there is a simplicial cocycle representing a non-trivial class in $H^1(M)$. We may find such a cocycle ϕ such that $\|\phi\|_1$ is at most an exponential function of t . Hence, the Thurston norm m of its Poincaré dual is at most an exponential function of t . In Theorem 1.6, ϕ and m will form part of our certificate. In both theorems, we will show how to certify that M is irreducible and that the Thurston norm of the Poincaré dual of $[\phi]$ is m .

The following is our certificate. This consists of various pieces of data, which are provided non-deterministically:

- (i) a solution (T) , which may be zero, to the normal surface equations in \mathcal{T} satisfying the quadrilateral constraints, and with weight at most 2^{2+80t^2} ; in fact, T will be the canonical tori for M ;
- (ii) a triangulation \mathcal{T}' for a 3-manifold M' with at most $1000t$ tetrahedra, and such that every boundary component is triangulated using just two triangles; in fact, this 3-manifold will be $M - \text{int}(N(T))$;
- (iii) a simplicial 1-cocycle ϕ' on \mathcal{T}' such that $\|\phi'\|_1 \leq 3^{800t} 1200t \|\phi\|_1$; it will in fact have the same Thurston norm as the dual of $[\phi]$;
- (iv) a decomposition of M' into two subsets M'_1 and M'_2 , which are unions of components of M' ; these will in fact be the atoroidal and Seifert fibred components, respectively; if any component of M' is both atoroidal and Seifert fibred, it is placed in M'_2 rather than M'_1 ;
- (v) two integers m_1 and m_2 that sum to m ; these will be the Thurston norms of classes in $H_2(M'_1, \partial M'_1)$ and $H_2(M'_2, \partial M'_2)$;
- (vi) a certificate, verifiable by the algorithm in Theorem 12.3, that the Thurston norm of the Poincaré dual of $\phi'|_{M'_2}$ is m_2 and that M'_2 is irreducible and has incompressible boundary;
- (vii) a solution S_1 for the normal surface equations in the triangulation of M'_1 , satisfying the compatibility constraints, such that each co-ordinate is at most $c^{t^2} \|\phi\|_1$; here c is some universal computable

constant;

- (viii) handle structures $\mathcal{H}_2, \dots, \mathcal{H}_{n+1}$ for sutured manifolds $(M_2, \gamma_2), \dots, (M_{n+1}, \gamma_{n+1})$; these must be positive and of uniform type and the number of handles in each \mathcal{H}_i is at most ct ; here n is at most ct ;
- (ix) for each i between 2 and n , a solution (S_i) for the boundary-regulated matching equations for \mathcal{H}_i , satisfying the compatibility constraints, such that each co-ordinate of each solution is at most c^t ;
- (x) for each i between 1 and n , a choice of transverse orientations on at most ct elementary discs of S_i ;
- (xi) a certificate that (M_{n+1}, γ_{n+1}) is a product sutured manifold, as provided by Theorem 12.1;
- (xii) certificates that certify that certain 3-manifolds with uniform handle structures are 3-balls, as provided by the work of Schleimer [25] or Ivanov [12]; the number of such 3-manifolds is bounded by a polynomial function of t ; the number of handles in each manifold is bounded by a linear function of t .

13.2. THE ALGORITHM TO VERIFY THIS CERTIFICATE

This is as follows:

- (1) Verification that T is a union of tori, using the Agol-Hass-Thurston algorithm [2].
- (2) Verification that \mathcal{T}' is the triangulation of $M' = M - \text{int}(N(T))$ provided by Theorem 11.4 and Proposition 10.3.
- (3) Verification that ϕ' is the cocycle provided by Addendum 11.5 and Proposition 10.3, where $[\phi']$ and $i^*([\phi])$ have duals with the same Thurston norm;
- (4) Verification that $m_1 + m_2 = m$.
- (5) Verification of the certificate in (vi) using Theorem 12.3.
- (6) Verification that no component of ∂S_1 bounds a disc in $\partial M'$, using Theorem 9.4.
- (7) The determination of the components of S_1 and their Euler characteristic, using the Agol-Hass-Thurston algorithm, and the verification that each component of S_1 has non-positive Euler characteristic.
- (8) Verification that $\chi_-(S_1) = m_1$.
- (9) For each i between 1 and n , the verification that each component of S_i is normally parallel to a unique component that contains an elementary disc with a pre-assigned transverse orientation given in (x), again using the Agol-Hass-Thurston algorithm.
- (10) The verification that S_i has a transverse orientation compatible with the transverse orientations given in (x) and (for $i > 1$) the transverse orientations on the discs intersecting ∂M_i , using Theorem 9.6 or Theorem 9.7.
- (11) Using Theorem 9.7, the computation of the signed intersection number between each oriented edge of \mathcal{T}' and S_1 .
- (12) Verification that S_1 is Poincaré dual to $\phi'|M'_1$, by checking that, for each simplicial loop ℓ in some generating set for $H_1(M'_1)$, $\phi(\ell) = \ell \cdot S_1$.
- (13) For each i between 2 and n , verification that no component of S_i is a planar surface that is disjoint from γ_i and that has all but one of its boundary curves trivial, using Corollary 9.5.
- (14) For each i between 2 and n , verification that each decomposing surface S_i is allowable, by checking that it is disjoint from the u-sutures of M_i and that any trivialising planar surface of a trivial curve

of ∂S_i is correctly oriented, using Corollary 9.5.

- (15) For each i between 1 and n , the application of the algorithm in Theorem 10.1 or 10.2 to \mathcal{H}_i and S_i . If the algorithm declares ‘not taut’, then the certificate is invalid. On the other hand, if the algorithm produces a handle structure \mathcal{H}'' , then the next stage is a verification that this is combinatorially equal to \mathcal{H}_{i+1} . In order to complete the algorithm in Theorem 10.1 or 10.2, certain 3-manifolds need to be certified as 3-balls. These are the manifolds supplied in (xii) and their certificates are verified using the algorithm of Schleimer [25] or Ivanov [12].
- (16) A verification of the certificate in (xi) using Theorem 12.1, together with a verification that no component of $R_{\pm}(M_{n+1})$ is planar and has each boundary component a u-suture.

13.3. PROOF OF THEOREMS 1.5 AND 1.6.

In order to prove Theorems 1.5 and 1.6, we need to establish two claims:

- (1) If there is a certificate that is verified by the above procedure, then the Thurston norm of the class in $H_2(M, \partial M)$ that is dual to $[\phi]$ is m , and M is irreducible.
- (2) If Thurston norm of the class in $H_2(M, \partial M)$ that is dual to $[\phi]$ is m and M is irreducible, then there is a certificate that is verified by the above procedure in time at most a polynomial function of t , $\log(m+2)$ and $\log \|\phi\|_1$.

We will establish these in the next two subsections.

13.4. A PROOF THAT THE CERTIFICATE DOES CERTIFY THURSTON NORM AND IRREDUCIBILITY

We start with (1). So, suppose that there is a certificate as above. The algorithm verifies in (1) of Section 13.2 that T is a union of tori. It also checks in (2) that \mathcal{T} is a triangulation of $M' = M - \text{int}(N(T))$. It also checks in (3) that the Poincaré duals of $i^*([\phi])$ and $[\phi']$ have the same Thurston norm. We will show below that the latter is m , and hence so is the former. By Proposition 11.2, this is equal to the Thurston norm of the Poincaré dual to ϕ , as long as $T \cup \partial M$ is incompressible. So, assuming that the certificate is verified, all that we need to establish is that $T \cup \partial M$ is incompressible, M is irreducible and the Thurston norm of the Poincaré dual to $[\phi']$ is m .

Now, $m = m_1 + m_2$, by (4). By (5), the Poincaré dual of $\phi'|M'_2$ has Thurston norm m_2 . So, we must check that the Poincaré dual of $\phi'|M'_1$ has Thurston norm m_1 . The algorithm checks in (11) and (12) that S_1 is dual to $\phi'_1|M'_1$. The algorithm also checks that $\chi_-(S_1) = m_1$ in (7) and (8). So, we must establish that S_1 has minimal Thurston complexity in its class in $H_2(M', \partial M')$. The algorithm checks in (6) that no component of ∂S_1 bounds a disc in $\partial M'$. Hence, by Lemma 2.1, S_1 has minimal Thurston complexity in its class in $H_2(M', \partial M')$ if and only if it has minimal Thurston complexity in its class in $H_2(M', N(\partial S_1))$. So, we check that the latter holds. In fact, we check that S_1 is taut.

Since no component of ∂S_1 bounds a disc in $\partial M'$ and $\partial M'$ has no u-sutures, the canonical extension $E(S_1)$ is equal to S_1 . We will show that, assuming that the certificate has been correctly verified, there is a hierarchy of the form

$$(M'_1, \emptyset) = E(M_1, \gamma_1) \xrightarrow{E(S_1)} E(M'_2, \gamma'_2) \xrightarrow{A_1} E(M_2, \gamma_2) \xrightarrow{E(S_2)} E(M'_3, \gamma'_3) \xrightarrow{A_2} \dots \xrightarrow{A_n} E(M_{n+1}, \gamma_{n+1}). \quad (*)$$

No component of $E(S_i)$ will be a disc disjoint from the sutures. Also, no boundary curve of $E(S_i)$ bounds a disc in $\partial E(M_i)$ disjoint from the sutures, by Lemma 3.2. Each surface A_i is a union of non-trivial annuli in M_i disjoint from the sutures γ_i . The final manifold $E(M_{n+1}, \gamma_{n+1})$ is a product sutured manifold, with no $S^2 \times I$ components, by (16). Hence, by Theorem 2.3, each manifold in this sequence is taut and each decomposing surface is taut. In particular, $E(S_1) = S_1$ is taut, and T and ∂M are incompressible. Also, as M' is irreducible and T is incompressible, then M is irreducible.

For $i > 1$, the manifolds (M_i, γ_i) are provided by giving their handle structures \mathcal{H}_i in (viii) of the certificate. For $i > 1$, each surface S_i is a boundary-regulated surface provided by (ix) in the certificate. It has a well-defined orientation, because (9) and (10) in the verification complete successfully. The manifold $(M'_{i+1}, \gamma'_{i+1})$ is obtained by decomposing (M_i, γ_i) along S_i . By (14), this decomposition is allowable. Hence, it has a canonical extension as in (*). By (13), no component of $E(S_i)$ is a disc disjoint from the sutures. The annuli A_i are provided by Theorem 10.1, together with the successful verification of (15).

Hence, we have shown that if the certified is verified successfully, then the Thurston norm of the dual of $[\phi]$ is m , and M is irreducible, as required.

13.5. THE EXISTENCE OF THE CERTIFICATE

We now check the other direction (2). We start with the 3-manifold M with the given triangulation \mathcal{T} , the cocycle ϕ and the integer m . We suppose that the Thurston norm of the dual of $[\phi]$ is m , and we have to show that there is a certificate as in Section 13.1.

By Theorem 11.3, the canonical tori T for M may be realised as a normal surface with at most 2^{80t^2} elementary discs. Hence, it has weight at most 2^{2+80t^2} . This forms part (i) in our certificate. Theorem 11.4 and Addendum 11.5 then provide a triangulation for $M' = M - \text{int}(N(T))$ with at most $200t$ tetrahedra. By Theorem 10.5, we may find such a triangulation with at most $1000t$ tetrahedra, and where each boundary torus is triangulated using just two triangles. This is (ii) in the certificate. Addendum 11.5 and Theorem 10.5 provide a 1-coycle ϕ' on \mathcal{T}' with $\|\phi'\|_1 \leq 3^{800t} 1200t \|\phi\|_1$, and such that the duals of $i^*[\phi]$ and $[\phi']$ have the same Thurston norm. This is (iii) in the certificate. By construction, steps (1), (2) and (3) in the verification may therefore be completed.

We let M'_2 be the union of the Seifert fibred components of M' , and let M'_1 be the union of the remaining components. Therefore M'_1 is atoroidal. This decomposition into M'_1 and M'_2 is (iv) of the certificate. We denote the inclusion maps of M'_1 and M'_2 into M by i_1 and i_2 . We denote the Thurston norms of the Poincaré duals of $i_1^*([\phi])$ and $i_2^*([\phi])$ by m_1 and m_2 . These form (v) in the certificate. Then $m_1 + m_2 = m$, by Proposition 11.2. So step (4) verifies successfully.

Note that Theorem 12.3 provides the certificate used in (vi) and this also is verifiable. This is step (5).

So, we now focus on the atoroidal components M'_1 of M' . Let ϕ'_1 be the restriction of ϕ' to M'_1 . By Theorem 8.3, there is a compact oriented lw-taut normal surface S_1 such that $[S_1, \partial S_1]$ is Poincaré dual to $[\phi'_1]$, satisfying $w(S_1) \leq k^{(1000t)^2} \|\phi'_1\|_1$, where k is a universal computable constant. This is at most $c^{t^2} \|\phi\|_1$ for some computable constant c . Because S_1 is lw-taut, whenever two components of S_1 are normally parallel, they are compatibly oriented, as otherwise the union of these two components is null-homologous. The vector (S_1) representing S_1 as a normal surface is (vii) in our certificate. Since S_1 is lw-taut, no component of ∂S_1 bounds a disc in $\partial M'$, and so step (6) in the verification may be completed. As S_1 is lw-taut, $\chi_-(S_1)$ is equal to the Thurston norm of $[S_1]$. So, $\chi_-(S_1) = m_1$, and step (8) in the verification completes successfully. Step (7) trivially completes successfully, because T and ∂M are incompressible, and therefore no component of S_1 can be a disc. Also, no component of S_1 is a sphere, because this would have to be null-homologous, by the irreducibility of M'_1 .

For each collection of normally parallel components of S_1 , one elementary disc from one of these components is chosen. The transverse orientations on these discs form part of the certificate. Note that there are at most $8000t$ of these discs, because this is an upper bound for the number of disjoint normal surfaces that are not normally parallel. The reason for including these transverse orientations as part of the certificate is to ensure that there is no ambiguity about the transverse orientation on S_1 . This collection of transverse orientations is (x) in our certificate, in the case $i = 1$. By the way that these elementary discs are chosen, step (9) in the verification completes successfully. Also, because they arose from the orientation on S_1 , so does step (10) in the case $i = 1$.

As $[S_1, \partial S_1]$ is Poincaré dual to $[\phi'_1]$, steps (11) and (12) also complete successfully.

We now need to establish the existence of the hierarchy as in (*). We will establish the existence of the manifolds (M'_i, γ'_i) and (M_i, γ_i) one at a time. We will also prove inductively that they have the following properties. Our first property is that (M_i, γ_i) and (M'_i, γ'_i) admit some allowable hierarchy. Our second property is that (M_i, γ_i) has a handle structure \mathcal{H}_i of uniform type. Our third property is that $E(M_i, \gamma_i)$ and $E(M'_i, \gamma'_i)$ are atoroidal and no component is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. Our fourth property is that, for each component X of (M'_i, γ'_i) , ∂X is not a single torus with no sutures, unless X is a component of M'_1 . The induction starts with (M'_2, γ'_2) . Note that Lemma 2.2 gives that (M'_2, γ'_2) is taut. It has no u-sutures, and so by Theorem 2.4, it does admit an allowable hierarchy. By Lemma 3.11, $E(M'_2, \gamma'_2)$ is atoroidal and no component is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. For each component Y of M'_1 such that $Y \cap S_1$ is non-empty, then $Y \cap M'_2$ cannot contain a component with boundary a single torus with no sutures. This is because such a component would have arisen from a toral component of S_1 or a union of annular components, which would have been homologically trivial in M'_1 . This contradicts the fact that S_1 is lw-taut.

We now consider the inductive step. So, consider the decomposition

$$(M_i, \gamma_i) \xrightarrow{S_i} (M'_{i+1}, \gamma'_{i+1}).$$

We apply Theorem 10.1 or 10.2 to obtain a decomposition

$$E(M'_{i+1}, \gamma'_{i+1}) \xrightarrow{A_{i+1}} E(M_{i+1}, \gamma_{i+1})$$

where A_{i+1} is a collection of oriented annuli disjoint from γ'_{i+1} and with non-trivial boundary curves. By Theorem 10.1 or 10.2, this decomposition is taut, and in particular, no component of (M_{i+1}, γ_{i+1}) is a solid torus with no sutures. Inductively, we are assuming that $(M'_{i+1}, \gamma'_{i+1})$ admits an allowable hierarchy, and hence, by Theorem 10.1 or 10.2, so does (M_{i+1}, γ_{i+1}) . By Theorem 10.1 or 10.2, $E(M_{i+1}, \gamma_{i+1})$ is atoroidal and no component is a Seifert fibre space other than a solid torus or a copy of $T^2 \times I$. By Lemma 3.10 and the tautness of (M_{i+1}, γ_{i+1}) , no component of (M_{i+1}, γ_{i+1}) has boundary a single torus with no sutures unless it is a component of M'_1 . Theorem 10.1 or 10.2 provides a handle structure \mathcal{H}_{i+1} for (M_{i+1}, γ_{i+1}) of uniform type. Each 0-handle of \mathcal{H}_{i+1} lies within a 0-handle of \mathcal{H}_i . Moreover, for each 0-handle H of \mathcal{H}_i , the 0-handles of \mathcal{H}_{i+1} lying within it have, in total, no greater complexity. Indeed, if the complexity is unchanged, then $\mathcal{H}_{i+1} \cap H$ contains a single 0-handle isotopic to H . Finally, \mathcal{H}_{i+1} has no parallelity handles other than 2-handles. We now apply Theorem 7.8 to (M_{i+1}, γ_{i+1}) , which is possible because (M_{i+1}, γ_{i+1}) admits an allowable hierarchy. Theorem 7.8 gives a fundamental boundary-regulated surface S_{i+1} , which also extends to an allowable hierarchy. Decomposition along S_{i+1} reduces the complexity of the handle structure by Theorem 7.8. The four properties of (M_{i+1}, γ_{i+1}) required by the induction are, as above, easily seen to hold.

By construction, steps (13) and (14) of the verification process complete successfully, because the surfaces S_i are part of an allowable hierarchy. Step (15) also completes successfully, because the handle structure \mathcal{H}_{i+1} is constructed using Theorem 10.1.

Repeat this process until we end with a product sutured manifold (M_{n+1}, γ_{n+1}) , no component of which is pre-spherical. Note that the process is guaranteed to terminate because the complexity of the handle structure \mathcal{H}_{i+1} is strictly less than that of \mathcal{H}_i , for each i . Thus, if we apply step (16) of the verification procedure to the final manifold, this also completes successfully.

In this way, we obtain the certificate, and we have shown that the verification completes successfully. We need to check that the verification can be completed in polynomial time. In order to do this, we will check that the length n for the hierarchy is at most kt , where k is a universal computable constant, and we will check that the algorithms applied to the i th stage of the hierarchy can be completed in time that is at most a polynomial function of t .

We first show that the length of the hierarchy is at most kt . Consider a 0-handle H of \mathcal{H}_2 , and the 0-handles of \mathcal{H}_i that lie inside H . As i increases, the complexity of these handles does not go up.

Moreover, if it stays constant, the handle remains unchanged up to isotopy. By Theorem 6.1, the number of different possible configurations for $\mathcal{H}_i \cap H$ is bounded by some universal computable constant k . Hence, the complexity can decrease at most k times. Now, at each stage of the hierarchy, the complexity must strictly decrease, for some 0-handle H by Theorem 7.8. So, the number of steps in the hierarchy is at most k times the number of 0-handles of \mathcal{H}_2 . But the number of 0-handles of \mathcal{H}_2 is at most $6000t$, since $1000t$ is an upper bound for the number of tetrahedra of \mathcal{T}' , and each tetrahedron of \mathcal{T}' contains at most six 0-handles of \mathcal{H}_2 that are not parallelity handles.

We now show that, at each stage of the hierarchy, the algorithms can be completed in time that is at most a polynomial function of t . The weight of each surface S_i is at most c^{t^2} . The logarithm of this is $t^2 \log(c)$. Whenever the algorithms of Section 9 or Theorem 10.1 are applied, it is this logarithm that appears in the argument. Hence, the running time is at most a polynomial function of t . The final step in the algorithm is the verification of the certificate that (M_{n+1}, γ_{n+1}) is a product sutured manifold with no pre-spherical components. The handle structure on (M_{n+1}, γ_{n+1}) is of uniform type and has at most ct handles. Hence, the time required to complete this verification is also at most a polynomial function of t .

This completes the proof of Theorems 1.5 and 1.6.

14. THE GENUS OF KNOTS IN 3-MANIFOLDS

In this section, we give a proof of Theorem 1.4.

Theorem 1.4. *If KNOT GENUS IN COMPACT ORIENTABLE 3-MANIFOLDS is NP, then NP = co-NP.*

This is a well known consequence of work of Agol, Hass and Thurston [2], but we give a proof here because it does not appear to have been put into print before, and because of its great relevance to the theme of this paper.

Recall that co-NP consists of the class of decision problems for which a negative answer can be certified in polynomial time. Since in practice, there is no reason to believe that certifying a negative answer is as easy or as hard as certifying a positive answer to a problem, then it is widely believed that $\text{NP} \neq \text{co-NP}$. However, if a single problem in co-NP were NP-complete, then NP and co-NP would be equal. It is this observation that is the basis for the proof of Theorem 1.4.

Proof. The main theorem of Agol, Hass and Thurston in [2] is that the problem of deciding whether a simplicial knot in a triangulated 3-manifold has genus at most g is NP-complete. (Somewhat confusingly, they called this problem ‘3-MANIFOLD KNOT GENUS’ but we do not do so here.) Therefore, if it is also in co-NP, then $\text{NP} = \text{co-NP}$. We will show that if KNOT GENUS IN COMPACT ORIENTABLE 3-MANIFOLDS is NP, then there is a method of certifying in polynomial time that the genus of a knot K in a compact orientable 3-manifold M is more than g . So, consider a simplicial knot K in a triangulated 3-manifold M . We may easily check, in polynomial time, whether K is homologically trivial. If it is not, then K does not bound a Seifert surface, and the genus of K is then, by convention, infinite. On the other hand if K is homologically trivial, then it is easy to prove, using normal surface theory, that the genus $g(K)$ of K is at most an exponential function of the number of tetrahedra t in the triangulation of M . Recall that we are given an integer g , and are trying to certify that the genus of K is greater than g . But assuming that KNOT GENUS IN COMPACT ORIENTABLE 3-MANIFOLDS is NP, there is a certificate that confirms the genus of K , which is verifiable in time that is bounded above by a polynomial function of t and the number of digits of $g(K)$ in binary. Since $g(K)$ is at most an exponential function of t , its number of digits is bounded by a linear function of t . Hence, the time required to verify the certificate is at most a polynomial function of t . Hence, we have indeed shown that the problem considered by Agol, Hass and Thurston is in co-NP, as required. \square

In this section, we give a proof of Theorem 1.7.

Theorem 1.7. *KNOTTEDNESS IN IRREDUCIBLE 3-MANIFOLDS is in $NP \cap co-NP$.*

Proof. The fact that this decision problem is in co-NP is well known. This argument is a very minor extension of the one given by Hass, Lagarias and Pippenger [10], who showed that UNKNOT RECOGNITION is in NP. Suppose that we are given a triangulation T , with t tetrahedra, of a compact orientable 3-manifold M with boundary a (possibly empty) union of tori, and that a knot K is given as a specified subcomplex of the 1-skeleton. If K is unknotted, we need a certificate that establishes this and that can be verified in polynomial time. One can build a triangulation T' for the exterior $X = M - \text{int}(N(K))$ by building the second derived subdivision of T , and then removing the interior of any simplex that has non-empty intersection with K . If K is unknotted, then it bounds an embedded disc, and we may arrange that the restriction of this disc to X is properly embedded and has boundary a curve that has winding number one around $N(K)$. One such disc D may be realised as a fundamental normal surface in X (see the proof of Theorem 4.1.13 in [22] for example). By Theorem 8.1, there is then an upper bound on the size of each co-ordinate of $[D]$, which is an exponential function of the number of tetrahedra in T' . Hence, the number of digits of each co-ordinate in binary is bounded above by a linear function of t . The certificate for unknottedness is this normal surface vector $[D]$. The verification of the certificate is as follows. Using the algorithm of Agol, Hass and Thurston, one verifies that D is a disc. One also verifies that its boundary has winding number one along $N(K)$.

We now present a proof that the decision problem is in NP. So, suppose that the knot K in the 3-manifold M is knotted. There are two possible scenarios: either K lies within a 3-ball in M , or it does not. Suppose first that K does lie within a 3-ball. Then $X = M - \text{int}(N(K))$ is reducible, and so there is a reducing sphere S which is a fundamental normal surface in the triangulation T' of X . The vector $[S]$ forms part of our certificate. The first step in the verification is to use the algorithm of Agol, Hass and Thurston to check that S is indeed a 2-sphere. The next step is to use Theorem 9.3 to produce a handle structure \mathcal{H} for the 3-manifold Y that is obtained by cutting X along S . Let Z be the component of Y containing $\partial N(K)$. Then Z is homeomorphic to the exterior of a non-trivial classical knot. Its genus is some positive integer g . This is bounded above by an exponential function of the number of handles of Z , because a minimal genus Seifert surface may be realised as fundamental. Hence, one can use the certificate provided by Theorem 1.5 to verify that the genus of this non-trivial knot is indeed g . Hence, we can certify in this case that K is indeed knotted.

The second scenario is when K does not lie within a 3-ball in M . In this case, $X = M - \text{int}(N(K))$ is irreducible, and the certificate that we use is provided by Theorem 1.6. Thus, in both cases, there is a certificate, verifiable in a polynomial time, which establishes that K is indeed knotted. \square

REFERENCES

1. I. AGOL, *Knot genus is NP*, Conference presentation (2002).
2. I. AGOL, J. HASS, W. THURSTON, *The computational complexity of knot genus and spanning area*, Trans. Amer. Math. Soc. 358 (2006) 3821–3850.
3. D. GABAI, *Foliations and the topology of 3-manifolds*. J. Differential Geom. 18 (1983) 445–503.
4. D. GABAI, *Foliations and genera of links*. Topology 23 (1984), no. 4, 381–394.
5. D. GABAI, *Genera of the arborescent links*. Mem. Amer. Math. Soc. 59 (1986), no. 339, i–viii and 1–98.
6. D. GABAI, *Essential laminations and Kneser normal form*. J. Differential Geom. 53 (1999), no. 3, 517–574.

7. O. GOLDBREICH, *Computational complexity. A conceptual perspective*, Cambridge University Press (2008).
8. W. HAKEN, *Theorie der Normalflächen*. Acta Math. 105 (1961) 245–375.
9. J. HASS, J. LAGARIAS, *The number of Reidemeister moves needed for unknotting*. J. Amer. Math. Soc. 14 (2001), no. 2, 399–428
10. J. HASS, J. LAGARIAS, N. PIPPENGER, *The computational complexity of knot and link problems*. J. ACM 46 (1999), no. 2, 185–211.
11. J. HASS, J. SNOEYINK, W. THURSTON, *The size of spanning disks for polygonal curves*. Discrete Comput. Geom. 29 (2003) 1–17.
12. S. IVANOV, *The computational complexity of basic decision problems in 3-dimensional topology*. Geom. Dedicata 131 (2008), 1–26.
13. W. JACO, U. OERTEL, *An algorithm to decide if a 3-manifold is a Haken manifold*. Topology 23 (1984), no. 2, 195–209.
14. W. JACO, E. SEDGWICK, *Decision problems in the space of Dehn fillings*. Topology 42 (2003) 845–906.
15. W. JACO, P. SHALEN, *Seifert fibered spaces in irreducible, sufficiently-large 3-manifolds*. Bull. Amer. Math. Soc. 82 (1976) 765–767.
16. W. JACO, J. TOLLEFSON, *Algorithms for the complete decomposition of a closed 3-manifold*. Illinois J. Math. 39 (1995), no. 3, 358–406.
17. K. JOHANNSON, *Homotopy equivalences of 3-manifolds with boundaries*. Lecture Notes in Mathematics, 761. Springer, Berlin, 1979.
18. G. KUPERBERG, *Knottedness is in NP, modulo GRH*, Adv. Math. 256 (2014), 493–506.
19. M. LACKENBY, *Dehn surgery on knots in 3-manifolds*, J. Amer. Math. Soc. 10 (1997) 835–864.
20. M. LACKENBY, *Exceptional surgery curves in triangulated 3-manifolds*, Pacific J. Math. 210 (2003), 101–163.
21. M. LACKENBY, *The crossing number of composite knots*, J. Topology 2 (2009) 747–768.
22. S. MATVEEV, *Algorithmic topology and classification of 3-manifold*, Algorithms and Computation in Mathematics, 9. Springer, Berlin, 2007.
23. A. MIJATOVIĆ, *Triangulations of fibre-free Haken 3-manifolds*, Pacific J. Math. 219 (2005) 139–186.
24. M. SCHARLEMANN, *Sutured manifolds and generalized Thurston norms*. J. Differential Geom. 29 (1989), no. 3, 557–614.
25. S. SCHLEIMER, *Sphere recognition lies in NP*, Low-Dimensional and Symplectic Topology, Proceedings of Symposia in Pure Mathematics, 82 (2011), 183–213.
26. W. THURSTON, *A norm for the homology of 3-manifolds*. Mem. Amer. Math. Soc. 59 (1986), no. 339, i–vi and 99–130.
27. J. TOLLEFSON, N. WANG, *Taut normal surfaces*. Topology 35 (1996), no. 1, 55–75.
28. F. WALDHAUSEN, *On irreducible 3-manifolds which are sufficiently large*. Ann. of Math. (2) 87 1968 56–88.

Mathematical Institute, University of Oxford,
 Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, United Kingdom.