

INCOMPRESSIBLE SURFACES, HIERARCHIES AND UNKNOT RECOGNITION

MARC LACKENBY

ABSTRACT. We present a new algorithm to determine whether a compact orientable surface properly embedded in a compact orientable 3-manifold is incompressible. As a special case, this provides a new algorithm to detect the unknot. The central technique is the use of hierarchies. Unlike previous algorithms, no challenging invariants need to be computed and no lengthy search procedures are required. This algorithm leads to a new proof that determining whether a compact orientable 3-manifold has incompressible boundary is in NP.

CONTENTS

1. Introduction	1
2. Hierarchies and boundary patterns	4
3. Essential hierarchies	8
4. Handle structures and normal surfaces	12
5. Decomposing a handle structure along a normal surface	16
6. The algorithm to decide whether a pattern is essential	23
7. An example	30
8. The algorithm to determine incompressibility	31
9. The number of steps	34
10. A refined algorithm using interior parallelity bundles	35
11. Handle structures of uniform type	38
12. Fundamental decomposing surfaces	43
13. Certifying an inessential pattern	45
14. Certifying an essential hierarchy	49
References	54

1. INTRODUCTION

Incompressible orientable surfaces occupy a central position in 3-manifold theory. It is therefore very natural to ask whether a given orientable surface properly embedded in a 3-manifold is incompressible. The following result provides an algorithm to determine this.

Theorem 1.1. *Let $\mathcal{T}$ be a triangulation of a compact orientable irreducible 3-manifold M , possibly with boundary. Let S be a compact orientable normal surface properly embedded in M . Then there is an algorithm to determine whether S is incompressible.*

This theorem is not new, but the algorithm is. The first algorithmic solution to the problem in Theorem 1.1 was given by Haken [3]. His procedure was to cut the manifold M along S and then triangulate the new 3-manifold $M \setminus S$. Then he searched for possible normal compression discs for the two copies of S in $M \setminus S$. In order to make this a finite procedure, he showed that it was possible to restrict to ‘fundamental’ normal discs. A more recent algorithm, due to the author [16], is more complicated, and uses sutured manifold hierarchies. Both these algorithms are useful and have interesting theoretical consequences. But they both have the downside that they require a search through a possibly very extensive search space, for example, the set of all fundamental normal surfaces. The algorithm presented here has the significant advantage that it does not require examination of a large search space. An interesting application of Theorem 1.1 is the following.

Corollary 1.2. *There is an algorithm to determine whether a given knot in the 3-sphere is trivial.*

This follows from Theorem 1.1 because the exterior of a knot in the 3-sphere has incompressible boundary if and only if the knot is non-trivial. Again, this result is not new. There are now many known algorithms to detect the unknot, the first of which was due to Haken [3]. Some require the computation of knot invariants, such as Khovanov homology [10] or Heegaard Floer homology [19, 20]. Others require a possibly extensive search: a search for a spanning disc [5, 16], a homomorphism from the fundamental group of the knot complement to a finite group with non-abelian image [11], sequences of Reidemeister moves [15] or sequences of grid diagrams [2]. As a result, none of these algorithms are likely to run in polynomial time. Indeed, in the case of Khovanov homology, there is an exponential lower bound on its worst-case running time [7], subject to some standard conjectures in complexity theory. The algorithm that we give seems amenable to possible speed-up, which we will explore in future work. Furthermore, in a companion paper [17], we will show how this algorithm can be extended to determine whether a given link in the 3-sphere is hyperbolic.

The algorithm generalises to 3-manifolds with a boundary pattern. Recall that a *boundary pattern* for a compact orientable 3-manifold M is a collection P of disjoint simple closed curves and trivalent graphs embedded in ∂M . A boundary pattern P is *essential* if, for every disc D properly embedded in M that intersects P transversely in at most three points away from the vertices of P , there is a disc D' in ∂M that intersects P in the empty set, a single arc, or a graph with a single 3-valent vertex in the interior of D' and three 1-valent vertices on $\partial D'$. (See Figure 1.) If there is a disc D properly embedded in M intersecting P transversely at most three times away from the vertices, but there is no corresponding disc D' , then D is termed a *violating disc*. In the final case, when D' intersects P in a graph with a single 3-valent vertex in the interior of D' and three 1-valent vertices on $\partial D'$, this graph is known as a *tripod*. Boundary patterns will be the focus of this paper. At this stage, we observe that when P is the empty pattern, then P is essential if and only if ∂M is incompressible.

Theorem 1.3. *Let (M, P) be a compact orientable irreducible 3-manifold with boundary pattern. Let $\mathcal{T}$ be a triangulation of M in which P is simplicial. Then there is an algorithm to determine whether P is essential. Moreover, if P is not essential, the algorithm provides a violating disc.*




FIGURE 1. Essential boundary pattern

When P is essential and M irreducible, the algorithm outputs an essential hierarchy for (M, P) . Such hierarchies, which are defined in Section 2, will play a central role in this paper. They are a finite sequence of decompositions along properly embedded orientable surfaces where the final manifold is a collection of 3-balls with an essential boundary pattern.

In the later parts of the paper, we will establish that the algorithms given above can be used to show that the decision problems are in NP. Moreover, we will drop the hypothesis that the 3-manifold is irreducible. Specifically, we will consider the following problems.

The decision problem INCOMPRESSIBLE BOUNDARY takes, as its input a triangulation of a compact orientable 3-manifold M and it asks whether ∂M is incompressible.

Theorem 1.4. *The decision problem INCOMPRESSIBLE BOUNDARY is in NP.*

This theorem was proved by the author in [16], using sutured manifold hierarchies, based on argument of Agol. The problem was also shown to be in co-NP by Ivanov [6].

This has applications for the famous problem KNOTTEDNESS. This takes, as its input, either a diagram of a knot K or a triangulation of its exterior, and it asks whether K is a non-trivial knot.

Corollary 1.5. *The decision problem KNOTTEDNESS is in NP.*

This was one of the main theorems in [16]. Previously, Kuperberg [11] showed that KNOTTEDNESS is in NP, conditionally on the Generalised Riemann Hypothesis.

The decision problem ESSENTIAL BOUNDARY PATTERN takes, as its input, a triangulation of compact orientable 3-manifold M with boundary pattern P that is simplicial, and it asks whether P is essential. The following is a new result.

Theorem 1.6. *The decision problem ESSENTIAL BOUNDARY PATTERN is in NP and co-NP.*

When P is inessential, the certificate is a violating disc in normal form. When P is essential and M is irreducible, the certificate is an essential hierarchy for (M, P) . The fact that essential hierarchies can be encoded efficiently and can be certified as essential will be extremely useful in future work [17].

Much of the theory of hierarchies as explained in Section 2 is due to Johannson [8]. The presentation here follows that of a graduate course [12] given by the author in 1999. One of the goals of that course was to provide a new proof of the loop theorem using hierarchies. The fact that the theory also has algorithmic applications was observed in a conversation with Ian Agol in 2013. The author developed and expanded the theory over the intervening years, resulting in this paper and its sequel [17].

The structure of the paper is as follows. In Section 2, we recall some basic material on hierarchies and boundary patterns. In Section 3, we explore essential hierarchies and prove one direction of Theorem 2.5, which asserts that when a 3-manifold admits an essential hierarchy, then it is irreducible and has essential boundary pattern. In Section 4, we introduce normal surfaces, in the context of handle structures with a boundary pattern. In Section 5, we examine how handle structures behave when decomposed along a normal surface. The goal here is to establish an upper bound on the number of decompositions in a hierarchy. This gives the other direction of Theorem 2.5, which gives a sufficient condition for a 3-manifold with boundary pattern to admit an essential hierarchy. In Section 6, we present the algorithm required by Theorem 1.3, and in Section 7, an implementation of the algorithm in a specific example is given. The algorithm to determine whether a surface is incompressible, as required by Theorem 1.1, is given in Section 8. A crude analysis of the number of steps in the algorithm is given in Section 9. The crucial quantity that controls an upper bound on the running time is the number of decompositions in the hierarchy. Therefore, in Section 10, we give a more complicated algorithm for constructing a hierarchy, that has the advantage that the number of decompositions is explicitly controlled. In Section 11, we show that the handle structures created in the hierarchy have uniform type, which means that each handle meets the pattern and its neighbouring handles in one of finitely many types, independent of the manifold. This is important for the later parts of the paper, in Sections 12, 13 and 14, where we show that the decision problem in Theorem 1.6, deciding whether a pattern is essential, is in NP and co-NP.

Remark 1.7. The author was partially supported by EPSRC grant EP/Y004256/1. For the purpose of open access, the author has applied a CC BY public copyright licence to any author accepted manuscript arising from this submission.

2. HIERARCHIES AND BOUNDARY PATTERNS

We will use the following terminology, which has become standard.

Definition 2.1. Let M be a manifold and let S be either a codimension one submanifold properly embedded in M or a compact codimension zero submanifold. Then the manifold $M \setminus S$ *obtained by cutting M along S* is defined as follows. Pick a Riemannian metric on M . Then $M - S$ inherits a Riemannian metric and $M \setminus S$ is defined to be metric completion of $M - S$. We also permit S not to be a subset of M as long as S and M are both subsets of some larger space and $S \cap M$ satisfies the above conditions, in which case $M \setminus S$ is defined to be $M \setminus (S \cap M)$. When M is compact and S is a codimension one submanifold, then the inclusion $M - S \rightarrow M$ induces a quotient map $M \setminus S \rightarrow M$.

Our algorithm makes use of hierarchies. Here, one starts with a compact orientable 3-manifold, and then one cuts along a properly embedded orientable surface,

forming a new 3-manifold. In this manifold, one finds another properly embedded orientable surface, and one cuts along this. The process continues until the final manifold is a collection of 3-balls. The manifolds in the hierarchy will inherit a boundary pattern, as follows.

Definition 2.2. Let (M, P) be a compact orientable 3-manifold with boundary pattern. Then a compact orientable surface S properly embedded in M is *transverse* to P if it is disjoint from the vertices and intersects the edges transversely. We then say that S is a *decomposing surface*. The manifold $M \setminus S$ inherits a boundary pattern P' , where P' is the union of $P \cap (M \setminus S)$ and the copies of ∂S in $\partial(M \setminus S)$. We write

$$(M, P) \xrightarrow{S} (M \setminus S, P')$$

for this decomposition. We also denote $(M \setminus S, P')$ by $(M, P) \setminus S$. (See Figure 2.)




FIGURE 2. Decomposing (M, P) along a surface S

Definition 2.3. Let (M, P) be a compact orientable 3-manifold with boundary pattern. Then a *partial hierarchy* for (M, P) is a sequence of decompositions

$$(M, P) = (M_1, P_1) \xrightarrow{S_1} (M_2, P_2) \xrightarrow{S_2} \dots \xrightarrow{S_\ell} (M_{\ell+1}, P_{\ell+1}).$$

It is a *hierarchy* if $M_{\ell+1}$ is a disjoint union of 3-balls.

This is a modification of the standard terminology, since we do not require the surfaces to be incompressible (although they will end up being so in the cases of interest to us).

Recall, from Section 1, the definition of an essential boundary pattern. At first sight, this is a strange definition, because it is not at all clear why one should be interested only in simple closed curves that intersect the pattern in at most three points. We will give more motivation for it in the results that follow.

Definition 2.4. An *essential hierarchy* for a compact orientable 3-manifold M with boundary pattern P is a sequence of decompositions

$$(M, P) = (M_1, P_1) \xrightarrow{S_1} (M_2, P_2) \xrightarrow{S_2} \dots \xrightarrow{S_\ell} (M_{\ell+1}, P_{\ell+1}).$$

where $(M_{\ell+1}, P_{\ell+1})$ is a collection of balls with essential boundary pattern.

Again, it is not immediately clear why one should only be concerned with whether the final pattern is essential. But this is explained by the following result.

Theorem 2.5 (Essential hierarchy iff essential pattern and irreducible). *Let M be a compact orientable 3-manifold with a boundary pattern P . Then (M, P) admits an essential hierarchy if and only if the following all hold:*

- (1) *either M contains a properly embedded orientable incompressible surface with no 2-sphere components, or ∂M is non-empty;*
- (2) *P is essential;*
- (3) *M is irreducible.*

It is relatively straightforward to determine quickly whether a boundary pattern for a 3-ball is essential. This can be rephrased in standard graph-theoretic terms as follows. Recall that a graph is *simple* if it has no edge loops and does not have multiple edges between pairs of vertices. A simple graph is *n -connected* for some positive integer n if it has at least $n + 1$ vertices and there is no set of $n - 1$ vertices whose removal disconnects the graph.

Proposition 2.6 (Essential pattern for a ball). *Let M be a 3-ball, and let P be a boundary pattern for M . Then P is essential if and only if P is empty, a simple closed curve or a connected graph, and in the latter case, the planar graph dual to P is either a simple 4-connected graph or the complete graph on 3 or 4 vertices.*

Proof. We first show that the given conditions on P imply that P is essential. If P is empty or a simple closed curve, then clearly P is essential. So suppose that P is a connected graph and that the graph G dual to P is simple and 4-connected or the complete graph on 4 vertices. Consider a simple closed curve C in ∂M that intersects P transversely in at most 3 points. If C is disjoint from P , then it must bound a disc in ∂M disjoint from P , since P is connected and ∂M is a sphere. If C intersects P at one point, this specifies an edge loop in G , contrary to our assumption that G is simple. If C intersects P at two points, then these two points must lie on the same edge of P , as otherwise we obtain two edges between the same pair of vertices of G . Hence, C bounds a disc in ∂M intersecting P in an arc. If C intersects P at three points, then this specifies a closed walk of length 3 in G . This is a cycle because otherwise G is not simple. Removing the vertices in this cycle does not separate G , because G is 4-connected or the complete graph on 3 or 4 vertices. Hence, C must bound a disc intersecting no regions of $\partial M - P$ other than the three that C visits. This implies that P intersects this disc in a tripod. Thus we have shown that P is essential.

Conversely, suppose that P is essential. We may suppose that P is non-empty and not a simple closed curve, as the implication is immediate in those cases. Then P is clearly connected as otherwise there is a simple closed curve in $\partial M - P$ that does not bound a disc in ∂M disjoint from P . Let G be its dual planar graph. An edge loop in G would give a simple closed curve intersecting the pattern once, contradicting the assumption that P is essential. More than one edge between a pair of vertices would give a simple closed curve in ∂M intersecting P in exactly two points that lie in different edges of P . Again this contradicts the assumption that P is essential. Consider a collection of at most three vertices in G . We will show that these do not separate G . Each vertex corresponds to a region of $\partial M - P$, and for these vertices to disconnect G , the union of these regions must disconnect ∂M . If they do then there must be a simple closed curve lying in the union of these regions that separates at least two other regions of $\partial M - P$. Since P is essential, this curve bounds a disc in ∂M intersecting P in an arc or a tripod. This implies that these three vertices did not in fact separate G . Note that removing these three vertices of G may create the empty graph, but in this case, G is the complete graph

on 3 vertices. Hence, we have shown that G is 4-connected or the complete graph on 3 or 4 vertices. $\square$

Note that it can be determined in polynomial time (as a function of the number of vertices and edges) whether a graph is simple and 4-connected or the complete graph on 4 vertices.

Theorem 2.5 is an extremely powerful tool for showing that a pattern is essential, as the following example demonstrates.

Example 2.7. The knot 5_2 is shown in Figure 3. Let M be its exterior and let P be the empty pattern. A hierarchy for (M, P) is given. The final manifold consists of two 3-balls, and the boundary pattern of one of them is shown. It is easy to check that this pattern is essential, and hence by Theorem 2.5, (M, P) is essential and so the knot 5_2 is not the unknot.




FIGURE 3. An essential hierarchy

Theorem 2.5 is not new. It is implicit in work of Johannson [8] and a version of it was used by author in lecture notes in 1999 [12]. However, because it is so

central to our work and because we will need details of the argument later in the paper, we give a full proof.

3. ESSENTIAL HIERARCHIES

Our aim is to prove Theorem 2.5. We first show that if (M, P) admits an essential hierarchy, then P is essential and M is irreducible. We will also establish that the decomposing surfaces have the following important property.

Definition 3.1. Let M be a compact orientable 3-manifold and let P be a boundary pattern. Let S be a compact surface properly embedded in M . Then a *pattern-compression disc* for S is a disc D embedded in M such that

- (1) $D \cap S$ is an arc α in ∂D ;
- (2) $\partial D \setminus \alpha = D \cap \partial M$ intersects P at most once; and
- (3) α does not separate off a disc from S intersecting P at most once.

If no such pattern-compression disc exists, then S is *pattern-incompressible*. (See Figure 4.) If D is a pattern-compression disc for S , then the surface obtained from S by cutting along $S \cap D$ and then attaching two parallel copies of D is the surface obtained by *pattern-compressing* along D .




FIGURE 4. A pattern-incompressible surface. (For the sake of clarity, we have assumed here that M is irreducible and P is essential.)

A key step in the proof of Theorem 2.5 is the following.

Lemma 3.2 (Essential pattern pulls back). *Let*

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

be a decomposition and suppose that P' is essential. Then P is essential and S is incompressible and pattern-incompressible.

Proof. We first show that P is essential. Let D be a disc properly embedded in M that is transverse to P and that intersects P at most three times. We wish to show that D is not a violating disc.

Perform a small isotopy to S , leaving P invariant, so that afterwards $S \cap D$ is a collection of simple closed curves and arcs properly embedded in D .

If there is a simple closed curve of $S \cap D$, there is one that is innermost in D . It bounds a disc D' in M' that is disjoint from P' . Since we are assuming that P' is essential, $\partial D'$ bounds a disc D'' in $\partial M'$ that is disjoint from the pattern. This disc is a subset of S . We can remove D' from D and replace it by a parallel copy of D'' . The result is violating for P if and only if D was violating. However, its intersection with S contains fewer simple closed curves.

So, we may assume that $D \cap S$ is a collection of arcs properly embedded in D . Say there are $n > 0$ of them. Then at the endpoints of these arcs, we see $4n$ points of intersection with P' . We also have at most 3 points of intersection between ∂D and P' . So, in the $n + 1$ discs $D \setminus (D \cap S)$, there are at most $4n + 3$ points of intersection with P' . Hence, one these discs, D_1 say, has at most three points of intersection with P' in its boundary. Necessarily, it has at least two points of intersection with P' . We are assuming that P' is essential. Hence, ∂D_1 bounds a disc D_2 in $\partial M'$. This disc D_2 intersects P' either in a single arc or a tripod. Therefore, $D_2 \cap S$ is a disc. We may remove D_1 from D and replace it by $D_2 \cap S$, and then perform a further small isotopy to give a disc D_3 properly embedded in M that is transverse to P and that has fewer components of intersection with S . This is a violating disc for M if and only if D was.

So we may assume that D is disjoint from S and hence lies in M' . Since P' is essential, ∂D bounds a disc D' in $\partial M'$ intersecting P' in the empty set, an arc or a tripod. Since $\partial D'$ is disjoint from S , all of D' must be disjoint from S . So, D' lies in ∂M . So, D was not a violating disc, and so we have proved that P is essential.

We now show that S is incompressible. Consider a disc D embedded in the interior of M such that $D \cap S = \partial D$. Then D is properly embedded in M' and is disjoint from the pattern P' . Since P' is essential, ∂D bounds a disc in $\partial M'$ disjoint from P' . This is a disc in S . So D is not a compression disc for S .

We now show that S is pattern-incompressible. Let D be a disc embedded in M such that $D \cap S$ is an arc in ∂D , and the remainder of ∂D is an arc in ∂M intersecting P at most once. Then D lies in M' and intersects P' at least twice and at most three times. Since P' is essential, ∂D bounds a disc D' in $\partial M'$ intersecting P' in an arc or a tripod. The intersection between D' and S is a disc separated off by $D \cap S$ that intersects P at most once. So D is not a pattern-compression disc. $\square$

Lemma 3.3 (Irreducible and essential pattern pulls back). *Let*

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

be a decomposition and suppose that M' is irreducible and P' is essential. Then M is irreducible and P is essential.

Proof. The fact that P is essential is a result of Lemma 3.2. We need to show that M is irreducible. Consider a 2-sphere S^2 properly embedded in M . This intersects S in a collection of simple closed curves. An innermost one bounds a disc D' in M' that is disjoint from P' . Since P' is essential, $\partial D'$ bounds a disc D'' in S and because M' is irreducible, we may isotope D'' across to D' and then a little further, to remove this curve of $S^2 \cap S$. So we may assume that S is disjoint from S^2 , and hence S^2 lies in M' . As M' is irreducible, S^2 bounds a ball in M' and this is also a ball in M . $\square$

In many circumstances, it is also the case that when P is essential and M is irreducible, then P' is also essential and M' is also irreducible. However, in order to guarantee this, we require that the decomposing surface S is incompressible and pattern-incompressible. We now investigate how to arrange this.

When we pattern-compress a surface, we would like to say that the new surface is, in some sense, simpler than the original one. We formalise this as follows, using the following measure of complexity.

Definition 3.4. The *pattern-complexity* $\chi_p(S)$ of a properly embedded surface S is defined to be $-4\chi(S) + |S \cap P|$.

A useful feature of pattern-complexity is the following immediate result.

Lemma 3.5. *If S is a compact connected orientable surface satisfying $\chi_p(S) < 0$, then S is a sphere or disc intersecting the pattern in at most three points.*

Lemma 3.6 (Pattern-compression simplifies). *Let S be a properly embedded connected orientable surface that is non-separating. Let S' be obtained from S by pattern-compressing or compressing. Then at least one component of S' is also non-separating and has smaller pattern-complexity than S .*

Proof. The lemma is straightforward and well known in the case of a compression. So suppose that we pattern-compress S . We decrease $-4\chi(S)$ by 4 and increase $|S \cap P|$ by at most 2. So χ_p certainly goes down by at least 2. Note that at least one component of S' is non-separating. Thus, the lemma is proved in the case where S' is connected. However, we need to consider the possibility that we have created a separating component S'_1 with negative pattern-complexity. Suppose that this occurs. Let S'_2 be the other component of S' which is necessarily non-separating. Note that S'_1 is a disc and it must intersect P at least twice by the definition of a pattern-compression. If S'_1 has three points of intersection with P , then $|S'_2 \cap P| < |S \cap P|$ and hence $\chi_p(S'_2) < \chi_p(S)$. If S'_1 has two points of intersection with P , then both of these must lie in S by the definition of a pattern-compression. Hence, again $\chi_p(S'_2) < \chi_p(S)$. $\square$

Corollary 3.7 (Simplest surface pattern-incompressible). *Let (M, P) be a compact orientable 3-manifold with a boundary pattern. Let S be a properly embedded connected orientable surface that is non-separating and that has smallest pattern-complexity among all such surfaces. Then S is pattern-incompressible and incompressible.*

Lemma 3.8 (Essential pushes forward). *Let M be a compact orientable 3-manifold with essential boundary pattern P . Let S be a connected pattern-incompressible incompressible surface properly embedded in M , which is not a disc intersecting P in at most 3 points. Then $M \setminus S$ inherits an essential boundary pattern P' .*

Proof. Let D be a disc properly embedded in $M \setminus S$ with $\partial D \cap P'$ at most three points. The curve ∂D may run through parts of $\partial(M \setminus S)$ coming from ∂M and parts coming from S . Note however the points where it swaps must be points of $\partial D \cap P'$, and that at most one side of any point of $\partial D \cap P'$ lies in S . Hence, at most one arc or simple closed curve of $\partial D \setminus P'$ lies in S .

Case 1. ∂D is disjoint from S .

Then $\partial D \subset \partial M$. Since P is essential, ∂D bounds a disc D' in ∂M intersecting P in the empty set, an arc or a tripod. If S intersects D' , then pick a simple closed

curve of $S \cap D'$ innermost in D' . The disc this bounds cannot be a compression disc for S . Hence, S must be a disc. By assumption, S intersects P in at least 4 points. Hence, it intersects some edge of $P \cap D'$ more than once. We can therefore find a pattern-compression disc for S , which is contrary to assumption. Hence, D' is disjoint from S , and therefore lies in $\partial(M \setminus S)$. This verifies that D is not a violating disc for P' .

Case 2. ∂D lies in S .

By the incompressibility of S , ∂D bounds a disc D' in S . This is a disc in $\partial(M \setminus S)$ disjoint from P' . So D is not violating.

Case 3. ∂D intersects S in an arc.

Then $\partial D \setminus S$ intersects P at most once. Since D is not a pattern-compression disc for S , $D \cap S$ separates off a disc D_1 of S intersecting P in at most one point. Then, $D \cup D_1$ is a disc properly embedded in M , intersecting P in at most two points. Since P is essential, there is therefore a disc D_2 in ∂M with $\partial D_2 = \partial(D \cup D_1)$ and with $D_2 \cap P$ either empty or an arc. Therefore, $D_1 \cup D_2$ is a disc in $\partial(M \setminus S)$ intersecting P' in an arc or a tripod.

This gives that P' is essential. $\square$

We also have following well known result.

Lemma 3.9 (Irreducible pushes forward). *Let M be a compact orientable irreducible 3-manifold. Let S be a properly embedded incompressible surface with no 2-sphere components. Then $M \setminus S$ is irreducible.*

We will occasionally want to cut along surfaces in the expectation that they are pattern-incompressible. An example of this is as follows.

Lemma 3.10 (Decomposition along square). *Let (M, P) be a compact orientable 3-manifold with an essential boundary pattern. Let S be a properly embedded disc that intersects P four times. Then S is pattern-incompressible unless ∂S bounds a disc D in ∂M that has one of the following forms:*

- (1) $D \cap P$ consists of two arcs; or
- (2) $D \cap P$ is a graph with two vertices joined by an edge, along with four edges running to the boundary of D .

Proof. Suppose that, on the contrary, S has a pattern-compression disc D' . Then $S \setminus D'$ has two components S_1 and S_2 , each of which is a disc intersecting P twice. For $i = 1$ and 2 , let D'_i be $S_i \cup D'$ isotoped a little so that it is disjoint from S . Then D'_i is properly embedded in M and intersects P either two or three times. If each of these discs intersects P twice, then using the fact that P is essential, we deduce that ∂S bounds a disc D in ∂M as in (1) above. If both D'_1 and D'_2 intersect P three times, then we deduce that ∂S bounds a disc D in ∂M as in (1) or (2) above. $\square$

Definition 3.11. We say that a disc S properly embedded in M intersecting the pattern four times is a *square*. It is *trivial* if ∂S bounds a disc D in ∂M satisfying (1) or (2) of Lemma 3.10. (See Figure 5.) Otherwise it is *non-trivial*.

We rephrase and extend Lemma 3.10 as follows.




FIGURE 5. Trivial squares

Lemma 3.12. *Let (M, P) be a compact orientable 3-manifold with a boundary pattern. Let S be a non-trivial square. Then a pattern-compression disc for S yields a violating disc for (M, P) . This violating disc is one of the components of the surface obtained by pattern-compressing S .*

We now have many of the elements required for the proof of Theorem 2.5. Let (M, P) be a compact orientable irreducible 3-manifold with an essential boundary pattern. Our aim is to build an essential hierarchy for (M, P) . If M is closed, then by assumption, it contains a properly embedded incompressible orientable surface with no 2-sphere components. Decomposing along this surface gives an irreducible 3-manifold with non-empty boundary and essential (empty) boundary pattern. So, we may assume that ∂M is non-empty. If ∂M is a union of spheres, then these bound 3-balls by the irreducibility of M , and so in this case there is nothing to prove. Hence, we may assume that $b_1(\partial M) > 0$. Since $b_1(M) \geq b_1(\partial M)/2$, we deduce that $b_1(M) > 0$. Therefore, M contains a properly embedded connected orientable non-separating surface S . By picking S to have smallest pattern-complexity, we may assume that S is incompressible and pattern-incompressible by Corollary 3.7. When we decompose along such a surface, Lemmas 3.8 and 3.9 imply that the resulting manifold is irreducible and has essential boundary pattern. We may therefore repeat with this manifold. However, we must show that this process can be arranged to terminate. In [12], a proof was given using Kneser's theorem bounding the number of disjoint non-parallel properly embedded incompressible boundary-incompressible surfaces in a compact 3-manifold. In the following section, we give an alternative proof. Although longer than the proof in [12], it will be more useful because it will lead to a good bound on the length of the resulting hierarchy.

4. HANDLE STRUCTURES AND NORMAL SURFACES

Our algorithm deals with partial hierarchies. The way that the surfaces in these partial hierarchies are encoded is as normal surfaces. Normal surface theory takes place in a 3-manifold equipped either with a triangulation or with a handle structure. We mostly use handle structures, as they behave better than triangulations when the 3-manifold is decomposed along a normal surface. Normal surfaces in handle structures have been used in the theory of hierarchies since the work of Haken [4] and Waldhausen [22]. One of their particular uses is as a method for controlling the length of hierarchies, in other words the number of decompositions until the manifold has been cut up into balls.

Definition 4.1. In this paper, a *handle structure* on a 3-manifold or a 2-manifold will satisfy the following conditions:

- (1) each i -handle $D^i \times D^{k-i}$ ($k = 2$ or 3) intersects the handles of lower index in $\partial D^i \times D^{k-i}$;

- (2) any two i -handles are disjoint;
- (3) in the case of a 3-manifold, the intersection of any 1-handle $D^1 \times D^2$ with any 2-handle $D^2 \times D^1$ is of the form $D^1 \times \alpha$ in $D^1 \times D^2$ where α is a collection of arcs in ∂D^2 , and of the form $\beta \times D^1$ in $D^2 \times D^1$ where β is a collection of arcs in ∂D^2 ;
- (4) any 2-handle runs over at least one 1-handle.

Definition 4.2. Let $\mathcal{H}$ be a handle structure for a 3-manifold and for $i \in \{0, 1, 2, 3\}$, let $\mathcal{H}^i$ denote the union of the i -handles. The *characteristic surface* for $\mathcal{H}$ is $\mathcal{H}^0 \cap (\mathcal{H}^1 \cup \mathcal{H}^2)$. We denote it by $\mathcal{F}(\mathcal{H})$, or simply by $\mathcal{F}$ when the context is clear. It has a handle structure consisting of just 0-handles and 1-handles, where the 0-handles are $\mathcal{F}^0(\mathcal{H}) = \mathcal{H}^0 \cap \mathcal{H}^1$ and the 1-handles are $\mathcal{F}^1(\mathcal{H}) = \mathcal{H}^0 \cap \mathcal{H}^2$. We will view the characteristic surface as a subsurface of $\partial \mathcal{H}^0$.

We denote the number of 0-handles in the handle structure $\mathcal{H}$ by $|\mathcal{H}|$.

Definition 4.3. A *handle structure* for a 3-manifold M with boundary pattern P is a handle structure for M satisfying Definition 4.1 such that P is disjoint from the 2-handles and, for each 1-handle $H_1 = D^1 \times D^2$, the intersection $H_1 \cap P$ is $D^1 \times X$ where X is a finite set of points in ∂D^2 .

We will be decomposing handle structures along normal surfaces. But first we deal with the following weaker notion.

Definition 4.4. Let $\mathcal{H}$ be a handle structure of a 3-manifold M with boundary pattern P . We say that a surface S properly embedded in M is *standard* with respect to $\mathcal{H}$ if the following hold:

- (1) it intersects each 0-handle in a collection of disjoint properly embedded discs;
- (2) it intersects each 1-handle $D^1 \times D^2$ in a surface of the form $D^1 \times C$ for some arcs C properly embedded in D^2 ;
- (3) it intersects each 2-handle $D^2 \times D^1$ in a surface of the form $D^2 \times X$ for some finite collection of points X in the interior of D^1 ;
- (4) it is disjoint from the 3-handles;
- (5) it intersects P transversely.

We have analogous definitions for 1-manifolds in surfaces.

Definition 4.5. Let $\mathcal{H}$ be a handle structure of a surface F . We say that a 1-manifold C properly embedded in F is *standard* with respect to $\mathcal{H}$ if the following hold:

- (1) it intersects each 0-handle in a union of disjoint properly embedded arcs;
- (2) it intersects each 1-handle $D^1 \times D^1$ in a union of arcs parallel to the core of the handle, that is, arcs of the form $D^1 \times X$ for some finite collection of points X in the interior of D^1 ;
- (3) it is disjoint from the 2-handles.

Definition 4.6. Let $\mathcal{H}$ be a handle structure for a 3-manifold M with boundary pattern P . A curve C in $\partial \mathcal{H}^0$ is *normal* if it intersects P and $\partial \mathcal{F}$ transversely, its intersection with $\mathcal{F}$ is standard and the following all hold:

- (1) C does not lie entirely in $\partial M - P$;
- (2) C does not bound a disc in $\partial \mathcal{H}^0 - \mathcal{F}$ that intersects P in a single arc or a tripod;

- (3) C intersects each 1-handle of $\mathcal{F}$ in at most one arc;
- (4) C intersects each component of $\partial\mathcal{H}^0 - (\mathcal{F} \cup P)$ in at most one arc;
- (5) C does not intersect components of $\mathcal{F}^1$ and $\partial\mathcal{H}^0 - (\mathcal{F} \cup P)$ that are adjacent;
- (6) if C intersects two components of $\partial\mathcal{H}^0 - (\mathcal{F} \cup P)$ that are incident along an arc of $P \cap \mathcal{H}^0$, then C also intersects this arc;
- (7) there is no disc in $\partial\mathcal{H}^0$ that has boundary equal to the union of a sub-arc of C and a sub-arc of $\partial\mathcal{F}^0$ is otherwise disjoint from C and $\mathcal{F}$, and that intersects P in at most two arcs, each of which runs from $\partial\mathcal{F}^0$ to C .

See Figure 6.




FIGURE 6. Configurations of the curve C in $\partial\mathcal{H}^0$ that violate normality. In cases (3)-(7), arcs α , β or γ are shown that are part of the boundary of a simplifying disc.

Definition 4.7. Let $\mathcal{H}$ be a handle structure for a 3-manifold M with boundary pattern P . A surface S properly embedded in M is *normal* with respect to $\mathcal{H}$ if it is standard and each component of $S \cap \partial\mathcal{H}^0$ is normal.

In the case where the pattern P is empty, Definition 4.6 becomes the usual definition of a normal surface in a handle structure, for example as in [18, Definition 3.4.1]. Most of the conditions in Definition 4.6 are natural generalisations to the case of general pattern. The possible exception is (7), which seems somewhat ad hoc. However, it will be used to establish that manifolds appearing in a hierarchy embed in the original manifold in a ‘simple’ way. (See Theorem 11.3.)

There is a way of understanding the definition of normality in terms of simplifying discs, as follows.

Definition 4.8. Let $\mathcal{H}$ be a handle structure for a 3-manifold M with boundary pattern P . Let S be a standard surface properly embedded in M . Then a *simplifying disc* for S is a disc D lying in a 0-handle H_0 of $\mathcal{H}$ satisfying one of the following:

- (1) ∂D is the concatenation of two arcs $\alpha \subset S$ and $\beta \subset \partial H_0$, where $\alpha = D \cap S$, and $\beta = D \cap \partial H_0$ is an arc on the boundary of a 2-handle of $\mathcal{H}$ that is vertical in its product structure (see Figure 7 and Figure 6 (3));
- (2) ∂D is disjoint from P and is the concatenation of three arcs $\alpha \subset S$, $\beta \subset \partial H_0 \cap \partial M$ and $\gamma \subset \partial H_0 \cap \mathcal{H}^2$, where $\alpha = D \cap S$, $\beta = D \cap \partial M$ and γ is an arc on the boundary of a 2-handle of $\mathcal{H}$ that is vertical in its product structure (see Figure 6 (5));
- (3) ∂D intersects P at most once and is the concatenation of two arcs $\alpha \subset S$ and $\beta \subset \partial H_0$, where $\alpha = D \cap S$, and $\beta = D \cap \partial H_0$ is an arc in ∂M that is not parallel to an arc of $S \cap \partial H_0 \cap \partial M$ intersecting P at most once (Figure 6 (4) and (6));
- (4) D is a subset of $\partial \mathcal{H}^0$ and ∂D is the concatenation of two arcs $\alpha = D \cap S$ and $\beta = D \cap \mathcal{F}^0$, and with the properties that D is otherwise disjoint from $\mathcal{F}$ and that D intersects P in at most two arcs, each of which runs from $\partial \mathcal{F}^0$ to C (see Figure 6 (7)).




FIGURE 7. A simplifying disc D

The following is immediate from the definition of normality.

Lemma 4.9. *Let $\mathcal{H}$ be a handle structure for a 3-manifold M with boundary pattern P . Then a surface S properly embedded in M is normal if and only if the following all hold:*

- (1) S is standard;
- (2) S has no simplifying disc;

- (3) no component of S is a disc lying in a 0-handle that is parallel to a disc in $\partial M \cap \partial \mathcal{H}^0$ intersecting P in the empty set, an arc or tripod.

We will want to place properly embedded surfaces into normal form. In general, we may need to make modifications to the surface in order to do this. These modifications will not increase the following quantities.

Definition 4.10. Let $\mathcal{H}$ be a handle structure of a compact orientable 3-manifold M , and let S be a standard surface properly embedded in M . The *weight* of S is the number of components of $S \cap \mathcal{H}^2$. The *extended weight* of S is $|S \cap \mathcal{H}^2| + |S \cap \mathcal{H}^1|$.

Lemma 4.11 (Normalising a surface). *Let M be a compact orientable 3-manifold with boundary pattern P . Let S be a compact orientable surface properly embedded in M transverse to P that is non-trivial in $H_2(M, \partial M)$. Then, after applying a sequence of isotopies preserving P , compressions and pattern-compressions, some component of the resulting surface is normal and either non-separating or a violating disc. Moreover, if S is standard in $\mathcal{H}$, then the weight and extended weight of the resulting normal surface are at most the weight and extended weight of S .*

Proof. We may readily ensure that S is standard. In particular, if S intersects some 0-handle H_0 in anything other than discs, then we may modify it, keeping $S \cap \partial H_0$ unchanged but ensuring that each of these curves bounds a disc in H_0 . This process is achieved by a sequence of isotopies and compressions.

We may further assume that S has no component that is a boundary-parallel disc, as we could remove this component and the surface would still be homologically non-trivial.

Suppose that S is not normal. Then by Lemma 4.9, S has a simplifying disc. The first and second types of simplifying disc can be used to pattern-isotope S to reduce its weight without increasing $|S \cap \mathcal{H}^1|$. The third type of simplifying disc is a potential pattern-compression disc for S . It satisfies (1) and (2) in Definition 3.1, but possibly not (3). We can pattern-compress S along this disc. If S was a violating disc, then some component of the resulting surface is again a violating disc, and we work with this instead. If S was non-separating, but the pattern-compression produces a violating disc, then we work with this violating disc instead. However, if the pattern-compression creates a component that is a boundary parallel disc, but not a violating disc, then discard it. The resulting surface has smaller extended weight and no greater weight than S . On the other hand, if the pattern-compression does not create a violating disc or a boundary-parallel disc, then some component of the resulting surface is non-separating, has smaller χ_p and no greater weight or extended weight. The fourth type of simplifying disc specifies an isotopy of S , which reduces the number of components of $S \cap \mathcal{H}^1$ and that leaves the weight of S unchanged. At each stage, we are either reducing χ_p or leaving it unchanged and reducing extended weight. So this process terminates. The resulting surface is normal, connected, and orientable. Moreover, it is either non-separating or a violating disc. $\square$

5. DECOMPOSING A HANDLE STRUCTURE ALONG A NORMAL SURFACE

We now wish to show that when (M, P) is decomposed along a normal surface, the resulting manifold is, in some sense, simpler than the original one. The notion of complexity that goes down is defined in terms of handle structures. Its definition needs to take account of the boundary pattern, as follows.

Definition 5.1. Let $\mathcal{H}$ be a handle structure for (M, P) , and let $\mathcal{F}(\mathcal{H})$ be its characteristic surface. For a 0-handle H of $\mathcal{H}$, its *index* $I(H)$ is defined to be the sum of $\chi_p(F)$, over all components F of $H \cap \mathcal{F}(\mathcal{H})$.

Note that $\chi_p(F)$ can be computed from the 0-handles D of F , as follows. Define the *index* of D to be

$$I(D) = \sum_D (-4 + 2|D \cap \mathcal{F}^1| + |D \cap P|).$$

Then $\chi_p(F)$ is $\sum_D I(D)$, as D ranges over each 0-handle of F .

Lemma 5.2 (Behaviour of index under decomposition). *Let $\mathcal{H}'$ be obtained from $\mathcal{H}$ by decomposing along a standard surface S . Let D be a 0-handle of $\mathcal{F}$. Then the sum, over all components D' of $D \cap \mathcal{H}'$, of $I(D')$ is equal to $I(D)$. Hence, for any component F of $\mathcal{F}(\mathcal{H})$, the sum, over all components F' of $F \cap \mathcal{H}'$, of $\chi_p(F')$ is equal to $\chi_p(F)$.*

Proof. Let D be 0-handle of $\mathcal{F}$. Its intersection with S is a collection of properly embedded arcs. Say that there are n of them. (See Figure 8.) Then decomposing D along these arcs gives $n+1$ discs and so this changes the index by $-4n$. The endpoint of each arc either lies in $\mathcal{F}^1(\mathcal{H})$ or in ∂M . In the former case, this component of $\mathcal{F}^1(\mathcal{H})$ is decomposed along the extension of the arc into $\mathcal{F}^1(\mathcal{H})$. Hence this gives an increase of 2 to the index. If the endpoint of the arc lies in ∂M , then we get two new points of intersection between $D \cap \mathcal{H}'$ and the boundary pattern. Hence, again this increases the index by 2. So each arc leads to an increase in index by 4. Therefore, in total, the index is unchanged. $\square$






FIGURE 8. Decomposition of $\mathcal{F}(\mathcal{H})$ along S

We will mostly want to deal with handle structures with the following property.

Definition 5.3. A handle structure $\mathcal{H}$ for (M, P) is *admissible* if the following both hold:

- (1) each 0-handle of $\mathcal{F}$ either has positive index or has zero index and is disjoint from P ;
- (2) for any simple closed curve C lying in a 0-handle of $\mathcal{H}$, that is disjoint from $\mathcal{F}$, and that intersects P at most 3 times, C bounds a disc in $\partial\mathcal{H}^0$ intersecting P in the empty set, an arc or a tripod.

FIGURE 9. The possible 0-handles of $\mathcal{F}$ with negative indexFIGURE 10. The possible 0-handles of $\mathcal{F}$ with zero index

Proposition 5.4 (CREATE ADMISSIBLE HANDLE STRUCTURE). *Let (M, P) be a compact orientable irreducible 3-manifold with a boundary pattern, and let $\mathcal{H}$ be a handle structure for (M, P) . Then there is an algorithm that either creates an admissible handle structure for a manifold (M', P') that is obtained from (M, P) by decomposing along disjoint non-trivial squares or provides a violating disc for (M, P) .*

Proof. The procedure that we will follow will not increase the index of any 0-handle of $\mathcal{F}$. It will strictly decrease the number of 1-handles of $\mathcal{H}$ at each stage, and so it is guaranteed to terminate. Suppose that there is a 0-handle D of $\mathcal{F}$ that either has negative index or has zero index and intersects P . Then D has at most one arc of intersection with $\mathcal{F}^1$. (See Figures 9 and 10.)

Suppose that it has exactly one such arc. Then D lies in a 1-handle H_1 of $\mathcal{H}$ and the arc lies in a 2-handle H_2 of $\mathcal{H}$, and we may remove $H_1 \cup H_2$ from the handle structure. There may be one or two arcs of P running along H_1 . If so, then we need to replace each component of $\mathcal{F}^1 \cap H_2$ with one or two arcs of P .

Suppose now that D is disjoint from $\mathcal{F}^1$. Then it is properly embedded in M . If D has negative index, it intersects P in at most 3 points. We can then determine whether ∂D bounds a disc D' in ∂M intersecting P in either the empty set, an arc or a tripod. If so, $D \cup D'$ bounds a 3-ball in M , by irreducibility. We remove this submanifold bounded by $D \cup D'$, and we insert into D the boundary pattern that was in D' . On the other hand, if ∂D does not bound such a disc D' , then D is a violating disc. It might not be normal, but in this case, we can find another properly embedded disc in $\mathcal{H}^0$, disjoint from $\mathcal{F}$, that also forms a violating disc and intersecting P fewer times. Hence, we may assume that D is normal, and the algorithm terminates with this disc. Suppose now that D has zero index, and hence is a square. If it is a trivial square, it bounds a disc D' in ∂M intersecting P as in (1) or (2) of Lemma 3.10. We can then remove the submanifold bounded by $D \cup D'$, and insert into D the boundary pattern that was in D' . If D is a non-trivial square, then we decompose along it.

Now suppose that there is a simple closed curve C lying in a 0-handle of $\mathcal{H}$, that is disjoint from $\mathcal{F}$, and that intersects P at most 3 times, but that C does not bound a disc in $\partial \mathcal{H}^0$ intersecting P in the empty set, an arc or a tripod. There is always such a curve C when $\partial \mathcal{H}^0 \setminus (\mathcal{F} \cup P)$ is not a union of discs. In that case, we

pick C to be boundary parallel in a component of $\partial\mathcal{H}^0 \setminus (\mathcal{F} \cup P)$ that is not a disc. We can determine whether C bounds a disc D' in ∂M intersecting P in the empty set, an arc or a tripod. If there is no such disc D' in ∂M , then a disc properly embedded in $\mathcal{H}^0$ bounded by C is a violating disc. On the other hand, if there is such a disc D' , then we can remove the submanifold bounded by $D \cup D'$ and insert into D the boundary pattern that was in D' .

We have performed a sequence of decompositions along squares. We need to explain why we can take these squares all to be disjoint. We also need to explain why a violating disc that is produced in the above process actually forms a violating disc for (M, P) . We note that when we decompose along a square, the copies of this disc in the resulting 3-manifold lie in 0-handles and are disjoint from handles of higher index. In particular, they are disjoint from the later decomposing squares. Moreover, they are disjoint from any later violating discs that are produced. Thus, the squares are all disjoint from each other, and any violating disc that is produced is actually a violating disc for (M, P) . $\square$

One of the reasons why we work with admissible handle structures is that they behave well when they are decomposed along a normal surface, as follows.

Lemma 5.5 (Decomposing an admissible handle structure). *Let (M, P) be a compact orientable 3-manifold with a boundary pattern, and let $\mathcal{H}$ be an admissible handle structure for (M, P) . Let S be a normal surface properly embedded in M , and let $\mathcal{H}'$ be the handle structure obtained by cutting along S . Then each 0-handle of $\mathcal{F}(\mathcal{H}')$ has non-negative index.*

Proof. Suppose that D' is a 0-handle of $\mathcal{F}(\mathcal{H}')$ with negative index. The possibilities for D' are shown in Figure 9.

Suppose first that D' intersects $\mathcal{F}^1(\mathcal{H}')$ in a single arc and is disjoint from P' . If $\partial D'$ is disjoint from S , then D' forms a 0-handle of $\mathcal{F}(\mathcal{H})$ that violates the assumption that $\mathcal{H}$ is admissible. Therefore, $\partial D'$ must intersect S and this intersection must be an arc, because $\partial D'$ is disjoint from P' . But this implies that S violated (3) in the definition of normality.

Suppose now that D' intersects $\mathcal{F}^1(\mathcal{H}')$ in a single arc and intersects P' in a point. Again, $\partial D'$ cannot be disjoint from S . It must therefore intersect S in an arc. But this implies that S violated (5) in the definition of normality.

Suppose that D' misses $\mathcal{F}^1(\mathcal{H}')$. Now D' is a subset of a 0-handle D of $\mathcal{F}^0(\mathcal{H})$, and its boundary alternates between arcs of $S \cap D$ and arcs lying in ∂D . Since D' intersects the pattern at most three times, there can be at most three such arcs. But in fact, there must be an even number of them. There must be at least one arc of $S \cap D$ lying in $\partial D'$, as otherwise $D' = D$, contradicting the assumption that $\mathcal{H}$ is admissible. Hence, $D' \cap S$ is an outermost arc in D . The remainder of $\partial D'$ is an arc A in ∂D intersecting P at most once. The arcs of $S \cap (\partial\mathcal{H}^0 \setminus (\mathcal{F} \cup P))$ incident to $\partial D'$ lie in the same component of $\partial\mathcal{H}^0 \setminus (\mathcal{F} \cup P)$ or in adjacent components of $\partial\mathcal{H}^0 \setminus (\mathcal{F} \cup P)$. Hence, by (4) and (6) in the definition of normality, they are either equal or they meet at a point of P . So the endpoints of $D' \cap S$ are joined by an arc of $S \cap (\partial\mathcal{H}^0 \setminus \mathcal{F})$. The union of this latter arc and A is a simple closed curve C that intersects P in at most two points. So by admissibility, C bounds a disc in $\partial\mathcal{H}^0$ intersecting P in an arc or the empty set. But then (7) in the definition of normality is contradicted. $\square$

In an admissible handle structure, we permit some 0-handles of $\mathcal{F}$ to have zero index. These give rise to the following types of handle of $\mathcal{H}$.

Definition 5.6. Let $\mathcal{H}$ be a handle structure for (M, P) . A handle H of $\mathcal{H}$ is an *interior parallelity handle* if one of the following holds:

- (1) it is a 0-handle disjoint from P , each 0-handle of $\mathcal{F} \cap H$ has index zero and $\mathcal{F} \cap H$ is connected;
- (2) it is a 1-handle disjoint from P and each 0-handle of $\mathcal{F} \cap H$ has index zero;
- (3) it is a 2-handle.

Each interior parallelity handle can be given a product structure $D^2 \times I$ so that the intersection with any other 0-handle, 1-handle or 2-handle is of the form $\alpha \times I$ where α is a union of arcs in ∂D^2 . These product structures are viewed as forming an I -bundle structure and this I -bundle structure can be chosen compatibly over all the parallelity handles. The union of the parallelity handles is the *interior parallelity bundle* for $\mathcal{H}$.

We will define more general versions of this structure in Section 14, which will give rise to ‘parallelity bundles’. These important subsets of a handle structure have had numerous applications, for example [14]. We recall the following terminology.

Definition 5.7. Let $\mathcal{B}$ be an I -bundle over a surface F . Then its *vertical boundary* $\partial_v \mathcal{B}$ is the I -bundle over ∂F , and its *horizontal boundary* $\partial_h \mathcal{B}$ is the (∂I) -bundle over F .

When we decompose a handle structure $\mathcal{H}$ along a normal surface, we would like the resulting handle structure $\mathcal{H}'$ to be simpler than $\mathcal{H}$, in some sense. The measure of complexity that we will use is as follows.

Definition 5.8. Define the q -complexity $c_q(\mathcal{H})$ to be

$$\sum_H (\max\{I(H), 0\})^2,$$

where the sum is taken over 0-handles H of $\mathcal{H}$. Here, ‘q’ stands for ‘quadratic’.

Proposition 5.9 (Decomposition and q-complexity). *Let $\mathcal{H}$ be an admissible handle structure for (M, P) . Let S be a compact orientable connected non-separating properly embedded normal surface. Let $\mathcal{H}'$ be the handle structure obtained from $\mathcal{H}$ by cutting along S . Then $c_q(\mathcal{H}') \leq c_q(\mathcal{H})$. Moreover, if this is an equality, then S is homologous in $H_2(M, \partial M)$ to a vertical surface in the interior parallelity bundle for $\mathcal{H}$.*

Proof. Consider any 0-handle H of $\mathcal{H}$. By Lemma 5.2, the sum, over all 0-handles H' of $H \cap \mathcal{H}'$, of $I(H')$ is equal to $I(H)$. By Lemma 5.5, each 0-handle H' has non-negative index. Hence,

$$\sum_{H'} I(H')^2 \leq I(H)^2,$$

with equality if and only at most one 0-handle of $H \cap \mathcal{H}'$ has positive index and the remaining components have zero index. We deduce that $c_q(\mathcal{H}') \leq c_q(\mathcal{H})$.

Suppose that this is an equality. We will show that S is homologous to a surface in $\mathcal{B}$. Consider a component S' of $S \setminus \mathcal{B}$. We will define a transverse orientation on S' and will show that S' with this transverse orientation is homologous to a 2-chain in $\partial M \cup \mathcal{B}$. Since S' is connected, this transverse orientation is either the

same as or opposite to the one it inherits from S . Hence, S' is homologous in $H_2(M \setminus \setminus \mathcal{B}, \partial(M \setminus \setminus \mathcal{B}))$ to a 2-chain in $\mathcal{B}$. Since this is true for each component of $S' \setminus B$, we deduce that S is homologous to a 2-chain in $\mathcal{B}$.

We first define the transverse orientation on the components of $S' \cap \mathcal{H}^0$. Each such component lies within a 0-handle H of $\mathcal{H}$ not lying in $\mathcal{B}$. This 0-handle cannot be disjoint from the 1-handles, since it would then form a 3-ball component of M , which cannot contain a properly embedded non-separating surface. It therefore has at least one 0-handle of $\mathcal{F}(\mathcal{H})$. It cannot be the case that all 0-handles of $H \cap \mathcal{F}(\mathcal{H})$ have zero index since H would then be part of the interior parallelity bundle, contrary to assumption. So $H \cap \mathcal{F}(\mathcal{H})$ has positive index. This is the same index as a unique 0-handle H' of $H \cap \mathcal{H}'$. The remaining 0-handles of $H \cap \mathcal{H}'$ have zero index. We orient this component of $S' \cap \mathcal{H}^0$ away from H' . This transverse orientation extends over the components of $S' \cap \mathcal{H}^1$ as follows. Each such component lies within a 1-handle H_1 that has two components of intersection with $\mathcal{F}^0(\mathcal{H})$. Since $\mathcal{H}$ is admissible and H_1 does not lie in the interior parallelity bundle, these two components of $\mathcal{F}^0(\mathcal{H})$ have positive index. They are divided up by S' into components of $\mathcal{F}^0(\mathcal{H}')$. In each of these two components of $\mathcal{F}^0(\mathcal{H})$, exactly one of these components of $\mathcal{F}^0(\mathcal{H}')$ has positive index. The adjacent discs of $S' \cap \mathcal{H}^0$ point away from this. Hence, we may extend the transverse orientation consistently over the components of $S' \cap \mathcal{H}^1$. Note that S' is actually disjoint from the 2-handles, since these are part of the interior parallelity bundle.

We now show that S' with this transverse orientation is homologous to a 2-chain in $\mathcal{B}$. Each component of $S' \cap \mathcal{H}^i$ is a disc. It lies within a handle H of $\mathcal{H}$. If we were to cut H along this component of $S' \cap \mathcal{H}^i$, we would obtain two 3-balls. We say that the 3-ball component into which this component points is *associated* with it. The union of these balls forms a 3-chain, the boundary of which is the union of S' and a 2-chain in $\partial M \cup \mathcal{B}$. Thus, S' is equal in $H_2(M \setminus \setminus \mathcal{B}, \partial(M \setminus \setminus \mathcal{B}))$ to a 2-chain in $\mathcal{B}$, and so S is equal in $H_2(M, \partial M)$ to a 2-chain in $\mathcal{B}$ with boundary in ∂M . By replacing this 2-chain with a 2-chain that is equal in $H_2(M, \partial M)$, we may assume that it is a surface properly embedded in $\mathcal{B}$ disjoint from $\partial_v \mathcal{B}$. Furthermore, we may assume that it is incompressible and also admits no boundary compression discs disjoint from $\partial_v \mathcal{B}$. Each component of such a surface is either vertical or horizontal in $\mathcal{B}$. But an orientable horizontal surface is homologically trivial. Hence, we have shown that S is homologous to a vertical surface in $\mathcal{B}$. $\square$

Lemma 5.10 (Cutting along a vertical surface). *Let (M, P) be a compact orientable 3-manifold with a boundary pattern. Let $\mathcal{H}$ be an admissible handle structure for (M, P) and let $\mathcal{B}$ be its interior parallelity bundle. Let S be a vertical surface in $\mathcal{B}$ that is properly embedded and non-separating in M . Then S may be placed in normal form and when $\mathcal{H}$ is cut along S , the resulting handle structure $\mathcal{H}'$ satisfies $c_q(\mathcal{H}') = c_q(\mathcal{H})$. Moreover, the following quantities do not increase and at least one strictly decreases: the total number of components of $\mathcal{F}^1$ in the 0-handles of $\mathcal{H}$ with positive index, and the number of interior parallelity 0-handles.*

Proof. The interior parallelity bundle is an I -bundle over a surface F . We give F a handle structure as follows. Each i -handle of $\mathcal{B}$ vertically projects onto an i -handle of F . In addition, we create a 0-handle of F for each component of $\partial_v \mathcal{B} \cap \mathcal{F}^0$ and a 1-handle of F for each component of $\partial_v \mathcal{B} \cap \mathcal{F}^1$. We call the latter handles *boundary handles*, since they form a regular neighbourhood of ∂F . They are added to ensure

that we really do have a handle structure for F . Without them, we would have 1-handles or 2-handles not completely attached to lower index handles.

The vertical surface S projects to a simple closed curve C in F . We isotope C so that it is normal in the handle structure of F . A corresponding isotopy of S makes it normal. When C runs over any boundary 0-handles of F , then the corresponding normal discs of S lie in 0-handles of $\mathcal{H}$ that are not parallelity 0-handles and hence that have positive index. When we decompose along these discs, then this reduces the number of components of $\mathcal{F}^1$ in these handles. When C runs through 0-handles in the interior of F , the corresponding normal discs of S decompose interior parallelity 0-handles of $\mathcal{H}$. This reduces the number of such 0-handles.

By Proposition 5.9, $c_q(\mathcal{H}') \leq c_q(\mathcal{H})$. In fact, this is an equality because the effect on $\mathcal{F}$ of the decomposition is to remove some lines of index zero 0-handles and 1-handles and replace them by two parallel arcs of pattern. Hence for any 0-handle H of $\mathcal{H}$, its index is unchanged by the decomposition. $\square$

The following is easily verified.

Lemma 5.11. *The algorithm CREATE ADMISSIBLE HANDLE STRUCTURE in Proposition 5.4 does not increase any of the following quantities: the q -complexity of $\mathcal{H}$, the total number of components of $\mathcal{F}^1$ in the 0-handles of $\mathcal{H}$ with positive index, and the number of interior parallelity 0-handles.*

We are now in a position to complete the proof of Theorem 2.5.

Proof. We have proved one direction of the theorem in Lemma 3.3. Therefore, let (M, P) be a compact orientable irreducible 3-manifold with an essential boundary pattern. If M is closed, then we are assuming that it contains a properly embedded orientable incompressible surface with no 2-sphere components. We must show that (M, P) admits an essential hierarchy.

If M is closed, we can decompose along the hypothesised surface, and the resulting manifold is irreducible by Lemma 3.9 and has non-empty boundary and essential (empty) pattern. So we may assume that ∂M is non-empty.

Let $\mathcal{H}$ be a handle structure for (M, P) . By Proposition 5.4, there is a decomposition along non-trivial squares, after which the manifold (which we continue to denote by (M, P)) has an admissible handle structure. By Lemma 3.12, these squares are pattern-incompressible and they are automatically incompressible because they are discs. Hence, by Lemma 3.8, the resulting pattern is essential and by Lemma 3.9, the manifold is irreducible.

If ∂M has any sphere components, these bound 3-balls. Thus, we may assume that ∂M is not a union of spheres. Therefore, $b_1(M) > 0$, and so M contains a connected properly embedded non-separating orientable surface. Let S be such a surface with smallest pattern-complexity. By Corollary 3.7, S is pattern-incompressible and incompressible. By Lemma 4.11, it may be placed into normal form.

By Lemma 3.8, the manifold M' obtained by decomposing along S inherits an essential boundary pattern and by Lemma 3.9, it is irreducible. By Proposition 5.9, the handle structure $\mathcal{H}'$ that it inherits has q -complexity at most that of $\mathcal{H}$. Moreover, its q -complexity is strictly less unless S is homologous to a vertical surface in the interior parallelity bundle.

In the case where S is homologous to a vertical surface in the interior parallelity bundle, we instead use Lemma 5.10 to decompose along such a surface. This leaves

$c_q(\mathcal{H})$ unchanged. But it decreases the quantity $m(\mathcal{H})$ which is the sum of the total number of components of $\mathcal{F}^1$ in the 0-handles with positive index, and the number of interior parallelity 0-handles.

Thus, if we consider the ordered pair $(c_q(\mathcal{H}), m(\mathcal{H}))$ and we compare such pairs using lexicographical ordering, then we have reduced this pair. We now apply Proposition 5.4 to turn $\mathcal{H}'$ into an admissible handle structure. By Lemma 5.11, this does not increase $c_q(\mathcal{H}')$ or $m(\mathcal{H}')$. Thus, we can repeat with this new handle structure. Since the complexity pairs are being reduced at each stage, this process must terminate, and so creates an essential hierarchy for (M, P) . $\square$

6. THE ALGORITHM TO DECIDE WHETHER A PATTERN IS ESSENTIAL

In this section, we present our algorithm required by Theorem 1.3. Thus, we are given a triangulation for a compact orientable irreducible 3-manifold M in which P is simplicial, and we must decide whether P is essential. We may assume that ∂M is non-empty, as otherwise, the pattern is trivially essential. The first step is to build a handle structure $\mathcal{H}$ for (M, P) . In our algorithm, we will store several lists:

- (1) handle structures $\mathcal{H}_1, \dots, \mathcal{H}_\ell$ for 3-manifolds $M_1, \dots, M_\ell$ with boundary patterns $P_1, \dots, P_\ell$;
- (2) normal surfaces $S_1, \dots, S_{\ell-1}$, where each S_i is a normal surface in $\mathcal{H}_i$ and where $(M_{i+1}, P_{i+1}) = (M_i, P_i) \setminus S_i$ with $\mathcal{H}_{i+1}$ its induced handle structure.

The algorithm is as follows:

- (1) Set $\mathcal{H}_1$ to be the initial handle structure $\mathcal{H}$ and set ℓ to be 0.
- (2) Apply the algorithm CREATE ADMISSIBLE HANDLE STRUCTURE in Proposition 5.4 to $\mathcal{H}_{\ell+1}$. If this produces a violating disc D , go to Step (12).
- (3) If Proposition 5.4 creates an admissible handle structure, then add any decomposing squares that have been used to the hierarchy. Increase ℓ by 1 (if no squares were used) or 2 (if squares were used).
- (4) Using linear algebra compute $b_1(M_\ell)$.
- (5) If $b_1(M_\ell) = 0$, then go to Step (9).
- (6) By reaching this step, $b_1(M_\ell)$ must be positive. If there is a vertical connected surface in the interior parallelity bundle for M_ℓ that is non-separating in M_ℓ , let S_ℓ be such a surface, isotoped so that is normal as in Lemma 5.10. Otherwise, let S_ℓ be any connected, non-separating, properly embedded, oriented normal surface S_ℓ in M_ℓ . (See Theorem 6.2 below for a simple and fast construction of such a surface.)
- (7) Is some component of S_ℓ a violating disc for (M_ℓ, P_ℓ) ? If yes, then set D to be this component and go to Step (12).
- (8) Set $(M_{\ell+1}, P_{\ell+1})$ to be $(M_\ell, P_\ell) \setminus S_\ell$ and set $\mathcal{H}_{\ell+1}$ to be its induced handle structure. Go to Step (2).
- (9) By reaching this step, $b_1(M_\ell)$ must be 0, and so ∂M_ℓ must be a union of spheres. Examine whether the intersection of P_ℓ with each such sphere is connected. If not, then we obtain a violating disc. Then examine all isotopy classes of simple closed loops in ∂M_ℓ that intersect P_ℓ transversely in at most 3 points. (Equivalently, consider whether the intersection between P_ℓ and each component of ∂M_ℓ is connected. If it is, examine all closed loops of length at most 3 in the dual graph.) For each such simple closed curve, determine whether it bounds a violating disc.

- (10) If no violating disc is found, then (M_ℓ, P_ℓ) is essential, and therefore by Theorem 2.5, so is (M, P) . In this situation, the algorithm terminates with a declaration that P is essential and it outputs the hierarchy that has been constructed.
- (11) By reaching this step, ∂M_ℓ is a union of spheres and the algorithm has found a closed loop in ∂M_ℓ that bounds a violating disc D .
- (12) Is D a violating disc for the initial manifold (M, P) ? If so, then the algorithm terminates by declaring that P is inessential and it outputs the violating disc D .
- (13) By reaching this step, D must run over at least one surface in the hierarchy. Let S_j be the last surface in the hierarchy that intersects D .
- (14) Is D a compression disc or pattern-compression disc for S_j ? If not, then $\partial D \cap S_j$ must in fact be a boundary-parallel arc in S_j , separating off a disc D' that intersects the pattern P_j at most once. Moreover, $D \cup D'$ must be a violating disc for (M_j, P_j) . In this case, replace D by $D \cup D'$ and go to Step (12).
- (15) By reaching this step, D forms a compression disc or pattern-compression disc for S_j in M_j . Compress or pattern-compress S_j along this disc to form a properly embedded surface. Perform isotopies and possibly compressions and pattern-compressions to this surface, and possibly discard components, to obtain a compact orientable connected properly embedded normal surface $\bar{S}_j$ that is non-separating in M_j or that is a violating disc for P_j . Replace S_j by $\bar{S}_j$. Discard all later surfaces in the hierarchy. Set ℓ to be equal to j . Go to Step (7).




FIGURE 11. The algorithm to determine whether a pattern is essential

Theorem 6.1 (The algorithm works). *This algorithm terminates with the correct declaration either that P is essential or that P is inessential. Moreover when P is essential, the algorithm outputs an essential hierarchy for (M, P) . When P is inessential, it outputs a violating disc.*

Proof. It is clear that if the algorithm terminates, then its output is correct. For if it terminates by declaring that P is inessential, then it does so because it has found a violating disc. If it terminates by declaring that P is essential, then it does so because it has found an essential hierarchy for (M, P) .

We must show that the algorithm terminates. We first note that it cannot remain in the hierarchy construction loop (Steps (2) to (8)) forever. Suppose that, on the contrary, it does. Each time we go through Step (6), we have an admissible handle structure $\mathcal{H}_\ell$. If there is no non-separating vertical surface in the interior parallelity bundle, then any connected normal orientable non-separating surface is chosen or a violating disc is found. If a violating disc is found, then we leave the hierarchy construction loop. If a connected normal orientable non-separating surface is chosen, then by Proposition 5.9, cutting along this surface strictly reduces the q -complexity $c_q(\mathcal{H}_\ell)$. When there is a vertical surface, then by Lemma 5.10, cutting along it leaves $c_q(\mathcal{H}_\ell)$ unchanged, but it reduces $m(\mathcal{H}_\ell)$. Then when the handle structure is made admissible in Step (2) using Proposition 5.4, none of these quantities increases. Hence, the ordered pair $(c_q(\mathcal{H}_\ell), m(\mathcal{H}_\ell))$ decreases at each stage. But this contradicts the fact that these pairs are well-ordered.

So, if the algorithm does not terminate, it must leave the hierarchy construction loop infinitely often. Index these occurrences by the natural numbers i . At the i th stage, let

$$(S_1^i, \dots, S_{\ell(i)}^i)$$

be the partial hierarchy that is currently stored. When the algorithm leaves the hierarchy construction loop and does not immediately then terminate, a violating disc D is found. Furthermore, D runs over at least one surface in the partial hierarchy, unless the algorithm immediately terminates. Let S_j^i be the last surface that it runs over. Suppose first that D is not a compression disc or pattern-compression disc for S_j^i . Then $D \cap S_j^i$ is either a simple closed curve that bounds a disc in S_j^i or an arc in S_j^i that separates off a disc in S_j^i intersecting P_j at most once. If $D \cap S_j^i$ is a simple closed curve bounding a disc D' in S_j^i , then $D \cup D'$ is a sphere. By the irreducibility of M , this bounds a ball in M . The boundary of this ball is disjoint from the surfaces $S_1^i, \dots, S_{j-1}^i$. Hence, each of these surfaces must be disjoint from the ball, by our assumption that each surface is either connected and non-separating or a disjoint union of non-trivial squares. Thus, $D \cup D'$ bounds a ball in M_j . None of the later surfaces $S_{j+1}^i, \dots, S_{\ell(i)}^i$ can enter this ball, again by our assumption that each is either connected and non-separating or a disjoint union non-trivial squares. Therefore, D' lies in $M_{\ell(i)}$ and its interior is disjoint from $P_{\ell(i)}$. This contradicts the assumption that D is a violating disc for $(M_{\ell(i)}, P_{\ell(i)})$. Suppose now that $D \cap S_j^i$ is an arc in S_j^i that separates off a disc D' in S_j^i intersecting P_j at most once. Then $D \cup D'$ is a disc properly embedded in M_j intersecting the pattern at most twice. Suppose that this is not a violating disc for (M_j, P_j) . Then $\partial(D \cup D')$ bounds a disc D'' in ∂M_j that intersects P_j either in a single arc of the empty set. The union of these three discs bounds a ball that none of the earlier surfaces or later surfaces can enter. Hence, $D' \cup D''$ contains an arc or tripod of $P_{\ell(i)}$. This contradicts the fact that D is a violating disc for $(M_{\ell(i)}, P_{\ell(i)})$. Hence, the algorithm does produce a violating disc for (M_j, P_j) , as claimed in Step (14).

Hence, the algorithm eventually produces a compression or pattern-compression disc for some S_j^i and the later surfaces are discarded. This reduces $\chi_p(S_j^i)$. Now, $\chi_p(S_1^i)$ can decrease only finitely many times. Hence, there is some point after

which S_1^i remains unchanged. Let S_1 be this final surface. After this point, the manifold obtained by cutting along S_1 remains unchanged. We therefore consider the surfaces S_2^i after this point. Again, these can change only finitely many times, and so there is some final surface S_2 . Repeating this process, we obtain an infinite sequence of surfaces $S_1, S_2, \dots$. We must show that this cannot occur. Let $\mathcal{H}_j$ be the handle structure obtained from $\mathcal{H}$ by cutting along $S_1, \dots, S_j$.

Now we note that if some surface in one of these hierarchies is a disc that intersects the pattern four times, and at some stage we decrease its pattern-complexity, then this creates a violating disc. This is then used to simplify an earlier surface in the hierarchy. Thus, none of the surfaces S_j in our infinite list is obtained from a square by pattern-compressing.

Thus, the surfaces $S_1, S_2, \dots$ and handle structures $\mathcal{H}_1, \mathcal{H}_2, \dots$ have the following property. Suppose that some $\mathcal{H}_j$ is admissible. If there is no non-separating vertical surface in the interior parallelity bundle, then S_j is some connected normal orientable non-separating surface. On the other hand, if there is such a vertical surface, then S_j is as in Lemma 5.10. Then the next surfaces after S_j are possible squares. The result is an admissible handle structure, and so on. As argued in the proof of Theorem 2.5, this is impossible, because the ordered pair $(c_q(\mathcal{H}_j), m(\mathcal{H}_j))$ decreases as we pass from one admissible handle structure to the next one. $\square$

One reason why the above algorithm works so well is that, in a compact orientable 3-manifold M with $b_1(M) > 0$, it is a very easy to construct a connected, non-separating, properly embedded, oriented surface. We make this precise in the following theorem. This seems to be well known, but it is so central to the efficiency of our algorithm that we give a proof.

Theorem 6.2 (BUILD NON-SEPARATING SURFACE). *There is an algorithm that takes, as its input, a triangulation $\mathcal{T}$ or handle structure $\mathcal{H}$ for a compact orientable 3-manifold M with boundary pattern P , and with $b_1(M) > 0$. In the case where ∂M is non-empty, we also assume that $\mathcal{H}$ has no 3-handles. The algorithm builds a connected, non-separating, properly embedded, oriented, normal surface S with weight at most $c^{|\mathcal{T}|}$ or extended weight at most $c^{|\mathcal{H}|}$. Here, c is a constant that is universal in the case of a triangulation. In the case of a handle structure, c depends on the number N , which is the maximum, over each 0-handle of $\mathcal{H}$, of the number of components of intersection between the 0-handle and each of the following: the 1-handles, the 2-handles, the vertices of P and the edges of P . The running time is bounded above by a polynomial function of $|\mathcal{T}|$ or $|\mathcal{H}|$, where the polynomial is universal for triangulations and depends on N in the case of handle structures.*

Proof. Let us initially assume that M is given via a triangulation $\mathcal{T}$, since this case is more straightforward.

We are assuming that $H^1(M)$ is non-zero. Any non-zero element of this group can be represented by a simplicial 1-cocycle, and indeed such a 1-cocycle is unique if we require that its evaluation on each edge in some given maximal tree in the 1-skeleton is zero. So, one can set up a system of linear equations, with a variable for each oriented edge not in this tree, and an equation specifying that the total evaluation around each face is zero. Then any non-zero real solution to these equations represents a non-trivial element of $H^1(M; \mathbb{R})$. In fact, one can ensure that the coefficients of this element are rational, by imposing enough extra constraints specifying that some edges have zero evaluation, until the solution space is

1-dimensional, and then requiring that the sum of the coefficients is 1. Then we can multiply through by a suitable integer to obtain a non-zero solution with integer coefficients, which is a 1-cocycle representing a non-trivial element of $H^1(M)$.

The size of the co-ordinates of this solution are bounded above by $2^{12|\mathcal{T}|}$ for the following reason. The system of equations has at most $n = E - V + 1$ variables, where E and V are the number of edges of $\mathcal{T}$. This is because a maximal tree has $V - 1$ edges in it. Note that $E - V + 1 \leq E \leq 6|\mathcal{T}|$. There is an equation for each face, as well as an equation that requires that the sum of the coefficients is 1. We also have extra conditions that require that some edges evaluate to zero. These equations are combined into a single matrix equation $Av = w$, where A represents the above equations, v is the solution, and w is the vector with all zeroes except the final entry, which is 1. We are assuming that the kernel of A is zero, and hence the rank of A is equal to the number of variables. There is a square submatrix B of A with the same number of columns of A having the same rank. Thus B is invertible. We have $Bv = w'$, where w' consists of some of the entries of w , and so all entries are zero apart from one. Then $v = B^{-1}w'$. Note that $B^{-1} = (\det(B))^{-1}\text{adj}(B)$, where $\det(B)$ is the determinant of B and $\text{adj}(B)$ is the adjugate matrix. We can obtain an integral solution by multiplying through by $\det(B)$. Each co-ordinate of this integral solution is an entry of the adjugate matrix, which is the determinant of a minor of B . The size of this is at most the product of the l^2 norms of the rows, which is at most $3^n\sqrt{n} \leq 2^{2n} \leq 2^{12|\mathcal{T}|}$, as required. Note that the procedure used to produce this solution runs in polynomial time, since it involves computing $(n - 1)^2$ determinants.

From this, one can build a properly embedded surface representing an element of $H_2(M, \partial M)$ dual to this 1-cocycle c , as follows. First, for each edge e of the triangulation, place a collection of $|c(e)|$ points along the interior of the edge. These will be the intersection points between e and the surface. Pick a circular ordering on the edges of each face. Since the evaluation of c around the face is zero, we may join the points in the edges in pairs by a collection of disjoint normal arcs, so that for each arc, one endpoint lies on an edge with positive evaluation and one endpoint lies on an edge with negative evaluation. These arcs form the intersection between the surface and the face. In the boundary of each tetrahedron, these arcs combine to form a collection of normal curves, each with length three or four, and we may fill these in with normal triangles and squares. The result is a compact oriented properly embedded surface dual to c . Any component of this surface can be used as the required surface S . Such a component can be found using the algorithm of Agol, Hass and Thurston [1, Corollary 17].

We now consider the case where M comes equipped with a handle structure $\mathcal{H}$. When ∂M is non-empty, we are assuming that $\mathcal{H}$ has no 3-handles. This determines a cell structure on M , as follows. Each handle becomes a 3-cell. The 2-cells arise from the components of intersection between handles, as well as the components of intersection between a handle and $\partial M \setminus P$. Some of these surfaces might not be discs, in which case we decompose them along arcs forming 1-cells, until they become discs. We will explain how to do this below. The remaining 1-cells arise from the components of intersection between any three handles, as well as the components of intersection between pairs of handles and $\partial M \setminus P$, and also the components of intersection between a handle and an edge of P . Finally, the 0-cells arise as follows: the components of intersection between four handles; the

components of intersection between any three handles and $\partial M \setminus P$; the components of intersection between any two handles and P ; the vertices of P . One arrangement when non-disc components of intersection between handles arises is when a 1-handle of $\mathcal{H}$ is disjoint from $\mathcal{H}^2$ and P . In that case, we insert a single vertical 1-cell running along the 1-handle. The other scenario where non-disc components arise is in the boundary of a 0-handle, since the intersection between the 0-handle and $\partial M \setminus P$ or $\mathcal{H}^3$ may have components other than discs. In that case, we insert properly embedded 1-cells into the planar surface, but without adding any new 0-cells.

We now pick a maximal tree in the 1-skeleton of this cell structure. Our approach will depend on whether M is closed or has non-empty boundary. When M is closed, we pick a maximal tree in the 1-skeleton of each 3-handle. The union of these trees forms a forest in M . In each 1-handle disjoint from the 2-handles, we also add the new vertical 1-cell to this forest. We then extend this to a maximal tree in all of M . The effect of this is that, for each 1-cocycle c that is zero on the edges in the maximal tree, the evaluation of each 1-cell in each 3-handle is zero. This is automatic for 1-cells in the maximal tree. But for any other 1-cell in a 3-handle, the evaluation of c on that edge is equal to the evaluation of the loop that runs along the edge, then through the maximal tree back to the start of the edge. This loop lies entirely in the 3-handle, and so is homotopically trivial, and hence its evaluation is zero, as claimed. When M has non-empty boundary, we pick a maximal tree in the 1-skeleton of each 2-handle and again any new vertical edges in a 1-handle are included in this forest. We also add each component of $P \cap \mathcal{H}^1$ to this forest. Furthermore, for each disc component of $\mathcal{H}^1 \cap \partial M$, we pick one of its arcs of intersection with $\mathcal{H}^0$ and we declare that all but one of the 1-cells in that arc lies in the forest. We again extend this to a maximal tree in M . This has the effect that the evaluation of each 1-cell in each 2-handle is actually zero.

As above, such cocycles c are solutions to a system of linear equations. The variables correspond to the oriented 1-cells not in the maximal tree, and there is an equation for each 2-cell. The argument above also gives that there is a non-zero integral cocycle with the property that its value on each 1-cell has modulus that is bounded above by an exponential function of $|\mathcal{H}|$, where the exponent only depends on the uniform type of $\mathcal{H}$.

We now construct a surface S that is dual to this 1-cocycle c . For each 1-cell e , we insert $|c(e)|$ points into its interior. These will be the points of intersection between S and e . Each 2-cell is a polygon, and we have specified the intersection between S and the boundary of this polygon. Since the total evaluation of the oriented edges in this boundary is zero, we can insert transversely oriented arcs into the interior of the polygon joining these points. Note that the 1-cells that are vertical in the 1-handles have zero evaluation, for the following reason. Each such 1-cell is either in the boundary of a 3-handle, or in the intersection between a 2-handle and ∂M , or is an inserted 1-cell in a 1-handle disjoint from the 2-handles. For each such 1-cell, its evaluation is zero. Each 2-cell in the vertical boundary of a 1-handle has four sides, and we have just shown that the two vertical sides have zero evaluation. Hence, the remaining two sides have the same evaluation. So the arcs that we insert into such a 2-cell can be isotoped to be vertical. We can also ensure that the arcs in the two polygons at the ends of this 1-handle correspond via the product structure. In the boundary of each handle, these arcs patch together to form closed curves. We can insert discs into the handle bounded by these curves. The result is a transversely

oriented, properly embedded surface dual to c . It is standard because it is disjoint from the 3-handles and respects the product structure of the other handles of $\mathcal{H}$. For 1-handles, we have already arranged this. For 2-handles in a closed 3-manifold M , this follows from the fact that the surface misses the 3-handles and intersects each 1-cell in points of the same sign. When M has non-empty boundary, the surface misses the 2-handles. Note also that all the points of intersection between the surface and any arc component of $\mathcal{H}^0 \cap \mathcal{H}^1 \cap \partial M$ have the same sign.

We need to arrange that this surface is in fact normal. We observe that when M is closed, we have in fact already arranged this. When M is closed, the only normality condition that is not automatic is (3) of Definition 4.6. However, for each disc that we have inserted into a 0-handle, it runs over each component of $\mathcal{F}^1$ at most once. For if it ran over a component of $\mathcal{F}^1$ more than once, then there would be two arcs of intersection in that component of $\mathcal{F}^1$ with opposite transverse orientations. However, all the arcs of intersection between S and any component of $\mathcal{F}^1$ are compatibly oriented.

We now deal with the case where M has non-empty boundary. Each component of $S \cap \mathcal{H}^0$ intersects each component of $\partial \mathcal{F}^0 \setminus (P \cup \mathcal{F}^1)$ and each component of intersection between the edges of P and $\mathcal{H}^0$ at most once, since S intersects each such component in points of the same sign. (This latter property is convenient at this stage, but we will not maintain it in the modifications that we make below.) So this results in finitely many types of disc in $S \cap \mathcal{H}^0$. For each such disc type D , we define its *complexity* to be $|D \cap P|^2 + |D \cap \mathcal{F}^1|^2$. This quantity is bounded above by a constant depending only on the type of $\mathcal{H}$. The *complexity* of S is equal to the sum of the complexities of its disc types. Note that in this sum, we only consider one disc of $S \cap \mathcal{H}^0$ of each type in each 0-handle. Hence, the complexity is bounded above by a constant times $|\mathcal{H}^0|$, where the constant only depends on the uniform type of $\mathcal{H}$.

We now examine the normality conditions (1)-(7) from Definition 4.6. Conditions (1) and (2) are automatically satisfied. In our setting, S is disjoint from the 2-handles, and so (3) and (5) are satisfied. If a curve of $S \cap \partial \mathcal{H}^0$ intersects a component of $\partial \mathcal{H}^0 - (\mathcal{F} \cup P)$ in more than one arc, contradicting (4), then we can pattern compress such a disc. We can pattern compress all discs of this type. For each such disc, this creates two discs D_1 and D_2 . Discs of these types may already exist in S , but this does not affect our argument. The points of intersection between D and P are shared out between D_1 and D_2 . The same is true for $D \cap \mathcal{F}^1$. Hence, the total complexity of S decreases. Note that when a pattern-compression is performed, then the class that S represents in $H_2(M, \partial M)$ is unchanged. A similar argument works when (6) in the definition of normality is violated. We pattern compress the disc D to discs D_1 and D_2 . This introduces two points of intersection with P , but the complexity of these discs still decreases. Finally note that condition (7) must be satisfied, since if it were not, then we obtain two points of intersection between S and a component of $\mathcal{H}^1 \cap \mathcal{H}^0 \cap \partial M$ that are oriented towards or away from each other, contrary to our construction. Thus we have constructed a normal surface. Any component of this is our required normal surface. $\square$

Remark 6.3. While this gives a simple construction of some surface S , mostly using only elementary linear algebra, it is not necessary to follow this construction in step (6) of the above algorithm to determine whether a boundary pattern is

essential. Indeed, any connected, non-separating, properly embedded, oriented, normal surface can be used.

Remark 6.4. When Theorem 6.2 is applied to a handle structure $\mathcal{H}$, both the upper bound on the extended weight of the surface and the running time of the algorithm depend on the quantity N . In most applications, N is bounded above by a universal constant. For example, this is true when $\mathcal{H}$ is dual to a triangulation of M in which P is simplicial. More generally, we will see in Section 11 that typically $\mathcal{H}$ can be arranged to be of ‘uniform type’ as one proceeds down a hierarchy. This term is defined in Definition 11.2, but it implies that N is bounded above by a universal constant.

7. AN EXAMPLE

We give an example of this algorithm in practice. We use it to show that the tripus 3-manifold has incompressible boundary. This manifold is shown in (i) of Figure 12. There, a 3-ball (represented by the outer circle) is shown, containing a properly embedded graph. The tripus manifold is the exterior of this graph. It is a famous example of a 3-manifold admitting a hyperbolic structure with totally geodesic boundary. Totally geodesic boundary components of a hyperbolic 3-manifold are always incompressible. However, the application of our algorithm gives a straightforward alternative proof of this.

The first decomposing surface is a thrice-punctured sphere as shown in (ii) of Figure 12. The exterior of this surface, with its boundary pattern, is shown in (iii). The next surface is the four discs shown in red in (iv). The exterior of these discs is two 3-balls. Their boundary patterns are shown in (v) and (vi). It is easily checked that they are essential. Hence, the empty boundary pattern for the tripus manifold is essential. In other words, the boundary is incompressible. Furthermore, the manifold is irreducible.

A skeptical reader may object that we have not quite followed the algorithm from Section 6, in several ways. We have not specified an initial handle structure $\mathcal{H}$ for the manifold. Of course we could have done so, and we could have realised the first surface as normal, and we could have decomposed $\mathcal{H}$ along it, thereby forming a new handle structure $\mathcal{H}_2$, and so on. Although a computer implementation would have used handle structures, we have chosen to avoid them in this example, for ease of exposition.

We deviated from the algorithm in another way, because the second surface is a union of discs which is separating, whereas the surfaces in Step (6) are required to be connected and non-separating. However, we could instead have chosen three of the four discs, and decomposed along each of these in turn. These are connected and non-separating. The resulting manifold would have been a single 3-ball. Its boundary pattern would have been essential, by an application of Lemma 3.2 and because decomposition along the fourth disc gives an essential pattern. We chose to decompose along four discs rather than three, again for ease of visualisation and exposition.

A reader may also object that we have been very fortuitous that we did not need to leave the hierarchy construction loop. So suppose instead that we had initially used a slightly different set of decomposing surfaces. Suppose that we used the same first surface, but that the second surface was a slightly different collection of discs S_2 , as shown in (i) of Figure 13. Then when we decompose along these




FIGURE 12. An essential hierarchy for the tripus manifold

discs, we get two three balls, with boundary patterns as shown in (ii) and (iii). The pattern in (ii) is inessential; a curve bounding a violating disc D is shown. This would have been discovered in step (9) of the algorithm. It runs over one of the discs in S_2 . When we pattern-compress S_2 along D , the result is five discs, one of which is boundary-parallel. Removing this from the collection, we obtain the four discs described in Figure 12 (iv). Hence, the algorithm would again terminate with the correct declaration that the boundary of the tripus manifold is incompressible.

8. THE ALGORITHM TO DETERMINE INCOMPRESSIBILITY

In this section, we complete the proof of Theorem 1.1.

Lemma 8.1. *Let S be a compact orientable surface properly embedded in a compact orientable 3-manifold M . Let D be a compression disc for ∂M that is disjoint from*




FIGURE 13. An alternative decomposing surface. A curve bounding a violating disc is shown.

S. Then *S* is incompressible in *M* if and only if it is incompressible in $M \setminus \setminus D$. Also, *S* is boundary-incompressible in *M* if and only if it is boundary-incompressible in $M \setminus \setminus D$.

Proof. One direction is clear. A compression or boundary-compression disc for *S* in $M \setminus \setminus D$ can be slid off the copies of *D* in $\partial(M \setminus \setminus D)$. It then becomes a compression or boundary-compression disc for *S* in *M*.

Suppose that *D'* is a compression or boundary-compression disc for *S* in *M*. The intersection between this and *D* is a collection of simple closed curves and properly embedded arcs. Consider a simple closed curve that is innermost in *D*. It bounds a disc in *D*. We may surger *D'* along this disc, and remove the resulting 2-sphere component. The result is a new compression or boundary-compression disc for *S* that intersects *D* a fewer number of times. Hence, we may assume that $D \cap D'$ contains no simple closed curves. When *D'* is a compression disc, this implies that $D \cap D' = \emptyset$, since if there were an arc of intersection, its boundary would lie in *S*, whereas *D* is assumed to be disjoint from *S*. So suppose that *D'* is a boundary-compression disc for *S*. Now consider an arc of $D' \cap D$ that is outermost in *D*. It separates off a sub-disc in *D* with interior disjoint from *D'*. We may surger *D'* along this disc, giving two discs D'_1 and D'_2 . One of these is a boundary-compression disc for *S*. Hence, again by minimising $|D \cap D'|$, we may assume that $D \cap D' = \emptyset$. But then *D'* lies in $M \setminus \setminus D$, and it forms a compression or boundary-compression disc for *S* in $M \setminus \setminus D$. $\square$

Theorem 8.2 (DETERMINING INCOMPRESSIBILITY AND BOUNDARY-INCOMPRESSIBILITY). *There is an algorithm that takes, as its input a compact orientable irreducible 3-manifold M and a compact orientable properly embedded surface S , and determines whether S is incompressible and boundary-incompressible. Furthermore, when S is compressible or boundary-compressible, the algorithm outputs a compression disc or a boundary-compression disc.*

Proof. We may remove the disc components of S without affecting whether S is compressible or boundary-compressible. So suppose that S has no disc components. Give M the empty boundary pattern, and let P be the pattern that $M \setminus S$ inherits. Running the algorithm provided by Theorem 1.3 on $(M \setminus S, P)$ establishes whether P is essential. Furthermore, when it is inessential, the algorithm provides a violating disc.

First suppose that P is essential. Then S is incompressible and boundary-incompressible, by Lemma 3.2. Hence, in this case, our algorithm terminates with the declaration that S is incompressible and boundary-incompressible.

Now suppose that P has a violating disc D . Then ∂D intersects P an even number of times, since P is where the copies of S in $\partial(M \setminus S)$ meet the parts of the boundary coming from ∂M . So, ∂D intersects P zero times or twice. If ∂D intersects P twice, then D forms a potential boundary-compression disc for S . If this is an actual boundary-compression disc, then the algorithm terminates and provides this disc. On the other hand, if $D \cap S$ separates off a subdisc D' from S , then $D \cup D'$ yields a violating disc for P that is disjoint from P . So we may suppose that ∂D is disjoint from P . It then lies in a copy of S or in $\partial M \setminus S$. In the former case, D is a compression disc for S . Suppose that D lies in $\partial M \setminus S$. If it forms a compression disc for ∂M , then we cut M along D and by Lemma 8.1, S is compressible or boundary-compressible in $M \setminus D$ if and only if it was beforehand. Thus, we can work with this new manifold, which has boundary with higher Euler characteristic. Suppose that D is not a compression disc for ∂M . Then ∂D bounds a disc D' in ∂M . The boundary of S must intersect this disc, since D is a violating disc. Pick an innermost curve of $\partial S \cap D'$ in D' . This bounds a disc in D' that forms a compression disc for S , since we are assuming that S has no disc components. $\square$

Theorem 8.3 (DETERMINING INCOMPRESSIBILITY). *There is an algorithm that takes, as its input, a compact orientable irreducible 3-manifold M and a compact orientable properly embedded surface S , and determines whether S is incompressible. Furthermore, when S is compressible, the algorithm outputs a compression disc.*

Proof. Again we may assume that S has no disc components. Give M the empty boundary pattern, and let P be the pattern that $M \setminus S$ inherits. Now replace each curve of P by two parallel curves, forming a new pattern P' . The algorithm given in Theorem 1.3 is then applied to $(M \setminus S, P')$. If P' is essential, then S is incompressible and the algorithm terminates with this declaration. So suppose that P' has a violating disc D . The boundary of $M \setminus S$ consists of parts coming from S and parts coming from ∂M , and between these lie annuli lying between parallel curves of P' . Hence, ∂D intersects P' an even number of times. Moreover, if it intersects P' twice, then one arc of $\partial D \setminus P'$ is boundary parallel in an annular component of $\partial(M \setminus S) \setminus P'$. In that case, we may isotope ∂D to form a violating disc disjoint from P' . So we may assume that D is disjoint from P' . It either

forms a compression disc for S or it is disjoint from S . In the former case, the algorithm terminates and provides this disc D . In the latter case, it either forms a compression disc for ∂M , in which case we compress along it and continue, or it can be used to find a compression disc for S lying in ∂M . $\square$

9. THE NUMBER OF STEPS

There are various steps in the algorithm from Theorem 1.3 that are not described with enough precision to be able to estimate their running time. So, we will initially content ourselves with an analysis of the number of times that the algorithm goes through any of the steps in Section 6. We do this in terms of two quantities:

- (1) the maximal length L of any of the hierarchies; thus the variable ℓ will be at most L throughout;
- (2) the maximal pattern-complexity g of any decomposing surface.

At present, we will not attempt to bound these two quantities. But we can bound the number of steps in the algorithm in terms of L and g .

When the algorithm is in the hierarchy construction loop, namely Steps (2)-(8), it continues until a violating disc is found or until it reaches a manifold $M_{\ell+1}$ with $b_1(M_{\ell+1}) = 0$. The number of times it can iterate this loop is at most L . By this stage, we have produced a list of surfaces $S_1, \dots, S_\ell$ with pattern-complexity $\chi_p(S_1), \dots, \chi_p(S_\ell)$. It may be the case that the final surface S_ℓ has negative pattern-complexity, in which case it is a violating disc (because it is assumed to be connected and non-separating), and we do not include it in this list. The algorithm then proceeds through the surface simplification process, namely Steps (9)-(15). If it does not terminate here, then we pattern-compress or compress some S_j and thereby reduce its pattern-complexity by at least one. We discard the later surfaces and return to Step (7) to create new ones.

In order to find an upper bound on the number of times we can leave the hierarchy construction loop, we define the (g, L) -complexity of the hierarchy $S_1, \dots, S_\ell$ to be

$$\sum_{i=1}^{\ell} \chi_p(S_i)(g+1)^{L-i}.$$

In other words, we extend the list $\chi_p(S_1), \dots, \chi_p(S_\ell)$ to one of length L by adding zeroes, and then we view this L -tuple of integers as the digits of a number in base $g+1$. This number is the (g, L) -complexity. Importantly, it goes down by at least one between successive occasions where we leave the hierarchy construction loop. This is where we compress or pattern-compress a surface S_j and we discard the later surfaces beyond S_j . When new ones are created to replace them, all we have is an upper bound of g on their pattern-complexity. Thus, the L -tuple of pattern-complexities might change as follows:

$$(\chi_p(S_1), \dots, \chi_p(S_j), \chi_p(S_{j+1}), \dots, \chi_p(S_L)) \longrightarrow (\chi_p(S_1), \dots, \chi_p(S_j) - 1, g, \dots, g).$$

But if we view these L -tuples as the digits of a number in base $g+1$, we see that this number has gone down by at least 1. Since the (g, L) -complexity is never more than $(g+1)^L$, we deduce that the number of times we leave the hierarchy construction loop is at most $(g+1)^L$. Hence, the number of times that we perform any step of the algorithm is at most $L(g+1)^L$. The extra factor L is because we may go round the hierarchy construction loop at most L times before leaving it. We summarise this as follows.

Proposition 9.1 (Number of iterations). *The number of times that the above algorithm can iterate through any step is at most $L(g+1)^L$, where L is the maximal length of any of the hierarchies and g is the maximum pattern-complexity of any of the surfaces.*

10. A REFINED ALGORITHM USING INTERIOR PARALLELITY BUNDLES

We saw in Proposition 9.1 that the quantity L , which provides an upper bound to the length of any of the partial hierarchies that we construct, is key to controlling the running time of the algorithm. With our current algorithm, it is not so straightforward to give good estimates on L . This is because the proof at the end of Section 5 that the hierarchies terminate uses the ordered pairs $(c_q(\mathcal{H}), m(\mathcal{H}))$. Although lexicographical ordering on these pairs is enough to establish termination, it does not establish how quickly this happens. Now under many circumstances $c_q(\mathcal{H})$ strictly decreases. The quantity $m(\mathcal{H})$ only becomes relevant when $\mathcal{H}$ has a non-separating vertical surface in its interior parallelity bundle. Therefore, in this section, we show how we may remove every component of the interior parallelity bundle that is not an I -bundle over a disc. This will lead to a somewhat more complicated algorithm. But it will mean that we can provide a linear bound on L in terms of the initial q -complexity of $\mathcal{H}$. The first change that we make to the algorithm is as follows.

First modification to the algorithm: After Step (2), once the handle structure has been made admissible, we then do the following:

- (2') Replace each component of the interior parallelity bundle that is an I -bundle over a disc by a single 2-handle.
- (2'') Cut along the vertical boundary components of the remaining I -bundles and further decompose these using squares into 3-balls with essential boundary patterns.

This ensures that the resulting 3-manifold has interior parallelity bundle that consists solely of 2-handles and hence cannot contain a non-separating vertical surface. Therefore, by Proposition 5.9, the q -complexity $c_q(\mathcal{H}_i)$ of the handle structure $\mathcal{H}_i$ strictly decreases when we decompose along a connected non-separating normal surface S_i .

However, this algorithm has an extra complication. In the previous algorithm described in Section 6, if a non-trivial square was found to be pattern-compressible, this led to a violating disc. This in turn was used to simplify an earlier surface in the hierarchy. Hence, pattern-compressed squares never survive in the hierarchy. But what about the annuli that are used in step (2'')? What if these are compressed or pattern-compressed? We now explain how to deal with this situation.

Proposition 10.1 ((PATTERN-)COMPRESSIBLE VERTICAL BOUNDARY). *Let M be a compact orientable irreducible 3-manifold other than a 3-ball, and let P be a boundary pattern for M . Let $\mathcal{H}$ be an admissible handle structure for (M, P) . Let B be a component of the interior parallelity bundle $\mathcal{B}$ that is not an I -bundle over a disc, and let A be a vertical boundary component of B . Suppose that A has a pattern-compression or compression disc E . Then there is an algorithm that, given $\mathcal{H}$, A and E , either creates a handle structure for (M, P) with smaller q -complexity or produces a violating disc for (M, P) .*

Proof. Suppose first that E is a pattern-compression disc. Pattern-compress A along E to form a properly embedded disc D . If this is a violating disc, then the algorithm produces this disc and stops. So, suppose that D is not violating. Its boundary therefore bounds a disc D' in ∂M that intersects P either in the empty set or a single arc. Since M is irreducible, $D \cup D'$ bounds a 3-ball in M . If the arc $D \cap \partial M$ is disjoint from D' , then we deduce that A is parallel to an annulus A' in ∂M that intersects P either in the empty set or a single core curve. On the other hand, if $D \cap \partial M$ lies in D' , then we deduce that A is compressible. Indeed one of the discs of $D' \setminus \partial A$ forms a compression disc for A when pushed a little into the interior of M . We will later examine the case where A is compressible, and so we now suppose that A is parallel to the annulus A' .

Consider first the case where the solid torus between A and A' contains B . Then we remove this solid torus from the handle structure, thereby removing this component of $\mathcal{B}$ and possibly others. If A' contained a core curve of P , then we transfer this across to A . This has the effect of replacing some components of $\mathcal{F}^0(\mathcal{H})$ and $\mathcal{F}^1(\mathcal{H})$ by an arc of P , which does not increase the q -complexity of $\mathcal{H}$. Moreover, the solid torus cannot be an entire component of B , and so it must contain a 0-handle of $\mathcal{H}$ of positive index. So removing this solid torus strictly reduced $c_q(\mathcal{H})$.

So, suppose that the interior of B is disjoint from this solid torus. If B is an I -bundle over a disc, then this implies that M is a 3-ball which is contrary to one of the hypotheses of the proposition. Suppose that B is an I -bundle over a surface F with a single boundary component other than a disc. This then contains a non-separating essential arc. The I -bundle over this arc, together with two meridian discs for the solid torus, forms a violating disc for (M, P) . Suppose now that B is an I -bundle over a surface with more than one boundary component. Pick an arc in the base surface with exactly one endpoint on the boundary component that this is the image of A . Let A'' be the vertical boundary component incident to the other endpoint of this arc. The I -bundle over this arc, together with a meridian disc for the solid torus, gives a pattern-compression disc for A'' . If we pattern-compress along this disc and we create a violating disc, then the algorithm outputs this. If it does not create a violating disc, then again there are two cases. It might lead to a compression disc for A'' , which we will consider later. If not, then A'' is parallel to an annulus in ∂M intersecting P either in the empty set or a single core curve. But now the region between A'' and this annulus contains B . Hence, we may remove this solid torus, thereby removing B from the handle structure.

We now consider the case where A has a compression disc E . This compresses D to two discs D_1 and D_2 . If one of ∂D_1 or ∂D_2 does not bound a disc in ∂M , then D_1 or D_2 is violating, and the algorithm outputs this disc. So we may suppose that D_1 and D_2 are parallel to discs D'_1 and D'_2 in ∂M . It may be that D'_1 and D'_2 are nested, in which case suppose that $D'_1 \subset D'_2$. It may also be the case that D'_1 contains boundary components of $\partial_h B$ other than $\partial D'_1$. But by changing our choice of A if necessary, we may assume that this is not the case. If one of D'_1 or D'_2 has non-empty intersection with the pattern, it forms a violating disc. So, suppose that D'_1 and D'_2 are disjoint from the pattern.

The first case we consider is where D'_1 and D'_2 are disjoint. Then $D'_1 \cup A \cup D'_2$ bounds a ball. This ball does not contain B because $\partial_h B$ has no boundary components in D'_1 , and B is not an I -bundle over a disc. We may modify the

handle structure in the ball by declaring that it is a single interior parallelity 0-handle. This reduces the q -complexity of the handle structure.

Suppose now that D'_1 and D'_2 are nested. Then D'_1 lies within D'_2 , by our earlier assumption. The fact that B is not an I -bundle over a disc imply that the interior of $\partial D'_1$ is disjoint from B . Hence, B lies in the region between A and $D'_2 \setminus D'_1$. We remove this region from the handle structure. This reduces the q -complexity of the handle structure. $\square$

Thus, the algorithm in Section 6 is also modified as follows.

Further modification to the algorithm: In Step (13), we again consider a violating disc D . This forms a pattern-compression disc for some surface S_j in the hierarchy. If S_j is a square, then this creates a new violating disc. If S_j is the vertical boundary of the interior parallelity bundle for $\mathcal{H}_j$, we apply the algorithm in Proposition 10.1. This either produces a violating disc or it simplifies $\mathcal{H}_j$. In all other cases, we apply Lemma 3.6 to create a non-separating connected orientable surface $\overline{S_j}$ with smaller pattern complexity. We discard the rest of the hierarchy, and then proceed to Step (14).

Proposition 10.2 (The refined algorithm works). *Let M be a compact orientable irreducible 3-manifold, and let P be a boundary pattern. The above algorithm terminates with the correct declaration either that P is essential or that P is inessential. Moreover, the length of any of the partial hierarchies produced by the above algorithm is at most $4c_q(\mathcal{H})$.*

Proof. As in the proof of Theorem 6.1, it is clear that if the algorithm terminates, then its output is correct.

We start by bounding the length of the partial hierarchies that the algorithm produces. Therefore consider one such partial hierarchy $S_1, \dots, S_\ell$, and let $\mathcal{H}_i$ be the handle structure on for the manifold (M_i, P_i) containing S_i . We will show that $c_q(\mathcal{H}_{i+4}) < c_q(\mathcal{H}_i)$ for each i . This will establish that $\ell \leq 4c_q(\mathcal{H})$.

Certainly, $c_q(\mathcal{H}_{i+1}) \leq c_q(\mathcal{H}_i)$ by Proposition 5.9 and Lemma 5.11. We must show that at least every 4 steps, this quantity strictly decreases. To establish this, we examine carefully how the algorithm works.

The first thing that the algorithm does is to make the initial handle structure admissible, using CREATE ADMISSIBLE HANDLE STRUCTURE in Proposition 5.4. It then modifies the handle structure to ensure that the interior parallelity bundle is just I -bundles over discs, or it finds a violating disc.

The algorithm then proceeds by iterating the hierarchy construction loop. At each stage through the loop, it cuts along a connected non-separating normal surface, then possibly cuts along non-trivial squares to make the handle structure admissible, then cuts along vertical boundary components of the interior parallelity bundle and then vertical squares. After this stage, the interior parallelity bundle again consists of I -bundles over discs.

If at any stage, the algorithm leaves the hierarchy construction loop and does not immediately terminate, it finds a compression disc or pattern-compression disc for some S_j . When this S_j is a connected non-separating normal surface, a component of the resulting surface is non-separating and it is normalised, and we work instead with that surface. On the other hand, when S_j is a square or the vertical boundary components of the interior parallelity bundle, then S_j is either used to produce a further violating disc or it is used to simplify the handle structure; in both cases,

it is discarded. Hence, in the list of surfaces $S_1, \dots, S_\ell$, either our given S_i is connected, normal and non-separating, or this is true for some S_j with $j \leq i + 3$. Therefore, $c_q(\mathcal{H}_{i+4}) < c_q(\mathcal{H}_i)$, as claimed.

We now show that the algorithm terminates. At the i th stage, let $(S_1^i, \dots, S_{\ell(i)}^i)$ be the partial hierarchy that is currently stored. As in the proof of Theorem 6.1, the surface S_1^i must eventually stabilise at some surface S_1 . This is because each time that S_1^i is modified, either $\chi_p(S_1^i)$ decreases (when a pattern-compression or pattern compression is performed) or the q -complexity of the handle structure containing S_1^i is reduced (as in (PATTERN-)COMPRESSIBLE VERTICAL BOUNDARY). We then consider the next surfaces S_2^i . Again they eventually stabilise at some surface S_2 . This is for the same reason as above, plus the observation that when S_2^i is used to create a violating disc, then either this leads to the algorithm terminating or it is used to reduce the pattern-complexity of a previous surface. However, the latter case does not arise because we are assuming that S_1^i does not change any further. Repeating in this way, we create an infinite sequence of surfaces $S_1, S_2, \dots$. However, this is impossible because we have just established an upper bound on the length of any partial hierarchy created by the algorithm. $\square$

11. HANDLE STRUCTURES OF UNIFORM TYPE

Our algorithm starts with a handle structure $\mathcal{H}$. We then decompose this handle structure along a sequence of surfaces. Suppose that each of the 0-handles of the initial handle structure $\mathcal{H}$ falls into one of finitely many ‘types’. For example, if $\mathcal{H}$ had been dual to an ideal triangulation, then for each 0-handle H , its intersection with $\mathcal{F}$ is a thickened version of the complete graph on 4 vertices. If $\mathcal{H}$ had been dual to a triangulation, then for each 0-handle H , $H \cap \mathcal{F}$ is a thickened version of a subgraph of the complete graph on 4 vertices. We want to maintain this local finiteness as we proceed down the hierarchy. In this section, we show how to achieve this.

Definition 11.1. Let H and H' be 0-handles in handle structures $\mathcal{H}$ and $\mathcal{H}'$ for (M, P) and (M', P') . We say that they have the same *type* if there is a homeomorphism $H \rightarrow H'$ sending $H \cap \mathcal{F}^0(\mathcal{H})$ homeomorphically to $H' \cap \mathcal{F}^0(\mathcal{H}')$, sending homeomorphically $H \cap \mathcal{F}^1(\mathcal{H})$ to $H' \cap \mathcal{F}^1(\mathcal{H}')$, and sending $H \cap P$ homeomorphically to $H' \cap P'$.

Definition 11.2. We say that a handle structure $\mathcal{H}$ has *uniform type* if there is a finite list of handle types such that any 0-handle of $\mathcal{H}$ is one of these types.

Of course, this definition only really makes sense if the list of handle types is given in advance, and then the requirement that $\mathcal{H}$ is of uniform type forces its 0-handles to conform to this list of types. But this slightly imprecise wording is useful, because it will be convenient occasionally to enlarge our finite list of types. An example is given below in Theorem 11.3.

It will be important that we only use handle structures of uniform type. This is for many reasons. For example, if we wish to write down a normal surface, we will typically do so using its normal vector, and it is therefore important that we can list the different types of normal disc that can arise within a handle.

The following result ensures that the handle structures in our partial hierarchies have uniform type.

Theorem 11.3 (Maintaining uniform type). *Let $\mathcal{H}$ be a handle structure of uniform type for a compact orientable irreducible 3-manifold M with boundary pattern P . Let $\mathcal{H}'$ be the handle structure of a 3-manifold M' with boundary pattern P' obtained from $\mathcal{H}$ by a finite sequence of the following moves:*

- (1) *decomposition along a normal surface that is disjoint from the interior parallel bundle, no component of which is a boundary-parallel disc, and applied to an admissible handle structure;*
- (2) *application of CREATE ADMISSIBLE HANDLE STRUCTURE in Proposition 5.4;*
- (3) *decomposition along vertical boundary components of the interior parallel bundle;*
- (4) *application of (PATTERN-)COMPRESSIBLE VERTICAL BOUNDARY in Proposition 10.1.*

Then $\mathcal{H}'$ is of uniform type, although possibly with a larger set of handle types. Moreover, when $\mathcal{H}'$ is admissible and its interior parallel bundle consists of I -bundles over discs, then for each 0-handle H of $\mathcal{H}$, there are only finitely many possibilities for the following subsets:

- (1) *the intersection between H and the 0-handles of $\mathcal{H}'$;*
- (2) *$H \cap \mathcal{F}^0(\mathcal{H}')$;*
- (3) *$H \cap \mathcal{F}^1(\mathcal{H}')$;*
- (4) *$H \cap P'$,*

up to the following operations:

- (1) *isotopy of H preserving $H \cap \mathcal{H}^1$, $H \cap \mathcal{H}^2$ and $H \cap P$;*
- (2) *removing a small collar neighbourhood of a subset of $\partial M'$;*
- (3) *re-embedding $\mathcal{H}' \cap H$ in H in a way that leaves $\mathcal{F}^0(\mathcal{H}')$, $\mathcal{F}^1(\mathcal{H}')$ and $P' \cap \partial \mathcal{F}(\mathcal{H}')$ unchanged.*

Remark 11.4. Note that we are only claiming fairly limited control over the location of the 2-handles of $\mathcal{H}'$. We are specifying their attaching locus via $H \cap \mathcal{F}^0(\mathcal{H}')$ and $H \cap \mathcal{F}^1(\mathcal{H}')$, but we do not specify where the remainder of the 2-handles lie.

The re-embedding of $\mathcal{H}' \cap H$ might require some further explanation. Suppose that we had applied a sequence of modifications to $\mathcal{H}$, giving the handle structure $\mathcal{H}'$, and then suppose that we made some further modifications to give another handle structure $\mathcal{H}''$. Suppose that $\mathcal{H}' \cap H$ and $\mathcal{H}'' \cap H$ differ solely by a Dehn twist about a disc properly embedded in M' and lying in $\mathcal{H}' \cap H$. Then $\mathcal{H}' \cap H$ and $\mathcal{H}'' \cap H$ would differ by operation (3) above, and so, for the purposes of this proposition, would be viewed as equivalent.

Proof. We define a temporary notion of complexity for a 0-handle H of a handle structure, called *t-complexity*. It will be a list of three integers. The first element of the triple is the sum of the indices of the components of $H \cap \mathcal{F}$ with positive index. The second element is $|H \cap \mathcal{F}^1|$. The third element is $|H \cap \mathcal{F}^0|$. We compare two such triples in the usual way, using lexicographical ordering. The set of triples is then well-ordered.

We will prove by induction on t-complexity that a 0-handle type H gives rise to only finitely many handle types of $\mathcal{H}' \cap H$ under a sequence of the operations in (1)-(4) above. Moreover, we will show that the union of the 0-handles of $\mathcal{H}' \cap H$ with positive index embed in H in one of only finitely many ways, up to the equivalence

described in (1)-(3) at the end of the proposition. This will also establish the second part of the proposition, because when $\mathcal{H}'$ is admissible and its interior parallelity bundle consists of I -bundles over discs, then all its 0-handles have positive index.

The induction starts with H having t -complexity $(0, 0, 0)$. Then H is an isolated 0-handle and none of the procedures in (1)-(4) affect it.

We now prove the inductive step. We consider when one of procedures (1)-(4) is applied to a 0-handle H . We may assume that the result is not a single 0-handle H' that is the same type as H and where H' is obtained from H by removing a small collar on a subset of ∂H , as in this case there is nothing to prove.

In the procedure `CREATE ADMISSIBLE HANDLE STRUCTURE`, there are three types of operation that are performed. One type is where we collapse a 1-handle and 2-handle pair, where the 1-handle has at most two arcs of P running along it. The effect on $H \cap \mathcal{F}^0$ is possibly to remove some components of $H \cap \mathcal{F}^0$ with non-positive index. The effect on $H \cap \mathcal{F}^1$ is to remove some components and replace them by at most two parallel arcs of $P \cap H$. This does not increase the index of any component of $H \cap \mathcal{F}$. If H is modified, then it reduces at least one of $|H \cap \mathcal{F}^1|$ and $|H \cap \mathcal{F}^0|$, and it does not increase either of them. There are only finitely such modifications that can be applied to H . Hence, the induction is proved in this case.

The second type of operation in `CREATE ADMISSIBLE HANDLE STRUCTURE` is the removal of isolated 0-handles of $\mathcal{F}$ with non-positive index. Each such operation does not affect the index of the other components of $\mathcal{F} \cap H$. It leaves $|H \cap \mathcal{F}^1|$ unchanged and it reduces $|H \cap \mathcal{F}^0|$. Hence, again the inductive step is proved.

The third type of operation in `CREATE ADMISSIBLE HANDLE STRUCTURE` is when there is a simple closed curve C in $\partial H \setminus \mathcal{F}$ that intersects the pattern at most three times. It is assumed that C does not bound a disc in ∂H but that it does bound a disc in ∂M intersecting P in the empty set, an arc or a tripod. When $\partial H \setminus (\mathcal{F} \cup P)$ is not a union of discs, then C is chosen to be boundary parallel in a component of $\partial H \setminus (\mathcal{F} \cup P)$ that is not a disc. Hence, in that case, there are only finitely many possibilities for C . On the other hand, when $\partial H \setminus (\mathcal{F} \cup P)$ is a union of discs, then again there are only finitely many possibilities for C . Now C bounds a properly embedded disc D in H , and the operation that is performed is to remove the 3-ball separated off by D , and possibly to insert into D an arc or tripod. The effect of this on H is to remove some components of $\mathcal{F}$ and to possibly insert some new pattern. This does not increase the index of the components of $H \cap \mathcal{F}$, it does not increase $|H \cap \mathcal{F}^1|$, and it strictly reduces $|H \cap \mathcal{F}^0|$. Hence again the inductive step is proved.

The effect when we decompose along vertical boundary components of the interior parallelity bundle is similar to the case of `CREATE ADMISSIBLE HANDLE STRUCTURE`. The effect is to replace some components of $H \cap \mathcal{F}^1$ by two parallel arcs of $H \cap P$. This leaves index unchanged, but reduces $|H \cap \mathcal{F}^1|$. It also removes some components of $H \cap \mathcal{F}^0$ that are disjoint from P and that intersect $\mathcal{F}^1$ in two arcs, and replaces them by two arcs of P . Again, this leaves index unchanged, and reduces $|H \cap \mathcal{F}^0|$.

In `(PATTERN-)COMPRESSIBLE VERTICAL BOUNDARY`, submanifolds of the manifold are removed. The effect on H is either to remove it entirely (in which case there is nothing to prove) or possibly to remove some components of $\mathcal{F}^1$ and replace them with an arc of P , and possibly to remove some components of $\mathcal{F}^0$ that intersect $\mathcal{F}^1$ in two arcs and are disjoint from P , and replace them with an arc of

P . As argued above, this does not increase the index of any component of $\mathcal{F}$ or $|H \cap \mathcal{F}^0|$, and it strictly reduces $|H \cap \mathcal{F}^1|$.

We now examine what happens when H is decomposed along a normal surface. We are assuming that the handle structure is admissible in this case. Hence, every 0-handle of $H \cap \mathcal{F}^0$ has non-negative index. So, H has non-negative index. Furthermore, if it has zero index, then it is an interior parallelity handle, and we are assuming that the normal surface is disjoint from the interior parallelity bundle. So, we may assume that the index of H is positive.

The normal surface intersects H in discs. We can decompose along these one at a time. So we focus on one such disc D . Note that there are only finitely normal disc types in H . This disc D divides H into two 0-handles H_0 and H_1 , both of which have non-negative index, by Proposition 5.5. The sum of these indices is $I(H)$. Hence, if they are both positive, then H_0 and H_1 have smaller index, and so we are done by induction in this case. So suppose that one 0-handle H_0 has zero index and the other handle H_1 has positive index. Since D is normal, each component of $\mathcal{F}^1 \cap H$ is either divided into two by D or lies wholly in one of the handles H_0 or H_1 . If one component of $\mathcal{F}^1 \cap H$ ends up wholly in H_0 , then H_1 has smaller t-complexity than H , and we are again done by induction.

If ∂D lies wholly in $\mathcal{F} \cap H$, then the component of $H_0 \cap \mathcal{F}$ incident to ∂D has zero index and so must be an annulus disjoint from the pattern. Hence, in this case, H_1 is a copy of H and, given H , there are only finitely many possible types for H_0 .

Suppose that ∂D is disjoint from $\mathcal{F} \cap H$. It is therefore properly embedded in M . Since we are assuming that the decomposing surface has no component that is a boundary-parallel disc, D must separate components of $\mathcal{F} \cap H$. Hence, in this case, H_1 has smaller index than H , which is contrary to assumption.

So suppose that ∂D intersects $\mathcal{F} \cap H$ in arcs. These must be boundary-parallel in $\mathcal{F}$.

Let D' be the component of $\partial H \setminus \partial D$ lying in H_0 . This intersects $\mathcal{F}$ in n components, say, each of which is a disc intersecting P twice. Let Γ be the graph $H_0 \cap P$. This is a graph where each vertex has degree 1 or 3. The vertices with degree 1 occur at the points of $P \cap \mathcal{F}$ in H_0 , and also at the points $\partial D \cap P$. Let v_1 and v_3 be the number of vertices with degree 1 and 3 respectively.




FIGURE 14. Left: the disc D' and its intersection with P and $\mathcal{F}$. Middle: a tree component of Γ disjoint from $\mathcal{F}$. Right: A tree component of Γ intersecting $\mathcal{F}$ at a single point

We claim that no component Γ' of Γ is a tree disjoint from $\mathcal{F}$ (see Figure 14 middle). A regular neighbourhood of Γ' in D' is a disc, and the part of the boundary of this disc not lying in $\partial D'$ is a collection of arcs. There must be at least one such

arc α that is not outermost in D' , in the sense that it does not separate off a sub-disc of D' disjoint from the other arcs and from $\mathcal{F}$. At its endpoints are two arcs of $\partial D \setminus (P \cup \mathcal{F})$ lying in the same component of $\partial H \setminus (P \cup \mathcal{F})$. By the normality of D , these are therefore the same arc of $\partial D \setminus (P \cup \mathcal{F})$. Hence, α was outermost which is a contradiction.

We now show that no component of Γ is a tree intersecting $\mathcal{F}$ in a single point (see Figure 14 right). Suppose that Γ' is such a component. Let β be the arc of $\partial \mathcal{F} \setminus P$ incident to Γ' and that is disjoint from ∂D . Let α be an arc in D' disjoint from Γ running from β to ∂D that runs close to Γ' . We may extend α along $\partial \mathcal{F}$ in two ways to form arcs α_1 and α_2 properly embedded in D' , each intersecting P once. Let α_1 be the arc that intersects Γ' . The endpoints of α_1 lie in components of $\partial H \setminus (\mathcal{F} \cup P)$ that are incident along an arc of $\partial H \cap P$. So, by (6) in the definition of normality, ∂D intersects this arc. By (4), ∂D must run straight between these two endpoints of α_1 , intersecting P and nothing else. We deduce that Γ' is a single edge. A similar argument applied to α_2 gives that it is also parallel to a sub-arc of ∂D . So, $D' \cap P$ consists of two arcs. These must emanate from the same 0-handle of $\mathcal{F}$, as otherwise ∂D violates (5) in the definition of normality. So, we deduce that D contradicts (7) in the definition of normality.

We deduce that each component Γ' of Γ satisfies $2\chi(\Gamma') - |\mathcal{F} \cap \Gamma'| \leq 0$. So, $v_1 - v_3 - |\mathcal{F} \cap \Gamma| = 2\chi(\Gamma) - |\mathcal{F} \cap \Gamma| \leq 0$. Now the number of vertices of $P \cap H_1$ is equal to the number of vertices of $P \cap H$, plus $|\Gamma \cap \partial D|$, minus v_3 . But $|\Gamma \cap \partial D| = v_1 - |\mathcal{F} \cap \Gamma|$. We deduce that the number of vertices of the graph $P \cap H_1$ is at most the number of vertices of $P \cap H$.

To summarise, we have shown that in each of the operations (1)-(4), the t-complexity of H strictly goes down, except possibly in the case of (1). When the t-complexity does not go down, H is decomposed along the normal disc D , the new handle H_1 has $H_1 \cap \mathcal{F}$ an exact copy of $H \cap \mathcal{F}$, and the new pattern P' intersects $H_1 \cap \mathcal{F}$ in the same location that P intersected $H \cap \mathcal{F}$. More precisely, there is a homeomorphism $H \rightarrow H_1$ taking $\mathcal{F}^0(H)$ homeomorphically to $\mathcal{F}^0(H_1)$, taking $\mathcal{F}^1(H)$ homeomorphically to $\mathcal{F}^1(H_1)$, and taking $P \cap \mathcal{F}(H)$ homeomorphically to $P' \cap \mathcal{F}(H_1)$. This homeomorphism need not take $P \cap H$ homeomorphically to $P' \cap H_1$, but we have shown that these graphs have the same number of vertices. Hence, after a sequence of steps (1)-(4), the t-complexity must go down at some stage, or the handles run through only finitely many types. In either case, the argument is complete. $\square$

Remark 11.5. The above proposition is crucial to the efficient certification that a boundary pattern P for a compact orientable 3-manifold M is essential, as required by Theorem 1.6. When M is irreducible, this is achieved by giving an essential hierarchy for (M, P) . Now (M, P) is given to us by means of a triangulation $\mathcal{T}$, and so Theorem 11.3 guarantees that the resulting handle structures for all the manifolds in the hierarchy are of uniform type. This is important at many points. For example, in Proposition 12.1 below, a constant k depends on the type of the handles. Hence, for all the handle structures that we consider, k is in fact bounded above by a universal constant independent of M and the hierarchy.

Another crucial conclusion from Theorem 11.3 is that when the handle structure $\mathcal{H}'$ for a manifold in our hierarchy is admissible and its interior parallelity bundle consists of I -bundles over discs, then the intersection between the 0-handles of $\mathcal{H}'$ and any tetrahedron of $\mathcal{T}$ is one of only finitely many types, where again this list of

types is universal. In particular, the number of 0-handles of $\mathcal{H}'$ is bounded above by a universal constant times the number of tetrahedra of $\mathcal{T}$. This is important because many quantities depend on $|\mathcal{H}'|$, for example again in Proposition 12.1. It is also central to the efficiency of our algorithm. For example, when a manifold in the hierarchy is created and all its boundary components are 2-spheres, then we need to examine whether the boundary pattern in these 2-sphere is essential. We can do so efficiently, since the handle structure on this manifold has uniform type and its number of handles is bounded above by a universal constant times $|\mathcal{T}|$.

Theorem 11.3 is an analogue of results known for sutured manifold hierarchies (see [13, Proposition 6.3] and [16, Theorem 6.1]).

12. FUNDAMENTAL DECOMPOSING SURFACES

Our goal in this section is to control the extended weight of our decomposing surfaces.

Proposition 12.1 (Extended weight bound). *Let M be a compact orientable irreducible 3-manifold, other than a 3-ball, with non-empty boundary and an essential boundary pattern P . Let $\mathcal{H}$ be a handle structure of (M, P) of uniform type. Then there is a constant k (depending only on the handle types of $\mathcal{H}$) with the following property. There is a compact orientable connected non-separating normal surface S properly embedded in M with smallest pattern complexity among all such surfaces and with extended weight at most $k^{|\mathcal{H}|}$.*

Our proof of this fact relies heavily on work of Tollefson and Wang [21]. A fairly minor complication is that they worked with triangulations rather than handle structures. Starting from $\mathcal{H}$, it is straightforward to build a triangulation $\mathcal{T}$ of M in which P is simplicial, satisfying the following conditions:

- (i) the number $|\mathcal{T}|$ of tetrahedra is at most $c_1|\mathcal{H}|$, where c_1 is a constant depending only on the handle types of $\mathcal{H}$;
- (ii) any incompressible, pattern-incompressible normal surface S in $\mathcal{T}$ can be isotoped, leaving P invariant, to form a standard surface in $\mathcal{H}$ with extended weight at most $c_2 w(S)$, where $w(S)$ is the weight of S in $\mathcal{T}$ and c_2 is a constant depending only on the handle types of $\mathcal{H}$.

Hence, we now work with a triangulation $\mathcal{T}$ satisfying these conditions.

We now recall the projective normal solution space associated with $\mathcal{T}$. There are $7|\mathcal{T}|$ types of normal triangles and squares in $\mathcal{T}$. We assign a variable x_i to each type. When S is a normal surface, these variables count the number of triangles and squares of each type in S , giving a list of $7|\mathcal{T}|$ non-negative integers for S called its *normal surface vector* and denoted (S) . Such vectors satisfy a system of linear constraints called the *matching equations*. There are three such equations for each face of $\mathcal{T}$ with tetrahedra on both sides. For each such face, the triangles and squares in the tetrahedron on one side of the face must patch together correctly with the triangles and squares in the tetrahedron on the other side. The triangles and squares on one side intersect the face in normal arcs, which come in three different types. The matching equations assert that for each such normal arc type, the number of triangles and squares on one side intersecting the face in arcs of that form is equal to the number of triangles and squares on the other side also intersecting that face in arcs of that form. The *projective normal solution space* is the subset of $\mathbb{R}^{7|\mathcal{T}|}$ subject to the matching equations, the conditions $x_i \geq 0$ for each

i , and the equation $\sum x_i = 1$. It is a convex compact polyhedron. Each normal surface S determines a point in the projective normal solution space by scaling all its co-ordinates x_i so that $\sum x_i = 1$. The unique closed face of minimal dimension in the projective normal solution space containing this point is called the face that *carries* S and is denoted C_S .

A normal surface is said to be the *normal sum* of two properly embedded normal surfaces G and H if its normal vector is $(G) + (H)$. We denote the surface by $G + H$. One cannot always form the normal sum of two normal surfaces, but this is possible if they are both carried by the same face of the projective normal solution space. A normal surface is said to be *fundamental* if it cannot be written as the normal sum of two non-empty normal surfaces.

Definition 12.2. A compact oriented incompressible pattern-incompressible surface S properly embedded in (M, P) is said to be *p-taut* if the following all hold:

- (1) $[S, \partial S]$ is non-trivial in $H_2(M, \partial M)$;
- (2) S minimises pattern-complexity in its homology class;
- (3) there is no homologically trivial union of components of S .

We say that S is *lw-p-taut* if it is p-taut and it has minimal weight among all p-taut surfaces representing $[S, \partial S]$.

This is an analogue of a definition of Tollefson and Wang [21, Section 1]. Instead of pattern-complexity, their surfaces minimised χ_- , which is defined to be the sum, over all components S' of S , of $\max\{0, -\chi(S')\}$. Their surfaces were known as lw-taut.

The following is an analogue of Theorem 3.3 of [21], but with their lw-taut surfaces replaced by lw-p-taut surfaces. The proof is essentially the same, and is omitted.

Theorem 12.3. *Let M be a compact orientable irreducible 3-manifold with $b_1(M) > 0$ and let P be an essential boundary pattern. Let $\mathcal{T}$ be a triangulation of M in which P is simplicial. Let S be a compact oriented incompressible pattern-incompressible properly embedded surface that is lw-p-taut. Let C_S be the face of the projective normal solution space that carries S . Then every surface G carried by C_S inherits an orientation and so represents an element $[G] \in H_2(M, \partial M)$. With this orientation, G is lw-p-taut. Moreover, for any two surfaces G and H carried by C_S , the normal sum $G + H$ satisfies $[G + H] = [G] + [H]$.*

Hence, we have the following result.

Theorem 12.4 (Fundamental surface). *Let M be a compact orientable irreducible 3-manifold with $b_1(M) > 0$ and let P be an essential boundary pattern. Let $\mathcal{T}$ be a triangulation of M in which P is simplicial. Then there is a fundamental normal surface S properly embedded in M satisfying the following:*

- (1) S is a compact connected orientable surface that is non-separating in M ;
- (2) S has smallest pattern-complexity among all surfaces satisfying (1).

Proof. Pick a compact orientable surface S satisfying (1) and (2) above. Such a surface exists because $b_1(M) > 0$. We may assume that it is incompressible and pattern-incompressible, by Corollary 3.7. Hence it can be isotoped, via an isotopy preserving the pattern, to a normal surface, also called S . We may further assume that S has least weight among all such surfaces. So it is lw-p-taut. Let C_S be the

face that carries S . Suppose that S is not fundamental. Then S can be written as a non-trivial normal sum $G + H$ of normal surfaces G and H . We may assume that G is connected, since if G_1 is any component of G , then $S = G_1 + ((G - G_1) + H)$. Both G and H lie in C_S . By Theorem 12.3, G and H inherit orientations which make them lw-p-taut. In particular, G is non-separating in M . Neither G nor H has negative pattern-complexity, since we are assuming that P is essential. Note that pattern-complexity and weight are both additive under normal summation. Hence, the pattern-complexity of G is at most that of S . But it has smaller weight, contrary to our assumption. $\square$

Proof of Proposition 12.1. We are given a handle structure $\mathcal{H}$ of (M, P) of uniform type. From this, construct a triangulation $\mathcal{T}$ of M in which P is simplicial and satisfying (i) and (ii) above. By Theorem 12.4, there is a fundamental normal surface S properly embedded in M satisfying (1) and (2) in the theorem. By a result of Hass, Lagarias and Pippenger [5, Lemma 6.1], the weight of S in $\mathcal{T}$ is at most $|\mathcal{T}|^{2 \cdot 2^{7|\mathcal{T}|+7}}$. Hence, S is isotopic, leaving P invariant, to a standard surface S in $\mathcal{H}$ with extended weight at most $c_2 |\mathcal{T}|^{2 \cdot 2^{7|\mathcal{T}|+7}}$. Then by Lemma 4.11, S can be isotoped, leaving P invariant, to be normal in $\mathcal{H}$ and without increasing its extended weight. Since, $|\mathcal{T}| \leq c_1 |\mathcal{H}|$, we obtain our required constant k and exponential bound on the extended weight of S . $\square$

13. CERTIFYING AN INESSENTIAL PATTERN

In this section, we establish one part of Theorem 1.6, by showing that ESSENTIAL BOUNDARY PATTERN is in co-NP. In the case where the given 3-manifold is irreducible, this is a consequence of the following result, which we will prove in this section.

Theorem 13.1 (Fundamental violating disc). *Let M be a compact orientable irreducible 3-manifold with a boundary pattern P . Let $\mathcal{T}$ be a triangulation for M in which P is simplicial. If P is inessential, then it has a violating disc that is normal and fundamental.*

The following result provides a violating disc in an important special case.

Theorem 13.2 (Fundamental compression disc). *Let M be a compact orientable irreducible 3-manifold with pattern P . Let $\mathcal{T}$ be a triangulation for M in which P is simplicial. If $\partial M \setminus P$ is compressible, then it has a compression disc that is normal and fundamental.*

This result was proved in [18, Theorem 4.1.13]. Although that theorem does not explicitly assert that the compression disc is fundamental, this is established in the proof. So, in our proof of Theorem 13.1, we may assume that $\partial M \setminus P$ is incompressible. In the terminology of Matveev [18], this means that (M, P) is *boundary-irreducible*.

We recall some of the standard terminology from normal surface theory. Let S be a normal surface that is the normal sum of two surfaces S_1 and S_2 . We may pattern-isotope one of the surfaces, preserving its normality throughout, so that afterwards S_1 and S_2 intersect transversely in a collection of simple closed curves and properly embedded arcs. Then each component of $S_1 \setminus S_2$ and $S_2 \setminus S_1$ is a *patch*. It is a *clean disc patch* if it is a disc that is disjoint from the pattern and is either disjoint from ∂M or intersects ∂M in a single arc.

A summation $S = S_1 + S_2$ is in *reduced form* if there do not exist normal surfaces S'_1 and S'_2 where each S'_i is pattern-isotopic to S_i , and $S = S'_1 + S'_2$, with $|S'_1 \cap S'_2| < |S_1 \cap S_2|$.

Let S be a compact surface properly embedded in M . A *clean boundary-compression disc* for S is a disc embedded in M such that ∂D is the concatenation of two arcs α and β , where

- (1) $\alpha = D \cap S$ does not separate off a disc from S that is disjoint from P ;
- (2) $\beta = D \cap \partial M$ is disjoint from P .

When S has no clean boundary-compression disc, it is said to be *cleanly boundary-incompressible*.

Note that a violating disc that intersects P as few times as possible is cleanly boundary-incompressible.

When S is a normal sum of two normal surfaces S_1 and S_2 in general position, it is well known (see for example [18, Section 3.2.2]) that S is obtained from $S_1 \cup S_2$ by resolving the intersections $S_1 \cap S_2$ as follows. A regular neighbourhood of $S_1 \cap S_2$ is removed from S_1 and S_2 . This regular neighbourhood is a collection of discs, annuli and Möbius bands. At each component of $S_1 \cap S_2$, a pair of discs or annuli are inserted, lying in a regular neighbourhood of that component of $S_1 \cap S_2$. These are known as *trace* discs and annuli. This procedure is called a *switch*. There are two possible ways to perform a switch. Only one choice of switch at each component of $S_1 \cap S_2$ leads to the normal surface S . These are called *regular switches*. The other type of switch is an *irregular switch*. If a switch is made at each component of $S_1 \cap S_2$ and at least one of these switches is irregular, then the resulting surface S' admits a disc D in some face of $\mathcal{T}$, such that $D \cap S'$ is an arc in ∂D , and the remainder of ∂D lies in an edge of $\mathcal{T}$. When D is disjoint from ∂M , this leads to an isotopy that reduces the weight of S' . When D intersects ∂M but not P , then it is a potential clean boundary-compression disc for S' . If it is not a clean boundary-compression disc then, assuming $\partial M \setminus P$ is incompressible, there is an isotopy of S' leaving P fixed, that reduces its weight. On the other hand, when D intersects P , then a small isotopy of D pushing it off the face makes it a potential clean boundary-compression disc. There are two possible directions in which to push D off the face and if neither direction gives a clean boundary-compression disc, then, again assuming $\partial M \setminus P$ is incompressible, S' must have been a disc parallel to a disc in ∂M intersecting P in a single arc. This observation leads to the following result.

Proposition 13.3. *Let M be a compact orientable irreducible 3-manifold with pattern P . Let $\mathcal{T}$ be a triangulation for M in which P is simplicial. Suppose that $\partial M \setminus P$ is incompressible. Let S be a compact connected orientable surface that is incompressible and cleanly boundary-incompressible and that is not a disc parallel to a disc in ∂M intersecting P in a single arc or the empty set. Suppose that S is the sum of two normal surfaces S_1 and S_2 in general position, and that this summation is in reduced form. Assume also that S has least weight, up to isotopy that preserves the pattern. Then no component X of $M \setminus (S_1 \cup S_2)$ is a product region, in the following sense: there is a homeomorphism from X to the product of a surface and an interval, the homeomorphism taking $P \cap X$, $N(S_1 \cap S_2)$ and $\partial M \cap X$ to unions of interval fibres.*

Proof. Note that ∂X contains at least one component of $S_1 \cap S_2$, since otherwise S is not connected. Resolve the intersections between S_1 and S_2 that lie in ∂X in such a way that none of the trace discs and annuli that are inserted are vertical in X . If all of these switches are regular switches, then the summation was not reduced, for the following reason. We can remove $S_1 \cap X$ from S_1 , and then insert $S_2 \cap X$, forming a surface S'_1 . We can also form S'_2 by removing $S_2 \cap X$ from S_2 , and then inserting $S_1 \cap X$. Then $S = S'_1 + S'_2$, but $|S'_1 \cap S'_2| < |S_1 \cap S_2|$. On the other hand, if not all the switches are regular switches, then we can perform these switches and then perform regular switches at the remaining components of $S_1 \cap S_2$. The resulting surface S' is isotopic to S , via an isotopy preserving the pattern. Since we are assuming that S is incompressible and cleanly boundary-incompressible and not a disc parallel to a disc in the boundary intersecting P in arc, we deduce that S' admits an isotopy preserving the pattern that reduces its weight. This contradicts our minimality assumption. $\square$

The following is [18, Lemma 4.1.8].

Theorem 13.4. *Let M be a compact orientable irreducible 3-manifold with pattern P . Let $\mathcal{T}$ be a triangulation for M in which P is simplicial. Suppose that $\partial M \setminus P$ is incompressible. Let S be a compact orientable normal surface properly embedded in M , that is incompressible and cleanly boundary-incompressible and that has no component that is a disc parallel to a disc in ∂M intersecting P in arc or the empty set. Suppose that S is a sum of two normal surfaces S_1 and S_2 in general position and that this summation is reduced. Suppose also that S has minimal weight up to isotopy preserving the pattern. Then no patch of the summation is a clean disc patch.*

Proof of Theorem 13.1. In the case where $\partial M \setminus P$ is compressible, then Theorem 13.2 provides a fundamental compression disc. So we may assume that $\partial M \setminus P$ is incompressible.

Let D be a violating disc that intersects the pattern as few times as possible. Then D is cleanly boundary-incompressible and it is automatically incompressible. So, by [18, Corollary 3.3.25], D can be isotoped, preserving the pattern, into normal form without increasing its weight. We may assume that D has least weight among all normal violating discs that intersect the pattern as few times as possible.

Suppose that D is a normal sum $S_1 + S_2$. We may assume that this summation is in reduced form. By [18, Theorem 4.1.36], neither S_i has a component that is a sphere, projective plane, or a disc intersecting the pattern at most once. Hence, each S_i satisfies $-\chi(S_i) + |S_i \cap P| \geq 0$. Note that Euler characteristic and the number of intersection points with the pattern are both additive under normal summation. Since $\chi(D) > 0$, then one of S_1 or S_2 has positive Euler characteristic, say S_1 . Then S_1 has a disc component. It must intersect the pattern at least twice. We may assume that S_1 is connected, by adding any other components of S_1 to S_2 if necessary. Now, S_2 must therefore have zero Euler characteristic and intersect the pattern at most once. Since no component of S_2 is a disc intersecting the pattern at most once, we deduce that every component of S_2 has zero Euler characteristic.

Our goal is to show that S_1 is a violating disc. Since S_1 has smaller weight than D , this will contradict our earlier minimality assumption. Suppose that, on the contrary, S_1 is a parallel to a disc in ∂M that intersects P in a single arc or a tripod.

Now S_2 cannot have any closed components, since the intersection between such a component and S_1 would lead to a simple closed curve of $S_1 \cap S_2$ and an innermost such curve in S_1 would create a clean disc patch. So, S_2 consists of annuli and Möbius bands.

We claim that each arc of $S_1 \cap S_2$ is essential in S_2 . If not, there is an inessential arc that is outermost in S_2 . It therefore separates off a disc D_2 in S_2 . This disc intersects P at most once, since S_2 intersects P at most once. Also $D_2 \cap S_1$ separates off a sub-disc D_1 of S_1 intersecting P at most once. The union of these discs creates a properly embedded disc in M intersecting P fewer times than D did. Hence, this disc is parallel to a disc in ∂M intersecting P in the empty set or an arc. But then $D_2 \cup (S_1 \setminus D_1)$ is a violating disc that intersects P at most as many times as D does and that has smaller weight than D , contradicting our minimality assumption about D .

We claim that each arc of $S_1 \cap S_2$ is essential in S_2 . If not, there is an inessential arc that is outermost in S_2 . It therefore separates off a disc D_2 in S_2 . This disc intersects P at most once, since S_2 intersects P at most once. Also $D_2 \cap S_1$ separates off a sub-disc D_1 of S_1 intersecting P at most once. The union of these discs creates a properly embedded disc in M intersecting P fewer times than D did. Hence, this disc is parallel to a disc in ∂M intersecting P in the empty set or an arc. But then D_1 and D_2 cobound a product region as in Proposition 13.3, contradicting that proposition.

Hence, $S_2 \setminus S_1$ consists of squares. Each square has boundary consisting of two arcs in S_1 and two arcs in ∂M . At least one square is in the ball bounded by S_1 and every such square is boundary-parallel in the ball. Consider an outermost such square. It is parallel to a square in S_1 . But the resulting product region contradicts Proposition 13.3.

Thus, we have shown that S_1 and S_2 cannot intersect, which is impossible, because S is connected. So S must have been fundamental. $\square$

We now complete the proof of Theorem 1.6. We are given a triangulation $\mathcal{T}$ of a compact orientable 3-manifold M , with boundary pattern P that is simplicial. When P is inessential, we need to provide a certificate that confirms this and that can be verified in polynomial time as function of $t = |\mathcal{T}|$, the number of tetrahedra. When M is irreducible, Theorem 13.1 shows that there is a violating disc D that is fundamental. This is the required certificate. The algorithm of Agol, Hass and Thurston [1] can be used to verify that D is indeed a disc intersecting the pattern at most 3 times. Also, by [16, Theorem 9.4], one can check whether ∂D bounds a disc in ∂M . If there is no such disc, then D is violating. So suppose that ∂D does bound a disc D' in ∂M . It may be the case ∂D bounds two discs, in which case we consider them both in turn. Again using [16, Theorem 9.4], one can check whether each component of $D' \setminus P$ is a disc intersecting $\partial D'$ in one arc and with at most one vertex of P in its boundary. Since D is violating, at least one component of $D' \setminus P$ is not of this form, and the algorithm provided by [16, Theorem 9.4] will confirm this.

We need to deal with the general case where M is reducible. Our certificate for establishing that P is essential is:

- (1) a vector (S) of a (possibly empty) normal surface S in $\mathcal{T}$ with weight at most 2^{740t^2} ; this will in fact be a collection of disjoint embedded spheres;

- (2) a triangulation $\mathcal{T}'$ for a 3-manifold M' with at most $1000t$ tetrahedra; this will be $M \setminus S$ with a 3-ball attached to each spherical boundary component;
- (3) a certificate that M' has inessential boundary pattern.

We now explain why this certificate exists. It is proved in [9, Lemma 4] that there is a maximal collection of disjoint non-parallel normal 2-spheres S in $\mathcal{T}$, none of which bounds a ball in M , with weight at most 2^{740t^2} . A triangulation for $M \setminus S$ with at most $1000t$ tetrahedra is produced using [16, Theorem 11.4]. Then a coned 3-ball is attached to each spherical boundary component, forming the triangulation $\mathcal{T}'$ for M' . The copy of P in M' is essential if and only if P was essential in M , because any violating disc for P in M' can be chosen to avoid the attached 3-balls, and so forms a violating disc in M , and conversely, any violating disc in M can be chosen to avoid S and so becomes a violating disc in M' . Since we are assuming that P is inessential, there is a certificate for (M', P) that establishes this.

The verification of the certificate is straightforward. One checks that S is a collection of 2-spheres, using the Agol-Hass-Thurston algorithm. Then $\mathcal{T}'$ is produced using the algorithm in [16, Theorem 11.4]. Finally the certificate showing that P is inessential in M' is checked in the way described above.

14. CERTIFYING AN ESSENTIAL HIERARCHY

In this section, we complete the proof of Theorem 1.6, by showing that ESSENTIAL BOUNDARY PATTERN is in NP. Thus, we are given a triangulation $\mathcal{T}$ of a compact orientable 3-manifold M with essential pattern P that is simplicial, and we must provide a certificate, verifiable in polynomial time, that establishes that P is indeed essential. We may assume that ∂M is non-empty, as otherwise there is nothing to prove. Using the techniques at the end of the previous section, it is not hard to reduce to the case where M is irreducible. We will describe this reduction in more detail at the end of this section, but for the moment, suppose that M is irreducible. Then, by Theorem 2.5, we know that (M, P) admits an essential hierarchy. Moreover, the existence of such a hierarchy guarantees that P is essential. So, our goal is to write down such a hierarchy and prove that it is essential. We know by Proposition 10.2 that there is such a hierarchy with length that is linearly bounded in terms of the number of tetrahedra in $\mathcal{T}$. Furthermore, by Theorem 11.3, the intermediate 3-manifolds have handle structures of uniform type. However, this only controls the local complexity of each handle, not the total number of handles. Thus, it is not immediately clear how to encode the handle structures of these manifolds. Nevertheless, we will establish the following result.

Theorem 14.1 (Efficient encoding of a hierarchy). *Let M be a compact orientable irreducible 3-manifold with a boundary pattern P . Let $\mathcal{H}$ be a handle structure for (M, P) of uniform type. If (M, P) admits an essential hierarchy, then there is a hierarchy for (M, P) that may be described by an amount of data that is at most a polynomial function of $|\mathcal{H}|$ and this hierarchy can be verified as essential by a polynomial-time algorithm. Furthermore, for each decomposition*

$$(M_i, P_i) \xrightarrow{S_i} (M_{i+1}, P_{i+1})$$

in this hierarchy, one of the following holds:

- (1) S_i is connected, orientable and non-separating and has smallest pattern-complexity among all such surfaces, apart from when M is closed and $i = 1$, when S_1 may be separating;

- (2) S_i is a union of non-trivial squares;
- (3) S_i is a union of properly embedded incompressible pattern-incompressible annuli disjoint from the pattern.

Recall that the hierarchy construction loop given in Sections 6 and 10 used the following steps iteratively (where the numbering is as given there):

- (2) Apply CREATE ADMISSIBLE HANDLE STRUCTURE.
- (2') Replace each component of the interior parallelity bundle that is an I -bundle over a disc by a single 2-handle.
- (2'') Cut along the vertical boundary components of the remaining I -bundles and further decompose these using squares into 3-balls with essential boundary patterns.
- (6) Let S be any connected, non-separating, properly embedded, oriented normal surface in the resulting 3-manifold.
- (8) Assuming that S is not a violating disc, then cut the handle structure along S .

We can view the iterations of this loop as running in the order (6), (8), (2), (2'), (2''), after an initial application of (2), (2'), and (2''). The handle structures that are obtained by repeating this loop can be constructed using the following result.

Proposition 14.2. *Let $\mathcal{H}$ be a handle structure of uniform type for the compact orientable irreducible 3-manifold M with essential boundary pattern P . Let S be a connected properly embedded normal surface in M . Then there is an algorithm that produces an admissible handle structure that is obtained by applying steps (8), (2), (2'), (2''). This runs in polynomial time as a function of the number of handles of $\mathcal{H}$ and the logarithm of the extended weight of S .*

It is a consequence of Theorem 11.3 that the handle structure obtained by decomposing $\mathcal{H}$ along S has uniform type. However, it might have a large number of handles, when S has very high extended weight. But most handles will be of the following form.

Definition 14.3. Let $\mathcal{H}$ be a handle structure for the compact 3-manifold M with boundary pattern P . Let S be a normal surface properly embedded in M . Let $\mathcal{H}'$ be the induced handle structure on $M \setminus S$. Then a handle of $\mathcal{H}'$ is a *parallelity handle* for S if it lies between parallel normal discs of S .

Each parallelity handle for S is a copy of $D^2 \times I$ where $(D^2 \times I) \cap S = \partial D^2 \times I$. We view this product structure as an I -bundle over D^2 . The I -bundle structures can be chosen so they agree on the intersection of any collection of parallelity handles. Hence, the union of the parallelity handles is an I -bundle, called the *parallelity bundle* $\mathcal{B}$ for S .

The vertical boundary $\partial_v \mathcal{B}$ of the parallelity bundle $\mathcal{B}$ is a union of annuli. However, these annuli need not be properly embedded. Instead the intersection with ∂M is a union of fibres in the bundle. We say that $\partial_v \mathcal{B} \setminus \partial M$ is the *inner vertical boundary*. It is a disjoint union of properly embedded annuli disjoint from the pattern, and squares in $M \setminus S$.

The following result is analogous to [16, Theorems 9.2 and 9.3]. It has essentially the same proof, which is omitted. The proof relies heavily on the algorithm of Agol, Hass and Thurston [1].

Theorem 14.4 (CUTTING ALONG A SURFACE). *There is an algorithm that takes, as its input,*

- (1) *a handle structure $\mathcal{H}$ of uniform type, with h handles, for a compact orientable 3-manifold M with boundary pattern P ;*
- (2) *a compact orientable properly embedded normal surface S ;*

and provides as its output, the following data. If (M', P') is the manifold with pattern obtained by decomposing along S , and $\mathcal{B}$ is its parallelity bundle for S , then the algorithm produces the handle structure that $M' \setminus \mathcal{B}$ inherits, and for each component B of $\mathcal{B}$, it determines:

- (1) *the genus and number of boundary components of its base surface;*
- (2) *whether B is a product or twisted I -bundle;*
- (3) *the way that $\partial_v B$ and $M' \setminus \mathcal{B}$ intersect.*

The algorithm runs in time that is bounded by a polynomial in h and the logarithm of the extended weight of S .

Proof of Proposition 14.2. Let $\mathcal{H}'$ be the handle structure that $M \setminus S$ inherits. This might have too many handles to be written down efficiently. Suppose instead that we are provided by the output of Theorem 14.4. In Proposition 14.2, the next operation that is applied to $\mathcal{H}'$ is CREATE ADMISSIBLE HANDLE STRUCTURE. If $\mathcal{B}$ has any components that are incident to ∂M , then an implementation of CREATE ADMISSIBLE HANDLE STRUCTURE will remove them, for the following reason. Let B be such a component of $\mathcal{B}$ that is incident to ∂M . There must be a 0-handle H of $\mathcal{H}'$ in $\mathcal{B}$ that is incident to ∂M . Since H is a parallelity handle for S that is incident to ∂M , its intersection with $\mathcal{F}(\mathcal{H}')$ is a line of 0-handles and 1-handles. Suppose first that there is more than one 0-handle of $\mathcal{F}(\mathcal{H}') \cap H$. Then one such 0-handle D intersects P' in two points and intersects $\mathcal{F}^1(\mathcal{H}')$ in a single arc. Hence, $\mathcal{H}'$ fails to be admissible. We could perform the operation described in CREATE ADMISSIBLE HANDLE STRUCTURE, which removes from $\mathcal{H}'$ the 1-handle and 2-handle that are incident to D and transfers arcs of P' incident to these handle across them. This 1-handle and 2-handle were in $\mathcal{B}$, and we view the new parallelity bundle for S as obtained from $\mathcal{B}$ by removing these two handles. In this way, we could continue to shrink B until it contained no 2-handles. Then $B \cap \mathcal{F}(\mathcal{H}')$ would consist of squares. For each such square that is non-trivial, we could decompose along it. For each square D that is trivial, it is parallel to a disc D' in $\partial M'$ that intersects P' as shown in Figure 5. We can remove the 3-ball bounded by $D \cup D'$ and transfer the boundary pattern in D' across to D . However, there is an alternative approach which is considerably faster to implement, and which has the same end result. Let A be a component of $\partial_v B$ that is incident to ∂M . Pick a maximal collection of vertical squares in B with vertical boundary lying in A , none of which is parallel to subset of $A \cap \partial M$, and no two of which are parallel. If any such square is non-trivial, we decompose along it. If any component D is a trivial square, then we identify the disc D' as in Figure 5 with the same boundary, remove the ball between these discs and then transfer the boundary pattern in D' across to D . These operations decompose B either to a 3-ball or to a collar on $\partial_v B - A$. In the latter case, we can remove this collar and transfer its boundary pattern to $\partial_v B - A$. This process can all be completed without knowing the precise handle structure on B , but instead just the information provided by Theorem 14.4. We repeat this for each component of $\mathcal{B}$ that is incident to ∂M . The resulting handle structure need not

be admissible, in which case we perform further operations in `CREATE ADMISSIBLE HANDLE STRUCTURE`. Initially, this does not affect the interior parallelity bundle. However, at some point, it may do so. This happens when the vertical boundary of a remaining component of $\mathcal{B}$ fails to be properly embedded. In that case, we can again remove this component by decomposing along its interior vertical boundary. Each of these operations either removes a component of the parallelity bundle for S , or removes a handle of $M' \setminus \mathcal{B}$. Hence, it can be completed in polynomial time, as a function of the number of handles of $\mathcal{H}$. The resulting object is a handle structure $\mathcal{H}''$ to which various I -bundles are attached along their vertical boundary. The actual result of running `CREATE ADMISSIBLE HANDLE STRUCTURE` is obtained from this by replacing the I -bundles by a union of parallelity handles in some way that we do not specify. The interior parallelity bundle for this handle structure includes these I -bundles but it may also include various other handles of $\mathcal{H}''$. We can compute this interior parallelity bundle, and then the steps in (2') and (2'') can be completed in polynomial time. Thus we have the algorithm required by Proposition 14.2. $\square$

The certificate required by Theorem 1.6 is as follows:

- (1) a sequence of handle structures $\mathcal{H}_1, \dots, \mathcal{H}_{\ell+1}$, where $\mathcal{H}_i$ is a handle structure for a 3-manifold M_i with boundary pattern P_i ;
- (2) for some of the $i \leq \ell$, a surface S_i or D_i properly embedded in M_i that is normal in $\mathcal{H}_i$.

These handle structures and surfaces are required to have the following properties:

- (1) the initial handle structure $\mathcal{H}_1$ is dual to the given triangulation $\mathcal{T}$;
- (2) for each 0-handle H of $\mathcal{H}_1$, the intersection between H and the 0-handles of $\mathcal{H}_i$, $\mathcal{F}^0(\mathcal{H}_i)$, $\mathcal{F}^1(\mathcal{H}_i)$ and P_i is one of the possibilities from Theorem 11.3;
- (3) each S_i and each D_i is normal in $\mathcal{H}_i$ and has extended weight at most $k^{|\mathcal{H}_i|}$, where k is the constant from Proposition 12.1;
- (4) when S_i is provided, then $\mathcal{H}_{i+1}$ is obtained from $\mathcal{H}_i$ using Proposition 14.2;
- (5) when D_i is provided, then it is a compression disc or pattern-compression disc for the vertical boundary of the interior parallelity bundle for $\mathcal{H}_i$, and then $\mathcal{H}_{i+1}$ is obtained from $\mathcal{H}_i$ by simplifying the handle structure according to the procedure in `(PATTERN-)COMPRESSIBLE VERTICAL BOUNDARY`;
- (6) when no D_i or S_i is provided, then $\mathcal{H}_{i+1}$ is obtained from $\mathcal{H}_i$ by applying `CREATE ADMISSIBLE HANDLE STRUCTURE`, and furthermore the interior parallelity bundle for $\mathcal{H}_{i+1}$ consists of 2-handles;
- (7) the length ℓ of this hierarchy is at most $4c_q(\mathcal{H}_1)$;
- (8) the final manifold is a collection of 3-balls with essential pattern and its handle structure is admissible with interior parallelity bundle consisting of 2-handles.

The existence of the certificate when M is irreducible and P is essential is established as follows. We run the algorithm given in Section 6 with the refinements from Section 10. In Step (6) of the algorithm in Section 6, we are free to pick any connected non-separating properly embedded orientable normal surface in the manifold. We pick one with smallest pattern complexity. Then Proposition 12.1 guarantees that there is such a surface with extended weight at most $k^{|\mathcal{H}_i|}$. Furthermore, since the surface has smallest pattern complexity, Corollary 3.7 ensures that it is incompressible and pattern-incompressible. Hence, if the algorithm leaves

the hierarchy construction loop and produces a violating disc, then this must form a compression or pattern-compression disc for some vertical boundary annulus for the interior parallelity bundle. This is provided in (5). Note that such a violating disc can be chosen to be fundamental in a triangulation refining $\mathcal{H}_i$ by Theorem 13.1, and hence have extended weight at most $k^{|\mathcal{H}_i|}$ in $\mathcal{H}_i$. Since M is assumed to be irreducible and P is essential, the algorithm terminates with an essential hierarchy, giving (8). Furthermore, the length of this hierarchy is at most $4c_q(\mathcal{H}_1)$ by Proposition 10.2, which gives (7). Finally, condition (2) is a consequence of Theorem 11.3.

We now explain how to verify the certificate. There are two main aspects to this. Firstly, we must check that each $\mathcal{H}_{i+1}$ really is obtained from $\mathcal{H}_i$ in the ways described above. This guarantees that we really do have a partial hierarchy. Secondly, we must check that the final manifold $M_{\ell+1}$ is balls with essential boundary pattern $P_{\ell+1}$. Since the final handle structure is admissible with interior parallelity bundle consisting of 2-handles, Theorem 11.3 implies that its number of handles is at most a uniform constant times $|\mathcal{H}_1|$. Furthermore each of its 0-handles is of uniform type. Hence, it is straightforward to verify, in polynomial time, that $\partial M_{\ell+1}$ is a collection of spheres and that its pattern $P_{\ell+1}$ is essential. This implies that each decomposing surface is incompressible by Lemma 3.2, and $M_{\ell+1}$ is irreducible by Lemma 3.9. So $M_{\ell+1}$ is a collection of balls.

So we now focus on verifying that each $\mathcal{H}_{i+1}$ is obtained from $\mathcal{H}_i$ as above. Suppose first that the normal surface S_i is given. We simply apply Proposition 14.2 and check that the resulting handle structure is $\mathcal{H}_{i+1}$. Note that S_i has extended weight at most $k^{|\mathcal{H}_i|}$. Since $|\mathcal{H}_i|$ is bounded above by a universal constant times $|\mathcal{H}_1|$ by (2) above, we deduce that the algorithm in Proposition 14.2 completes in polynomial time.

Suppose now that the disc D_i is given in $\mathcal{H}_i$. This is a compression disc or pattern-compression disc for a vertical boundary component A_i of the interior parallelity bundle $\mathcal{B}_i$ for $\mathcal{H}_i$. We may remove $\mathcal{B}_i$ from $\mathcal{H}_i$, forming a new handle structure $\mathcal{H}'_i$ containing D_i . It is of uniform type and its number of 0-handles satisfies $|\mathcal{H}'_i| \leq |\mathcal{H}_i|$. So we may confirm that D_i is a compression disc or pattern-compression disc for A_i . Then the procedure in (PATTERN-)COMPRESSIBLE VERTICAL BOUNDARY may be implemented in polynomial time and the verifier may check that the resulting handle structure is $\mathcal{H}_{i+1}$.

Finally, when neither D_i nor S_i is given, then the procedure CREATE ADMISSIBLE HANDLE STRUCTURE may be implemented and the verifier can check that the resulting handle structure is $\mathcal{H}_{i+1}$.

Thus, the verifier can check that this does give an essential hierarchy for (M, P) . This completes the proof of Theorem 1.6 in the case where M is irreducible and P is essential.

We now complete the proof of Theorem 1.6. We need to deal with the general case where M is reducible. Our certificate for establishing that P is essential is:

- (1) a vector (S) of a normal surface S in $\mathcal{T}$ with weight at most 2^{740t^2} ; this will in fact be a collection of disjoint embedded spheres;
- (2) a triangulation $\mathcal{T}'$ for a 3-manifold M' with at most $1000t$ tetrahedra; this will be $M' \setminus S$ with a 3-ball attached to each spherical boundary component;
- (3) a certificate that M' has essential boundary pattern.

Parts (1) and (2) of the certificate are verified in the same way as in Section 13. One checks that S is a collection of 2-spheres, using the Agol-Hass-Thurston algorithm. A triangulation for $M \setminus S$ is produced using [16, Theorem 11.4], and then a coned 3-ball is attached to each spherical boundary component, forming T' . The copy of P in M' is essential if and only if P was essential in M , because any violating disc for P in M' can be chosen to avoid the attached 3-balls, and so forms a violating disc in M , and conversely, any violating disc in M can be chosen to avoid S and so becomes a violating disc in M' . Now M' is irreducible. So we have a certificate, verifiable in polynomial time, that (M', P) is essential.

REFERENCES

- [1] I. Agol, J. Hass, and W. Thurston. The computational complexity of knot genus and spanning area. *Trans. Amer. Math. Soc.*, 358(9):3821–3850, 2006.
- [2] I. A. Dynnikov. Arc-presentations of links: monotonic simplification. *Fund. Math.*, 190:29–76, 2006.
- [3] W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961.
- [4] W. Haken. Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I. *Math. Z.*, 80:89–120, 1962.
- [5] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999.
- [6] S. V. Ivanov. The computational complexity of basic decision problems in 3-dimensional topology. *Geom. Dedicata*, 131:1–26, 2008.
- [7] F. Jaeger, D. L. Vertigan, and D. J. A. Welsh. On the computational complexity of the Jones and Tutte polynomials. *Math. Proc. Cambridge Philos. Soc.*, 108(1):35–53, 1990.
- [8] Klaus Johannson. *Homotopy equivalences of 3-manifolds with boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer, Berlin, 1979.
- [9] Simon A. King. The size of triangulations supporting a given link. *Geom. Topol.*, 5:369–398, 2001.
- [10] P. B. Kronheimer and T. S. Mrowka. Khovanov homology is an unknot-detector. *Publ. Math. Inst. Hautes Études Sci.*, 113:97–208, 2011.
- [11] Greg Kuperberg. Knottedness is in NP, modulo GRH. *Adv. Math.*, 256:493–506, 2014.
- [12] Marc Lackenby. Three-dimensional manifolds, lecture notes. Available at <https://people.maths.ox.ac.uk/lackenby/>, 1999.
- [13] Marc Lackenby. Exceptional surgery curves in triangulated 3-manifolds. *Pacific J. Math.*, 210(1):101–163, 2003.
- [14] Marc Lackenby. The crossing number of composite knots. *J. Topol.*, 2(4):747–768, 2009.
- [15] Marc Lackenby. A polynomial upper bound on Reidemeister moves. *Ann. of Math. (2)*, 182(2):491–564, 2015.
- [16] Marc Lackenby. The efficient certification of knottedness and Thurston norm. *Adv. Math.*, 387:Paper No. 107796, 142, 2021.
- [17] Marc Lackenby. Link hyperbolicity is in NP. In preparation, 2026.
- [18] Sergei Matveev. *Algorithmic topology and classification of 3-manifolds*, volume 9 of *Algorithms and Computation in Mathematics*. Springer, Berlin, second edition, 2007.
- [19] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and genus bounds. *Geom. Topol.*, 8:311–334, 2004.
- [20] Sucharit Sarkar and Jiajun Wang. An algorithm for computing some Heegaard Floer homologies. *Ann. of Math. (2)*, 171(2):1213–1236, 2010.
- [21] Jeffrey L. Tollefson and Ningyi Wang. Taut normal surfaces. *Topology*, 35(1):55–75, 1996.
- [22] Friedhelm Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Ann. of Math. (2)*, 87:56–88, 1968.

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