

THE WEBER-SEIFERT DODECAHEDRAL SPACE IS NON-HAKEN

BENJAMIN A. BURTON, J. HYAM RUBINSTEIN, AND STEPHAN TILLMANN

ABSTRACT. In this paper we settle Thurston’s old question of whether the Weber-Seifert dodecahedral space is non-Haken, a problem that has been a benchmark for progress in computational 3-manifold topology over recent decades. We resolve this question by combining recent significant advances in normal surface enumeration, new heuristic pruning techniques, and a new theoretical test that extends the seminal work of Jaco and Oertel.

1. INTRODUCTION

1.1. Motivation. In recent decades, 3-manifold topology has seen the resolution of many important decision problems, from Haken’s unknot recognition algorithm in the early 1960s [9] through to the difficult and multifaceted homeomorphism algorithm that was finally tied together with Perelman’s proof of the geometrisation conjecture in 2002. Many of these resolutions, however, are in theory only. Decision algorithms in 3-manifold topology are often highly complex and extremely inefficient, and only a handful have ever been implemented for practical use. A key motivation in computational topology is to elevate such algorithms from hypothetical procedures to practical tools that can be employed in other theoretical applications.

In 1980, Thurston asked whether the Weber-Seifert dodecahedral space is a Haken manifold [1] (we define these concepts in Section 1.2). Resolving this question has been a symbolic goal for computational topologists: In 1984, Jaco and Oertel devised an algorithm to decide whether a 3-manifold is Haken [11], and the only barrier to solving Thurston’s question has been improving, implementing, and running the Jaco-Oertel algorithm. Furthermore, resolving this specific question has broader implications—many of the improvements to the Jaco-Oertel algorithm have a wider impact, since this algorithm uses Haken’s machinery of *normal surface theory*, a toolset that has now become ubiquitous in 3-manifold decision problems.

Here we resolve Thurston’s question by proving the following theorem:

Theorem 1.1. *The Weber-Seifert dodecahedral space is non-Haken.*

Received by the editors March 3, 2010 and, in revised form, July 10, 2010.

2010 *Mathematics Subject Classification.* Primary 57N10.

Key words and phrases. Haken manifold, Weber-Seifert dodecahedral space, normal surface, incompressible surface.

The first author was supported under the Australian Research Council’s Discovery funding scheme (project DP1094516).

The second and third authors were partially supported under the Australian Research Council’s Discovery funding scheme (projects DP0664276 and DP1095760).

©2011 American Mathematical Society
Reverts to public domain 28 years from publication

The proof is essentially computational; as noted earlier, the main difficulty lies in redeveloping and extending the underlying algorithms to the point where we can implement and execute them in a feasible timeframe. We achieve this through a combination of techniques:

- We implement and exploit recent advances in the core machinery of normal surfaces, including the Q -theory of Tollefson [19] and the enumeration algorithms of the first author [5, 7].
- We extend the work of Jaco and Oertel, essentially showing that we can reduce the computational workload by searching for compatible *pairs* of normal surfaces instead of individual normal surfaces (Theorem 1.2 below).
- We develop new heuristic tests for identifying compressible surfaces; in theory these do not guarantee conclusive results, but in practice we find that they allow us to completely circumvent the most difficult step of the Jaco-Oertel algorithm (testing surfaces for compressibility).

1.2. Definitions and key results. The closed, irreducible 3-manifold M is *Haken* if it contains an embedded, injective surface different from the 2-sphere or the projective plane. A surface S in M is *injective* if the inclusion $S \subset M$ induces a monomorphism between the fundamental groups. The boundary of a regular neighbourhood of an injective surface is two-sided and (geometrically) incompressible.

The *Weber-Seifert dodecahedral space* is formed by identifying opposite faces of a dodecahedron with a $3/10$ twist, and it was one of the first known examples of a hyperbolic 3-manifold [20]. For our computations we use a triangulation of this manifold with 23 tetrahedra. There are at least three distinct triangulations with this number of tetrahedra, and they are the smallest triangulations of the Weber-Seifert dodecahedral space known to the authors.

The setting for this paper is the following. Let M be a closed, orientable, irreducible 3-manifold with a fixed triangulation \mathcal{T} . Work of Jaco and Oertel [11], Tollefson [18], and Oertel [16] studies the set of isotopy and projective isotopy classes of closed, injective surfaces in M and how they are represented in the projective solution space \mathcal{P} of normal surface theory. Here, two embedded, two-sided surfaces in M are in the same *projective isotopy class* if there exist multiples of each which are isotopic.

A natural environment for algorithmic topology on a 3-manifold is a *0-efficient triangulation* [13]. A triangulation of a closed 3-manifold is 0-efficient if the only normal 2-spheres are the links of the vertices of the triangulation (that is, frontiers of small regular neighbourhoods of these vertices). Standard facts about normal surfaces are recalled in Section 2, and the only non-standard term that needs to be clarified is the following: A surface in M is termed a *vertex surface* if it is a connected, two-sided normal surface, and the ray from the origin through its normal coordinate passes through a vertex of \mathcal{P} . Both the nomenclature and the definition of this concept vary widely in the literature; the present is chosen so that every vertex surface is two-sided and there is a unique vertex surface associated with each admissible vertex of \mathcal{P} (an *admissible* vertex has at most one non-zero quadrilateral coordinate per tetrahedron). In this terminology, Jaco and Oertel [11] show that M contains an embedded, injective surface if and only if one of the vertex surfaces is incompressible. (As usual, a sphere is not incompressible.)

In Section 2 we prove the following key result regarding vertex surfaces:

Theorem 1.2. *Suppose that the closed, orientable, irreducible 3-manifold M is Haken, and let \mathcal{T} be a 0-efficient triangulation of M . Then one of the following cases holds:*

1. (General case.) *There are two distinct, incompressible vertex surfaces which are compatible.*
2. (Exceptional case.) *There are at most finitely many projective isotopy classes of injective surfaces in M , and at least one of the following holds:*
 - (a) *the manifold M fibres over the circle, and there is a unique projective isotopy class of injective surfaces—in particular, $\text{rk } H^1(M; \mathbb{R}) = 1$ and the class is represented by a fibre; or*
 - (b) *there is a solid torus in M which contains at least three edges of \mathcal{T} and has normal boundary.*

Note that, even if there are at most finitely many projective isotopy classes of injective surfaces in M , the triangulation may also satisfy the general case.

The following corollary follows immediately from the above result. It is also a consequence of work of Oertel [16] and Tollefson [18], where faces of the projective solution space are studied using branched surfaces.

Corollary 1.3. *Let M be a closed, orientable, irreducible 3-manifold, and let \mathcal{T} be a 0-efficient triangulation of M . If $\text{rk } H^1(M; \mathbb{R}) \geq 2$, then there are two distinct, incompressible vertex surfaces which are compatible.*

In Section 3 we introduce some simple tests that can help identify when a two-sided surface is compressible. These tests are merely heuristic techniques—there is no guarantee for any particular surface that they will give a conclusive result, nor can they prove a surface to be incompressible. However, these techniques are found to be surprisingly effective in practice. Indeed, combined with some human intervention and the original algorithm of Jaco and Oertel, they can completely resolve the question of whether the Weber-Seifert dodecahedral space is non-Haken (Theorem 1.1), as noted at the end of Section 4.

Combining Theorem 1.2 with these heuristic techniques, we obtain the following new test to identify whether the given manifold M is non-Haken:

Test 1.4 (Is M non-Haken?). *First check that $\text{rk } H^1(M; \mathbb{R}) = 0$, since M is otherwise Haken. Then compute a 0-efficient triangulation and enumerate all vertex surfaces. Check that each vertex surface of zero Euler characteristic misses at most two edges. If this is not the case, the test is inconclusive. Otherwise, let $\mathcal{S} = \emptyset$. For each vertex surface, either determine a compression disc (for instance, using the heuristic techniques mentioned above) or add it to the set \mathcal{S} . Last, check that no two members of \mathcal{S} are compatible. If this is the case, M is non-Haken; otherwise, the test is inconclusive.*

This naïve test suffices to prove Theorem 1.1. The projective solution space of the chosen triangulation with 23 tetrahedra has 1751 admissible vertices: one is represented by a 2-sphere which links the vertex, 24 are represented by tori which link the edges, and the remaining 1726 vertices are represented by surfaces of negative Euler characteristic. Heuristic pruning finds compressing discs for *all but 16* of them, and no two of the remaining surfaces are compatible. The total computation time was a little over six hours; details are given in Section 4.

Other applications of Theorem 1.2 are given in Section 2. All routines developed in this paper are implemented in version 4.90 of the open-source software package *Regina* [2, 4].

2. COMPATIBLE, INJECTIVE VERTEX SURFACES

The original algorithm of Jaco and Oertel [11] to determine whether a manifold is Haken requires the enumeration of all vertex surfaces, and for each such surface one needs to check whether it is incompressible or not. Deciding incompressibility is computationally very expensive, and so this section uses the work of Jaco and Oertel to improve on their algorithm.

2.1. Triangulation. The notation and terminology of [13] will be used in this paper. Hence a triangulation \mathcal{T} consists of a union of t pairwise disjoint 3-simplices, $\tilde{\Delta}$, a set of face pairings, Φ , and a natural quotient map $p: \tilde{\Delta} \rightarrow \tilde{\Delta}/\Phi = M$. This is often referred to as a *semi-simplicial* or *singular* triangulation since not all simplices are necessarily embedded in M . The space $\tilde{\Delta}$ has a natural simplicial structure with four vertices for each 3-simplex. The quotient map p is required to be injective on the interior of each simplex of each dimension. The image of a simplex in $\tilde{\Delta}$ under p is a singular simplex in M . It is customary to refer to the image of a 3-simplex as a *tetrahedron in M* (or *of the triangulation*) and to refer to its faces, edges and vertices with respect to the pre-image. Similarly for images of 2-, 1-, and 0-simplices, which will be referred to as *faces*, *edges* and *vertices in M* (or *of the triangulation*) respectively. If a singular simplex is contained in ∂M , then it is termed *boundary* (such as a *boundary edge* or a *boundary face*); otherwise, it is termed *internal*. Notice that an internal singular simplex need not be disjoint from ∂M . A *normal isotopy of M* is an isotopy which leaves the image of the interior of every simplex in $\tilde{\Delta}$ invariant. The quotient space M is a manifold if the link of each vertex in M is a sphere or a disc.

2.2. Normal surfaces. This terminology again follows [13]. A normal surface in the triangulated 3-manifold M meets every tetrahedron in a pairwise disjoint, finite union of discs which are *normal triangles* or *normal quadrilaterals*. A normal surface is hence a properly embedded surface in M . The *normal coordinate* is a point in \mathbb{R}^{7t} that records the number of discs of each type in a normal surface. It satisfies a system of integral, linear equations, termed the *matching equations*. The set of all solutions with non-negative coordinates to this system is intersected with the affine subspace consisting of all points whose coordinates sum to one to give the *projective solution space \mathcal{P}* . This is a (bounded) polytope whose vertices have rational coordinates. Given any normal surface, its normal coordinate determines a unique point in \mathcal{P} .

A point in \mathbb{R}^{7t} is *admissible* if all of its coordinates are non-negative and at most one quadrilateral coordinate from each tetrahedron is non-zero. Each integral admissible solution to the matching equations determines a unique normal surface and vice versa. Two normal surfaces are said to be *compatible* if they do not meet a tetrahedron in quadrilateral discs of different types. This is the case if and only if the sum of their normal coordinates is admissible.

A surface in M is termed a *vertex surface* if it is a connected, two-sided normal surface and the ray from the origin through its normal coordinate passes through

a vertex of \mathcal{P} . Any ray from the origin through an admissible vertex of \mathcal{P} contains a unique vertex surface. A vertex surface is sometimes termed a *vertex normal surface* or a *fundamental edge surface* in the literature. It should not be confused with a *vertex linking surface* (which is the boundary of a small neighbourhood of a vertex in the triangulation).

2.3. 0–efficiency. The triangulation \mathcal{T} is 0–*efficient* if the only normal 2–spheres are vertex linking. It is shown in [13] that any triangulation of a closed, orientable, irreducible 3–manifold can be modified to a 0–efficient triangulation unless $M = S^3, \mathbb{R}P^3$ or $L(3, 1)$. Moreover, the conversion algorithm is implemented in *Regina* [4]. The algorithm typically takes only marginally longer than the time required to enumerate the admissible vertex solutions to the so-called Q –matching equations of Tollefson [19]. Detailed time trials can be found in [7].

2.4. The work of Jaco and Oertel. Jaco and Oertel [11] give an algorithm to decide whether a triangulated manifold is Haken. Surfaces are analysed using handle decompositions, but since their arguments are topological, the results carry over to triangulations. In this subsection, one of their key results will be restated in a topological version. This version follows verbatim from the proof of Theorem 2.2 in [11].

Let M be a closed, irreducible 3–manifold, and F_1 and F_2 be embedded surfaces in general position. Then $F_1 \cap F_2$ is a finite union of pairwise disjoint curves. A component of $F_1 \cap F_2$ is termed a *switch curve*. Let γ be a switch curve. A regular neighbourhood $N(\gamma)$ of γ is chosen such that $N(\gamma) \cap F_i$ is a regular neighbourhood of γ in F_i for each i . Since $N(\gamma)$ is either a solid Klein bottle or a solid torus, it follows that $\partial N(\gamma) \setminus (F_1 \cup F_2)$ consists either of two Möbius bands or of four annuli, called *switch bands* or *switch annuli* respectively. Two switch annuli are said to be *opposite* if they do not share a boundary component. A *switch along γ* consists of deleting the portion of $F_1 \cup F_2$ inside $N(\gamma)$ and connecting the free boundary components by a switch band or by two opposite switch annuli (depending on whether $\partial N(\gamma)$ is a Klein bottle or torus respectively). It follows that there are two possible switches along γ .

Denote by $F_1 + F_2$ the surface obtained from $F_1 \cup F_2$ where at each component of $F_1 \cap F_2$ one of the two possible switches has been chosen. If F_1 and F_2 are compatible normal surfaces with respect to a handle decomposition or a triangulation, then there is a natural choice at each switch curve, called a *regular switch*, such that $F_1 + F_2$ is again a normal surface.

The surface $F = F_1 + F_2$ is said to be in *reduced form* if it cannot be written as $F = F'_1 + F'_2$, where F'_i is isotopic to F_i in M and $F'_1 \cap F'_2$ has fewer components than $F_1 \cap F_2$. It should be noted that in these two sums, the embedding of F in M is the same (these are not equalities up to isotopy), and that any sum can be changed to a sum in reduced form.

We will also denote by $F_1 +_\gamma F_2$ the surface obtained from $F_1 \cup F_2$ by choosing the same switches as for $F_1 + F_2$ except for the curve γ , where the other switch possibility is chosen.

Theorem 2.1 (Jaco-Oertel). *Let M be a closed, irreducible 3–manifold and F be an embedded, two-sided and incompressible surface. If $F = F_1 + F_2$ is in reduced*

form, then either

- (1) F_1 and F_2 are incompressible, or
- (2) there exists $\gamma \in F_1 \cap F_2$ such that the surface $F_1 +_\gamma F_2$ has two components, F' and T , with the property that T is a torus which bounds a solid torus in M and has longitude isotopic to γ , and F' is isotopic to F via an isotopy which fixes the complement of a neighbourhood of the union of the solid torus and $N(\gamma)$. Moreover, $T = A \cup A'$, where A' is a switch annulus not contained in F and A is an annulus contained in F .

The second possibility is illustrated in Figures 6 and 8 of [11]. Under the additional assumption that F is a normal surface of least weight (with respect to a triangulation or handle decomposition), one sees that the second alternative is not possible and concludes that the summands are incompressible. Under these circumstances, one can also omit the hypothesis that the surface be in reduced form. This is the result stated in [11].

2.5. Proof of Theorem 1.2. Since M is Haken, the 0-efficient triangulation \mathcal{T} has a single vertex, v . Suppose S is a connected, injective surface in M . By possibly replacing S with the boundary of a regular neighbourhood of S , one may assume that S is a two-sided, geometrically incompressible surface. Recall that the weight of an embedded surface in M which is in general position with respect to the triangulation is the cardinality of its intersection with the 1-skeleton, written $\text{wt}(F) = |F \cap \mathcal{T}^{(1)}|$. After performing an isotopy, one may also assume that S is a normal surface of least weight in its isotopy class. If S is not a vertex surface, then, as in [11, Corollary 3.4], it follows from Theorem 2.1 that there are two compatible, injective vertex solutions. Since \mathcal{P} has finitely many vertices, it follows that if there are infinitely many projective isotopy classes of injective surfaces in M , then there are two compatible, injective vertex solutions. This is the general case (case 1) of the theorem and it remains to show that otherwise we are in the exceptional case (case 2a or 2b).

Hence suppose that there are at most finitely many projective isotopy classes of injective surfaces in M and that S is a vertex surface. We first produce a second normal surface isotopic but not normally isotopic to S by performing a *finger move*. That is, we push a portion of the surface along an edge and across the vertex in a controlled fashion. The details are as follows.

Let e be an edge of the triangulation. An intersection point $p \in S \cap e$ is *outermost* if one of the two components of $e \setminus \{p, v\}$ does not contain any other intersection points. Since S is not a vertex linking sphere, it follows that there is an edge e_0 having an outermost point $p_0 \in S \cap e_0$ which is incident with a quadrilateral disc in S . Orient e_0 such that travelling from p_0 in the positive direction to v , one does not meet S . Since S is two-sided, give S a transverse orientation which agrees at p_0 with the orientation of e_0 . Let $N(e_0)$ be a small regular neighbourhood of e_0 . Since M is orientable, this is a torus, and $S \cap N(e_0)$ consists of meridian discs since S is normal. Denote by D the connected component of $S \cap N(e_0)$ passing through p_0 . We now perform an isotopy of S which fixes $S \setminus D$ and moves p_0 along e_0 in the positive direction just past v . The resulting surface, S_1 , will not be normal and can be chosen such that

$$\text{wt}(S_1) = \text{wt}(S) + 2(E - 1),$$

where E is the number of edges of the triangulation. Now push S_1 slightly off S in the positive direction wherever they agree, giving a surface S_2 which is disjoint from S and has the same weight as S_1 .

Since no face in the triangulation is a cone [13] or a dunce hat [12], and S is pushed in the direction of the oriented edge e_0 , it follows from analysing the resulting isotopy of a quadrilateral disc in S meeting p_0 , that there is at least one face of the triangulation which S_2 meets in a return arc. Since S_2 is incompressible, it can be normalised by isotopies giving a normal surface S_3 , which satisfies

$$\text{wt}(S_3) \leq \text{wt}(S) + 2(E - 2).$$

Note that S_3 is disjoint from S , since S acts as a barrier surface for the normalisation of S_2 . If S_3 is normally isotopic to S , then there is a product region between the two surfaces which does not contain the unique vertex of the triangulation. However, there also is a product region between S_3 and S containing the vertex which arises from the first isotopy. It follows that M fibres over the circle with fibre S . If every connected, injective surface in M is the fibre in some fibration over the circle, then there is a unique projective isotopy class of injective surfaces in M (since we assume that there are at most finitely many). This is the first exceptional case stated in Theorem 1.2, case 2a.

Hence, assume that S is not a fibre in a fibration of M over the circle, which means that S_3 is not normally isotopic to S . If S_3 is a vertex surface, we have the general case of the theorem. Note also, by our assumption that S is least weight in its isotopy class, it follows that $\text{wt}(S_3) \geq \text{wt}(S)$. If $\text{wt}(S_3) = \text{wt}(S)$, then we can apply [11] to conclude that either S_3 is a vertex surface or S_3 is a sum of vertex surfaces which are all incompressible. In the latter case, each of the vertex surfaces is compatible with S and they cannot all be copies of S and so we again have the general case of the theorem.

Hence suppose that $\text{wt}(S_3) > \text{wt}(S)$, and that S_3 is not a vertex surface. We may further assume that S_3 has minimal weight amongst all normal surfaces that are *disjoint* from S and are isotopic to S but not normally isotopic to S . Since S_3 is not a vertex surface, there is a positive integer n and there are vertex surfaces V_i and positive integers n_i , such that

$$nS_3 = \sum n_i V_i,$$

where mF signifies m pairwise disjoint, parallel copies of F (each normally isotopic to F), and the sum uses the usual *regular switches* with respect to the triangulation. Since S_3 is compatible with S , each V_i is compatible with S . Moreover, at least one V_i is not normally isotopic to S since otherwise nS_3 is isotopic to $(\sum n_i)S$. We can therefore write

$$nS_3 = V + W,$$

where V is a vertex surface distinct from but compatible with S . There are normal surfaces V' and W' (isotopic in M to V and W respectively) such that $nS_3 = V' + W'$ in reduced form.

According to Theorem 2.1, we have two cases to consider. In the first case, V' (and hence V) is incompressible, and therefore S and V are two distinct compatible, injective, vertex surfaces. This is the general case of the theorem.

In the second case, following [11], there exists $\gamma \in V' \cap W'$ such that the surface $V' +_\gamma W'$ is the disjoint union of two surfaces, one of which is isotopic to nS_3 and

the other is a compressible torus T . Since the switch curve γ is two-sided in V' and W' , it follows that $V' +_\gamma W'$ has $(n - 1)$ components which are normally isotopic copies of S_3 . In addition, there is the torus T and a component X which is isotopic to S_3 . The latter normalises to give a surface S_4 .

If S_4 is disjoint from S , then it must be normally isotopic to S since its weight is strictly less than the weight of S_3 . This gives an isotopy from S_3 to S which must pass through the vertex since otherwise M is a bundle. It follows that the vertex is contained in the solid torus bounded by T . Considering weight, we have

$$n \operatorname{wt}(S_3) > (n - 1) \operatorname{wt}(S_3) + \operatorname{wt}(S) + \operatorname{wt}(T),$$

whence

$$\operatorname{wt}(S) + 2(E - 2) \geq \operatorname{wt}(S_3) > \operatorname{wt}(S) + \operatorname{wt}(T),$$

giving

$$2(E - 2) > \operatorname{wt}(T).$$

Since T is separating, it must meet each edge an even number of times, whence the solid torus bounded by it contains at least three edges. The boundary of a regular neighbourhood of one of the edges is an embedded torus and a barrier surface (see [13]). Since M is irreducible, the process of normalising T in the complement of this torus shows that T either shrinks to a normal surface using isotopies or it shrinks to a 2-sphere embedded in a tetrahedron using isotopies and a single compression. In the latter case, M is a lens space, contradicting the fact that it is Haken. Since normalisation does not increase the weight, we have the second exceptional case, case 2b, stated in Theorem 1.2.

It remains to analyse the possibility that S_4 is not disjoint from S . In this case, X cannot be disjoint from S since otherwise S is a barrier for the normalisation of X to S_4 . But if $S \cap X \neq \emptyset$, then S meets a neighbourhood of the union of the solid torus bounded by T and $N(\gamma)$ in a union of annuli and can be isotoped to be disjoint from this and X . (See Figure 8 of [11] for an illustration of the isotopy.) This would reduce the weight of S , which is a contradiction. Therefore S does not meet X and hence does not meet S_4 . This concludes the proof of the theorem. \square

2.6. Applications of Theorem 1.2. Throughout this subsection, M denotes a closed, orientable, irreducible 3-manifold with fixed 0-efficient triangulation \mathcal{T} .

Definition 2.2 (Large normal torus). A normal torus disjoint from at least three edges will be termed a *large normal torus*.

For instance, the boundary of a regular neighbourhood of a layered solid torus subcomplex in M shrinks to a large normal torus unless M is a lens space (see, for instance, [3, 12] for a definition of this subcomplex). Such a subcomplex appears in the natural triangulations of Dehn fillings of knot complements. Large normal tori therefore occur in many natural triangulations of both Haken and non-Haken manifolds.

Algorithm 2.3 (Large normal torus recognition). *To check whether there is a large normal torus in M , it is necessary and sufficient to verify that each vertex surface of Euler characteristic zero is disjoint from at most two edges.*

This algorithm follows immediately from the fact that (i) edge weights and Euler characteristics are additive, and (ii) the only normal 2-spheres are vertex linking and hence can be made disjoint from any other normal surface.

If M is atoroidal and \mathcal{T} is 0-efficient, then every normal torus bounds a solid torus [12]. It follows that in this case, Algorithm 2.3 can be used to decide whether there is a solid torus in M which contains at least three edges and has normal boundary. In the general case, Algorithm 2.3 only helps in certain cases to identify when such a solid torus does not exist.

This is also the philosophy in the tests below. Whilst in theory there are algorithms to determine whether a given manifold is non-Haken or a tiny bundle, they often turn out to be impractically slow. Below are some simple tests that allow an answer in a feasible amount of time, even though they may not always be conclusive. We begin with the new non-Haken test described in the introduction.

Test 1.4 (Is M non-Haken?). *First check that $\text{rk } H^1(M; \mathbb{R}) = 0$, since M is otherwise Haken. Then enumerate all vertex surfaces. Use Algorithm 2.3 to check that there is no large normal torus. If there is a large normal torus, the test is inconclusive. Otherwise, let $\mathcal{S} = \emptyset$. For each vertex surface, either determine a compression disc (for instance, using the heuristic techniques described in Section 3) or add it to the set \mathcal{S} . Last, check that no two members of \mathcal{S} are compatible. If this is the case, then M is non-Haken; otherwise, the test is inconclusive.*

Several routines are described in Section 3 that search for compression discs that are computationally easy to find. These routines are often sufficient to keep the list \mathcal{S} relatively short. There are, however, many classes of triangulated non-Haken manifolds which the above approach will not recognise as non-Haken. For instance, if the triangulation contains a layered solid torus subcomplex, then there is a large normal torus as noted earlier.

Definition 2.4 (Tiny bundle). The closed, orientable, irreducible 3-manifold M is a *tiny bundle* if M fibres over the circle and there is a unique projective isotopy class of injective surfaces in M .

Test 2.5 (Is M a tiny bundle?). *The computational steps are as in Test 1.4, except that one first checks that $\text{rk } H^1(M; \mathbb{R}) = 1$.*

Test 2.6 (Is there a finite number of projective isotopy classes?). *First show that M is Haken. Then enumerate all vertex surfaces. Let $\mathcal{S} = \emptyset$. For each vertex surface, either determine a compression disc or add it to the set \mathcal{S} . Last, check that no two members of \mathcal{S} are compatible. If this is the case, then there are at most finitely many projective isotopy classes of injective surfaces; otherwise, the test is inconclusive.*

3. HEURISTIC PRUNING

In this section we introduce some simple tests that can help identify when an embedded surface within a closed 3-manifold is compressible. These are merely heuristic techniques—there is no guarantee for any particular surface that they will give a conclusive result. Moreover, they work in one direction only—they can never show a surface to be *incompressible*.

With these tests, we are able to take a list of *potential* incompressible surfaces (such as the vertex normal surfaces within a triangulation) and filter out irrelevant surfaces. This leaves us fewer surfaces on which we must run more expensive procedures, such as the conclusive but extremely slow incompressibility algorithm of Jaco and Oertel [11].




FIGURE 1. Cutting along a face with three boundary edges

The key idea behind these heuristic tests is to search within a bounded 3-manifold triangulation for embedded discs with simple combinatorial structures. Lemmata 3.1 and 3.2 describe these structures, and Algorithm 3.4 shows how they can be used effectively to test for compressibility. Although these tests are simple in theory, we see in Section 4 that they can be surprisingly effective in practice.

Lemma 3.1. *Let \mathcal{T} be a triangulation of a bounded 3-manifold M (that is, a compact 3-manifold with non-empty boundary). Let F be a non-boundary face of \mathcal{T} , and suppose that all three edges of F lie entirely within the boundary ∂M , as illustrated in the leftmost diagram of Figure 1. Note that neither the edges of F nor the vertices of F are required to be distinct.*

Suppose that we “unglue” the two tetrahedra on either side of F (that is, we remove the corresponding pair of tetrahedron faces from the list of face identifications that make up \mathcal{T}), as illustrated in the rightmost diagram of Figure 1. Then the result is a new triangulation \mathcal{T}' of some 3-manifold M' , which is homeomorphic to M sliced along a properly embedded disc.

Note that we do not describe this operation as “slicing \mathcal{T} along the face F ”, since F might have self-intersections and therefore might not be embedded. However, whether or not F is embedded, the act of ungluing the two tetrahedra on either side of F is well defined and simple to perform. Proving that self-intersections of F do not matter is in fact the main point of this lemma.

Proof. If the edges and vertices of F are all distinct, then this result is straightforward—the face F forms a properly embedded disc in M , and the new triangulation \mathcal{T}' is just M sliced along this disc.

Consider then the case where different edges and/or vertices of F are identified. We first prove that \mathcal{T}' is indeed a 3-manifold triangulation, and then we show that the corresponding manifold M' has the required property.

The only situations in which \mathcal{T}' might *not* be a 3-manifold triangulation are (i) where some edge of \mathcal{T}' is identified with itself in reverse, and (ii) where some vertex of \mathcal{T}' does not have a small closed neighbourhood that is a 3-ball.

The first situation is easily eliminated, since ungluing the tetrahedra on either side of F cannot create any new edge identifications (though it can remove them). Consider then some vertex V of the face F in \mathcal{T} , and let $\text{nb}(V)$ be a small closed neighbourhood of V . Since the edges and vertices of F all lie on the boundary ∂M , the neighbourhood $\text{nb}(V)$ must be a 3-ball with V on its boundary, as illustrated in the leftmost diagram of Figure 2.

Now consider what happens when we unglue the tetrahedra on either side of F . The intersection $F \cap \text{nb}(V)$ consists of up to three “triangular” discs in $\text{nb}(V)$ (one for each corner of F that meets V). Note that these discs might be joined along




FIGURE 2. The ungluing operation in the vicinity of a vertex of F




FIGURE 3. A “bad” hole where V has no 3-ball neighbourhood

sections of their boundaries (corresponding to edges of F) to form larger discs or even branched structures. Examples of such discs are shown in the centre diagram of Figure 2.

The interiors of these individual triangular discs are disjoint and embedded in $\text{int}(\text{nb}d(V))$, and the boundaries of these discs lie in the boundary $\partial \text{nb}d(V)$. Although the disc boundaries might intersect on ∂M (as a result of identifications between edges of F), they can never intersect in $\text{int}(M)$ —in other words, the portions of the disc boundaries within $\text{int}(M)$ are also disjoint and embedded.

We find that, when we unglue the tetrahedra on either side of F , we effectively slice $\text{nb}d(V)$ along these discs, as illustrated in the rightmost diagram of Figure 2. This divides $\text{nb}d(V)$ into several smaller 3-balls, splitting V into several different vertices as a consequence. Note that any disc edges that are pinched together on ∂M will fall apart, since only the face F was holding them together (this happens with the rightmost disc in Figure 2).

Because of the well-behaved manner in which these discs are placed within $\text{nb}d(V)$, we see that every resulting vertex of \mathcal{T}' has a small closed neighbourhood that is a 3-ball (that is, no “bad” holes have been cut out of $\text{nb}d(V)$ as illustrated in Figure 3). Therefore \mathcal{T}' is indeed a 3-manifold triangulation, and we denote the corresponding 3-manifold by M' . Figure 4 illustrates the entire transformation from \mathcal{T} to \mathcal{T}' in the case where all three vertices of F are identified (here the single vertex V in \mathcal{T} splits into four vertices in \mathcal{T}').

Now that we know that \mathcal{T}' is indeed a triangulation of the 3-manifold M' , the remainder of the lemma is straightforward to prove. Define N to be the 3-manifold obtained by removing a small open neighbourhood of the boundary ∂M from M . Likewise, let N' be obtained by removing a small open neighbourhood of $\partial M'$ from M' . It is clear that $F \cap N$ is a properly embedded disc in N , and that the 3-manifold N' is obtained by slicing N along this disc. Since M and M' are homeomorphic to N and N' , respectively, it follows that M' is obtained by slicing along a properly embedded disc in M . \square




FIGURE 4. Cutting along a face whose vertices are all identified




FIGURE 5. Cutting along a disc surrounding an edge of degree one

Lemma 3.2. *Let \mathcal{T} be a triangulation of a bounded 3-manifold M . Let e be a non-boundary edge in \mathcal{T} of degree one, let Δ be the (unique) tetrahedron containing e , and suppose that the edge opposite e in Δ lies entirely within the boundary ∂M , as illustrated in the leftmost diagram of Figure 5.*

Suppose we replace Δ with two tetrahedra Δ_1 and Δ_2 , each with two faces folded together to form an edge of degree one, as illustrated in the rightmost diagram of Figure 5. Of the two portions of \mathcal{T} that were originally joined to Δ along the shaded faces, we join one of these portions to Δ_1 and the other to Δ_2 . We leave the remaining faces of Δ_1 and Δ_2 as boundary faces.

Then the result is a new triangulation \mathcal{T}' of some 3-manifold M' , which is homeomorphic to M sliced along a properly embedded disc.

Proof. In contrast to the previous result, this lemma contains no unusual cases. The edge opposite e in Δ always bounds a properly embedded disc in M (running directly through the centre of the tetrahedron Δ), and so \mathcal{T}' triangulates a 3-manifold M' that is obtained by slicing along this disc. □

In order to take full advantage of Lemmata 3.1 and 3.2, it helps to have as many faces and edges of a triangulation exposed to the boundary as possible. The following operation assists us in this regard.

Lemma 3.3. *Let \mathcal{T} be a triangulation of a bounded 3-manifold M . Let F be a non-boundary face of \mathcal{T} , and suppose that precisely two of the three edges of F lie within the boundary ∂M , as illustrated in the leftmost diagram of Figure 6. Once again, neither the edges nor the vertices of F are required to be distinct.*

Suppose that we unglue the two tetrahedra on either side of F , exposing these tetrahedra to the boundary as illustrated in the rightmost diagram of Figure 6. Then the result is a new triangulation \mathcal{T}' of the same 3-manifold M .

We refer to this operation as a book opening move.




FIGURE 6. Performing a book opening move




FIGURE 7. New types of discs in the intersection $F \cap \text{nbnd}(V)$

Proof. The proof is almost identical to Lemma 3.1, and we do not repeat the details. The only differences are

- Instead of slicing the manifold M along a properly embedded disc, we slice it along a “half-properly embedded” disc. By this, we mean an embedded disc whose boundary consists of (i) an arc in ∂M , and (ii) an arc in $\text{int}(M)$. Slicing along such a disc will never change the underlying manifold M .
- When we examine the neighbourhood of a vertex V , the intersection $F \cap \text{nbnd}(V)$ can include new types of discs that are “half-properly embedded” in $\text{nbnd}(V)$, as illustrated in Figure 7. Such discs are harmless however, and do not change the key fact that the resulting vertices in \mathcal{T}' all have 3-ball neighbourhoods.

We refer to the proof of Lemma 3.1 for the full details. □

We can now pull together all of these operations to build a heuristic algorithm for detecting compressible surfaces.

Algorithm 3.4 (Heuristic pruning). *Let \mathcal{T} be a triangulation of the closed 3-manifold M , and let S be an embedded surface in \mathcal{T} . We can potentially show that S is a compressible surface through the following procedure:*

- (1) *Cut the triangulation \mathcal{T} along the surface S , and retriangulate the resulting bounded 3-manifold M' (which may be disconnected). Let \mathcal{T}' denote this new bounded triangulation, let β denote the number of boundary components, and let σ denote the number of boundary spheres (so $\sigma \leq \beta$).*
- (2) *Perform local simplification moves on \mathcal{T}' to reduce the number of tetrahedra (such as Pachner moves [17] or boundary shellings). There is no need to produce a minimal triangulation; it suffices to reach a point where there are no further immediate simplifications that can be done.*
- (3) *Perform book opening moves on \mathcal{T}' as described by Lemma 3.3 until no more can be done.*
- (4) *Search for all locations within \mathcal{T}' at which the preconditions of Lemmata 3.1 and 3.2 are satisfied. That is, search for internal faces whose edges are all*

boundary, and search for internal degree one edges whose opposite edges are boundary.

- (5) For each such location, temporarily perform the corresponding operation upon \mathcal{T}' that slices M' along a properly embedded disc.

If any of the sliced triangulations obtained in step 5 still has β boundary components or still has σ boundary spheres, then the original surface S is compressible in M .

Proof. Given Lemmata 3.1–3.3, the only part of this algorithm that remains to be proven is the final claim that, if some sliced triangulation from step 5 has either β boundary components or σ boundary spheres, then S is compressible in M .

Let D be the properly embedded disc in M' that we slice along in step 5. If D is *not* a compressing disc for S , then ∂D bounds a disc in $\partial M'$, whereupon slicing along D' produces a new 2–sphere boundary component but otherwise leaves all existing boundary components unchanged. That is, we obtain $\beta + 1$ boundary components, $\sigma + 1$ of which are 2–spheres.

Note that a compressing disc can produce a new boundary sphere without a new boundary component (for instance, slicing along the meridional disc of a solid torus), or a new boundary component without a new boundary sphere (for instance, slicing along a separating disc with non-trivial topology on each side). This is why we must count both boundary spheres *and* boundary components in order to detect compressing discs. \square

We finish this section with some notes regarding both the structure and implementation of Algorithm 3.4.

- The reason for steps 2 and 3 is to increase our chances of meeting the preconditions of Lemmata 3.1 and 3.2. In particular, simplifying the triangulation increases our chances of finding a region of \mathcal{T}' that is only “one tetrahedron thick”, and book opening moves help expose more edges and vertices to the boundary.
- The precise local simplification moves of step 2 are left up to the reader. Many moves of this type are documented in the literature (particularly by authors involved in census enumeration); see [3, 14] for some examples.

The moves that we use in the following section with the Weber-Seifert dodecahedral space include Pachner moves (also called bistellar moves [17]), collapsing edges between distinct vertices, removing tetrahedra through boundary shellings, simplifying triangulations in the vicinity of low-degree edges and vertices, and the book opening move and its inverse (the *book closing* move).

- Step 1, in which we cut along the surface S in the triangulation \mathcal{T} , causes a number of difficulties. The most severe problem is that it can generate a very large number of tetrahedra—in the case of the Weber-Seifert dodecahedral space, we frequently find tetrahedra numbering in the thousands. It is therefore critical to have a simplification procedure that is both fast and effective.

Moreover, cutting along a surface is messy for a programmer to implement, since tetrahedra can be subdivided into many different pieces of up to eleven distinct shapes (see Figure 8 for some examples). Each of these shapes must be individually retriangulated (typically by the programmer as




FIGURE 8. Sample pieces obtained by cutting along a normal surface

he or she implements the routine), and the code must then be able to automatically adjust the triangulations of these pieces so that the quadrilaterals, pentagons, and hexagons on their boundaries can be glued together.

We can make this cutting operation simpler if we first remove a small neighbourhood of each vertex of the original triangulation. This reduces the number of different shapes from eleven to four, which is significantly easier for a programmer to manage. The boundaries of these pieces are also simpler to handle, with only quadrilaterals and hexagons to worry about. Of course we must not forget to glue the missing 3-balls back onto the boundary of the new triangulation once we are finished.

4. THE WEBER-SEIFERT DODECAHEDRAL SPACE

To conclude this paper, we apply the new Test 1.4 to resolve an outstanding conjecture of Thurston. The *Weber-Seifert dodecahedral space* is formed by identifying opposite faces of a dodecahedron with a $3/10$ twist, and it was one of the first known examples of a hyperbolic 3-manifold [20]. It was conjectured by Thurston that the Weber-Seifert dodecahedral space is non-Haken [1], and here we prove this to be true.

By building on Haken's earlier work, Jaco and Oertel gave an algorithm in 1984 to determine whether a given 3-manifold M has the Haken property [11]. This algorithm has since been improved by other authors [10, 19], though the basic framework remains the same:

- (1) Enumerate all vertex surfaces in some triangulation \mathcal{T} of M .
- (2) For each vertex surface $S \subset M$, test whether S is incompressible:
 - (i) Cut M along the surface S and retriangulate.
 - (ii) For each component \mathcal{T}' of the resulting triangulation, enumerate all fundamental normal surfaces in \mathcal{T}' and test whether any of these is a compressing disc.

As noted in the introduction, all normal surfaces considered here are embedded. In particular, we do not consider the more general case of immersed and/or singular surfaces within \mathcal{T} .

In theory, proving the Weber-Seifert dodecahedral space to be non-Haken should be a simple matter of running the Jaco-Oertel algorithm. However, this algorithm is extremely slow in practice. If t is the number of tetrahedra in \mathcal{T} , then step 1 can grow exponentially slow in t and produce exponentially many surfaces, and each triangulation \mathcal{T}' in step 2 can contain exponentially many tetrahedra. Even worse, each enumeration in step 2(ii) can grow exponentially slow in the size of \mathcal{T}' , which

TABLE 1. The pairwise identifications of tetrahedron faces in the triangulation \mathcal{T}_{WS}

	123	124	134	234		123	124	134	234
<i>A</i>	<i>E</i> : 123	<i>D</i> : 124	<i>C</i> : 134	<i>B</i> : 234	<i>M</i>	<i>D</i> : 123	<i>L</i> : 142	<i>F</i> : 132	<i>G</i> : 432
<i>B</i>	<i>H</i> : 123	<i>G</i> : 124	<i>F</i> : 134	<i>A</i> : 234	<i>N</i>	<i>J</i> : 423	<i>U</i> : 431	<i>G</i> : 314	<i>E</i> : 234
<i>C</i>	<i>K</i> : 123	<i>J</i> : 124	<i>A</i> : 134	<i>I</i> : 234	<i>O</i>	<i>I</i> : 132	<i>R</i> : 234	<i>E</i> : 134	<i>Q</i> : 314
<i>D</i>	<i>M</i> : 123	<i>A</i> : 124	<i>H</i> : 134	<i>L</i> : 234	<i>P</i>	<i>L</i> : 143	<i>E</i> : 124	<i>W</i> : 134	<i>S</i> : 431
<i>E</i>	<i>A</i> : 123	<i>P</i> : 124	<i>O</i> : 134	<i>N</i> : 234	<i>Q</i>	<i>U</i> : 213	<i>F</i> : 124	<i>O</i> : 324	<i>V</i> : 134
<i>F</i>	<i>M</i> : 143	<i>Q</i> : 124	<i>B</i> : 134	<i>I</i> : 134	<i>R</i>	<i>G</i> : 123	<i>T</i> : 123	<i>W</i> : 423	<i>O</i> : 124
<i>G</i>	<i>R</i> : 123	<i>B</i> : 124	<i>N</i> : 314	<i>M</i> : 432	<i>S</i>	<i>J</i> : 123	<i>V</i> : 123	<i>P</i> : 432	<i>H</i> : 234
<i>H</i>	<i>B</i> : 123	<i>T</i> : 124	<i>D</i> : 134	<i>S</i> : 234	<i>T</i>	<i>R</i> : 124	<i>H</i> : 124	<i>J</i> : 134	<i>W</i> : 142
<i>I</i>	<i>O</i> : 132	<i>K</i> : 134	<i>F</i> : 234	<i>C</i> : 234	<i>U</i>	<i>Q</i> : 213	<i>L</i> : 213	<i>N</i> : 421	<i>K</i> : 234
<i>J</i>	<i>S</i> : 123	<i>C</i> : 124	<i>T</i> : 134	<i>N</i> : 231	<i>V</i>	<i>S</i> : 124	<i>K</i> : 124	<i>Q</i> : 234	<i>W</i> : 123
<i>K</i>	<i>C</i> : 123	<i>V</i> : 124	<i>I</i> : 124	<i>U</i> : 234	<i>W</i>	<i>V</i> : 234	<i>T</i> : 243	<i>P</i> : 134	<i>R</i> : 341
<i>L</i>	<i>U</i> : 214	<i>M</i> : 142	<i>P</i> : 132	<i>D</i> : 234					

becomes *doubly* exponential in t . For these reasons, the Jaco-Oertel algorithm has to date never been successfully applied to the Weber-Seifert dodecahedral space.

Our proof begins in the same manner as step 1 of the Jaco-Oertel algorithm—we triangulate the Weber-Seifert dodecahedral space and then obtain a list of all vertex surfaces with the help of recent developments in normal surface enumeration algorithms [5, 7]. However, the doubly exponential enumeration of step 2 remains out of our reach, and so instead we use Theorem 1.2 and heuristic pruning to show that each of the surfaces in our list has a compressing disc.

All of the computation in this proof was carried out using the open-source software package *Regina* [2, 4]. The supporting data for this proof (including the relevant triangulations and normal surfaces) is contained in the file `weber-seifert.rga`, which readers can download from the *Regina* website at <http://regina.sourceforge.net/data.html>.

Definition 4.1. Let \mathcal{T}_{WS} denote the 23-tetrahedron triangulation described in Table 1, which lists the pairwise identifications between the 4×23 faces of 23 individual tetrahedra. Here the tetrahedra are labelled A, \dots, W and the vertices of each tetrahedron are numbered $1, \dots, 4$.

Alternatively, we can describe \mathcal{T}_{WS} using the dehydration notation of Callahan, Hildebrand and Weeks. The dehydration string of \mathcal{T}_{WS} is

`xppphocgaeaaahimnkontspmuuqrsvuwvtvwxwjjsvvcxxjjqattdworrko,`

from which we can recover the full structure of \mathcal{T}_{WS} using the the rehydration procedure described in [8].

To read Table 1, each row gives the face identifications for a single tetrahedron, and each column indicates one of the four faces. For instance, the cell in the bottom left corner indicates that face 123 of tetrahedron L is identified with face 214 of tetrahedron U (with vertices 1, 2 and 3 of L identified with vertices 2, 1 and 4 of U respectively).

Lemma 4.2. \mathcal{T}_{WS} is a triangulation of the Weber-Seifert dodecahedral space.




FIGURE 9. A pentagonal cone triangulated with five tetrahedra




FIGURE 10. Moves to simplify the triangulation \mathcal{T}_{WS}

Proof. The triangulation \mathcal{T}_{WS} can be constructed as follows:

- (1) Build a regular dodecahedron by joining together twelve pentagonal cones, with the twelve apices meeting at the centre of the dodecahedron and the twelve pentagonal bases forming the boundary of the dodecahedron.
- (2) Triangulate each pentagonal cone with five tetrahedra, as illustrated in Figure 9.
- (3) Identify opposite faces of the dodecahedron with a $3/10$ twist, giving a closed 3-manifold triangulation with 60 tetrahedra.
- (4) Simplify this triangulation by first collapsing edges between distinct vertices, and then applying 3-2 Pachner moves (also called bistellar moves [17]). These operations are illustrated in Figure 10.

It is clear from this construction that \mathcal{T}_{WS} triangulates the Weber-Seifert dodecahedral space as claimed. \square

In fact, we conjecture that \mathcal{T}_{WS} is a *minimal* triangulation of the Weber-Seifert dodecahedral space, i.e., that the space cannot be triangulated with 22 tetrahedra

TABLE 2. Counting vertex surfaces according to genus

Genus	0	1	2	3	4	5	6	7	8
Quad. vertex surfaces		24	43	82	135	30		300	36
Standard vertex surfaces	1	24	187	465	387	115	32	318	54
Genus (continued)	9	10	11	12	13	14	15	16	Total
Quad. vertex surfaces	36	12							698
Standard vertex surfaces	54	36	30		18	12	12	6	1751

or fewer. It should be noted that \mathcal{T}_{WS} is not the only 23-tetrahedron triangulation of the Weber-Seifert dodecahedral space; through a repeated application of 2–3 and 3–2 Pachner moves, we can obtain at least two distinct¹ alternatives, with the following dehydration strings:

$$(4.1) \quad \begin{aligned} & \text{xppphjgbgaaaaginlnslmupurtqsqrwwvwmwfcvuvkekaldmphextbvfw} \\ & \text{xppfpnkalaaaamkionrmtnpqrsutsqwuvvwxdwvgtvqkpwpxalnjcrkfn} \end{aligned}$$

Now that we are equipped with a triangulation, we embark on the first step of the Jaco-Oertel algorithm—the enumeration of vertex surfaces. At the core of this enumeration is a linear programming problem that takes place in a vector space of dimension $7t = 161$ (where $t = 23$ is the number of tetrahedra in \mathcal{T}_{WS}).

A direct enumeration in \mathbb{R}^{161} remains out of our reach computationally, and so we take an indirect approach instead. Tollefson [19] describes a smaller vector space of dimension $3t$, in which we consider only the quadrilateral discs in each normal surface. We refer to this smaller vector space \mathbb{R}^{3t} as *quadrilateral coordinates*, and in contrast we refer to the original \mathbb{R}^{7t} as used by Haken and then Jaco and Oertel as *standard coordinates*.

Our plan is (i) to enumerate all vertex surfaces in quadrilateral coordinates, and then (ii) to convert this result into a list of all vertex surfaces in standard coordinates. Note that the latter step is not just a matter of changing between coordinate systems, since the “vertex surface” property is not preserved between coordinate systems—instead we must apply the complex (though extremely fast) conversion procedure described in [5].

In fact, step (ii) is not necessary for the Jaco-Oertel algorithm, since Tollefson proves that some two-sided incompressible surface must appear as a vertex surface in quadrilateral coordinates, if such a surface exists at all. However, because our proof relies on Theorem 1.2, we must work in standard coordinates. For us then, quadrilateral coordinates are simply a means to an end.

By running the streamlined normal surface enumeration algorithm described in [7], we obtain the following result through direct computation (recalling that vertex surfaces are defined here to be two-sided and connected):

Lemma 4.3. *The triangulation \mathcal{T}_{WS} has 698 vertex surfaces in quadrilateral coordinates. The genera of these surfaces are distributed according to the first row of Table 2.*

The computation required to prove Lemma 4.3 is not trivial—on a 2.4 GHz Intel Core 2 CPU, the enumeration takes a little under $5\frac{1}{2}$ hours. Without recent

¹By *distinct*, we mean that one triangulation cannot be obtained from another simply by relabelling tetrahedra and their vertices.

TABLE 3. The 16 vertex surfaces that remain after heuristic pruning

Genus	4	5	7	8	9	Total
Surfaces remaining	1	3	8	1	3	16

improvements to the normal surface enumeration algorithm [7], this computation could take orders of magnitude longer, and without Tollefson's quadrilateral coordinates it would remain completely infeasible.

Although Table 2 only lists the genus of each surface, complete descriptions of all 698 surfaces can be found in the file `weber-seifert.rga`, as noted at the beginning of this section.

We now make our move into standard coordinates. By running the 698 surfaces of Lemma 4.3 through the quadrilateral-to-standard conversion algorithm described in [5], we obtain the following result:

Lemma 4.4. *The triangulation \mathcal{T}_{WS} has 1751 vertex surfaces in standard coordinates. The genera of these surfaces are distributed according to the second row of Table 2.*

In contrast to the full enumeration in quadrilateral coordinates, the conversion algorithm of [5] is extremely fast, taking just over 1 second on the same 2.4 GHz Intel Core 2 CPU. As before, complete descriptions of all 1751 surfaces can be downloaded in the file `weber-seifert.rga`.

We pause here to make some observations about the low-genus surfaces in our list. By examining the individual normal discs that make up the spheres and tori in Table 2, we obtain the following result:

Lemma 4.5. *The only vertex normal sphere in \mathcal{T}_{WS} is the frontier of a small regular neighbourhood of the single vertex of \mathcal{T}_{WS} . Likewise, the only vertex normal tori in \mathcal{T}_{WS} are the frontiers of small regular neighbourhoods of the 24 edges of \mathcal{T}_{WS} .*

Using the nomenclature of Jaco and Rubinstein [13], these surfaces are called *vertex links* and *thin edge links*, respectively, and the triangulation \mathcal{T}_{WS} is both 0-efficient and 1-efficient as a result.

Now that we have a full list of vertex surfaces at our disposal, we can bring in the techniques of Theorem 1.2 and heuristic pruning to prove our final result.

Theorem 1.1. *The Weber-Seifert dodecahedral space is non-Haken.*

Proof. Suppose the Weber-Seifert dodecahedral space does contain a two-sided incompressible surface. Cases 2a and 2b of Theorem 1.2 are easily eliminated through a homology computation and Lemma 4.5, and so it follows from Theorem 1.2 that there must be two distinct, compatible, incompressible vertex surfaces S_1, S_2 in \mathcal{T}_{WS} . We therefore run through our list of 1751 vertex surfaces in search of such a pair S_1, S_2 .

We can eliminate the vertex linking sphere and the 24 vertex linking tori immediately. Running Algorithm 3.4 over the 1726 remaining surfaces shows that 1710 of these contain a compressing disc. That is, heuristic pruning eliminates *all but 16* of these vertex surfaces. Those surfaces that remain are summarised in Table 3 (once again, see `weber-seifert.rga` for their full descriptions).

It follows that, if they exist at all, the surfaces S_1 and S_2 must belong to this smaller list. However, comparing quadrilateral types for all $\binom{16}{2}$ pairs shows that

no two of these surfaces are compatible, and so by Theorem 1.2 the Weber-Seifert dodecahedral space cannot be Haken. \square

We finish with a handful of observations regarding the different elements used in the proof of Theorem 1.1.

- It was mentioned earlier that the number of vertex surfaces can grow exponentially in the number of tetrahedra t . The best theoretical bounds known to date are based on the upper bound theorem of McMullen [15], yielding theoretical limits of $O(4^t)$ in quadrilateral coordinates and $O(15^t)$ in standard coordinates [7]. It is therefore surprising in our case with $t = 23$ to find just 698 and 1751 vertex surfaces respectively. Such discrepancies between theory and practice are common and are discussed in greater detail in the paper [6].
- We can recall from Section 3 that heuristic pruning involves two distinct tests: one for internal faces with three boundary edges (Lemma 3.1) and one for discs surrounding edges of degree one (Lemma 3.2). It is worth comparing the relative effectiveness of these tests.

Of the 1726 vertex surfaces upon which we attempt heuristic pruning, 1695 can be eliminated using Lemma 3.1 but only 88 can be eliminated through Lemma 3.2. These are success rates of approximately 98% and 5% respectively. It appears therefore that testing for faces with three boundary edges is significantly more powerful in practice.

- Tollefson proves that quadrilateral coordinates are sufficient for running the original Jaco-Oertel algorithm [19], whereas in this paper we use standard coordinates instead. It is worth noting that this choice does not lead to any significant loss of efficiency or power:
 - Assuming that we already have a list of vertex surfaces in quadrilateral coordinates, creating a list of vertex surfaces using the conversion algorithm of [5] is extremely fast, taking only a matter of seconds of processing time.
 - Although we have more surfaces to deal with in standard coordinates (1751 instead of 698), heuristic pruning eliminates these differences entirely. That is, applying heuristic pruning to the vertex surfaces in *quadrilateral* coordinates leaves us with precisely the same 16 surfaces that we describe in Table 3. Similar behaviour is seen when working with the alternate triangulations described by (4.1).
- Although the triangulation \mathcal{T}_{WS} was chosen arbitrarily, in hindsight this was a fortuitous choice. If we attempt to apply the method used in Theorem 1.1 to either of the alternative triangulations described by (4.1), we do not arrive at a conclusive proof.

Specifically, if we (i) eliminate vertex and thin edge links, (ii) eliminate surfaces through heuristic pruning, and then (iii) eliminate surfaces without compatible partners according to Theorem 1.2, some surfaces still remain. For the first alternative we are left with one compatible pair of genus 7 surfaces, and for the second alternative we are left with three compatible pairs of genus 7 surfaces. All of these leftover surfaces are vertex surfaces in both standard and quadrilateral coordinates. They can eventually be eliminated, but only with additional manipulation of the corresponding bounded triangulations.

- It is interesting to compare the relative power and efficiency of Theorem 1.2 and heuristic pruning as individual techniques. Ignoring the vertex link and thin edge links, consider the remaining 1726 vertex surfaces in \mathcal{T}_{WS} . As we have seen already, heuristic pruning alone eliminates 1710 of these 1726 surfaces (around 99%). On the other hand, if we use Theorem 1.2 as a filtering tool (by removing all surfaces with no compatible partner as required by Theorem 1.2), we can eliminate 1227 of these 1726 surfaces (around 71%).

Although this suggests that Theorem 1.2 is less effective as a filtering tool than heuristic pruning, it is significantly faster to use—filtering by Theorem 1.2 takes just a few seconds, whereas running all 1726 surfaces through heuristic pruning takes about 40 minutes (primarily because we must cut along each surface, which can produce triangulations with thousands of tetrahedra to simplify and test). It follows that Theorem 1.2 could be used as a very fast initial filter, leaving a smaller set of surfaces to run through the more expensive heuristic pruning.

A more sophisticated variant of this idea is to repeatedly call upon Theorem 1.2 *throughout* the heuristic pruning process. That is, every time a surface S is eliminated through heuristic pruning, we immediately eliminate every other surface that has S as its only compatible partner. Although this should further improve the efficiency of elimination, the final result (i.e., the set of leftover surfaces) will of course remain the same.

Finally, we note that it is possible to prove the Weber-Seifert dodecahedral space to be non-Haken without employing Theorem 1.2 at all, although the relevant computations require significantly more human intervention. If we begin with the first alternative triangulation of (4.1), we obtain 1909 vertex surfaces. With heuristic pruning this list reduces to just nine surfaces: eight surfaces S_1, \dots, S_8 of genus 7, and one surface S_9 of genus 8.

For each genus 7 surface S_i ($1 \leq i \leq 8$), cutting along S_i gives at least one bounded triangulation \mathcal{T}_i with free fundamental group. By repeatedly applying Pachner moves, we can recognise the underlying manifold as a genus 7 handlebody, showing the original surface S_i to be compressible.

The final surface S_9 is more difficult to deal with. We cut along S_9 to obtain bounded triangulations \mathcal{T}_9 and \mathcal{T}'_9 ; although we are not able to identify either component, with sufficiently many Pachner moves we can nevertheless manufacture a compressing disc in the form described by Lemma 3.1.

REFERENCES

- [1] Joan S. Birman, *Problem list: Nonsufficiently large 3-manifolds*, Notices Amer. Math. Soc. **27** (1980), no. 4, 349.
- [2] Benjamin A. Burton, *Regina: Normal surface and 3-manifold topology software*, <http://regina.sourceforge.net/>, 1999–2009.
- [3] ———, *Face pairing graphs and 3-manifold enumeration*, J. Knot Theory Ramifications **13** (2004), no. 8, 1057–1101. MR2108649 (2005h:57003)
- [4] ———, *Introducing Regina, the 3-manifold topology software*, Experiment. Math. **13** (2004), no. 3, 267–272. MR2103324 (2005g:57042)
- [5] ———, *Converting between quadrilateral and standard solution sets in normal surface theory*, Algebr. Geom. Topol. **9** (2009), no. 4, 2121–2174. MR2551665

- [6] ———, *The complexity of the normal surface solution space*, SCG '10: Proceedings of the Twenty-Sixth Annual Symposium on Computational Geometry, ACM, 2010, pp. 201–209. MR2742952
- [7] ———, *Optimizing the double description method for normal surface enumeration*, Math. Comp. **79** (2010), no. 269, 453–484. MR2552235
- [8] Patrick J. Callahan, Martin V. Hildebrand, and Jeffrey R. Weeks, *A census of cusped hyperbolic 3-manifolds*, Math. Comp. **68** (1999), no. 225, 321–332. MR1620219 (99c:57035)
- [9] Wolfgang Haken, *Theorie der Normalflächen*, Acta Math. **105** (1961), 245–375. MR0141106 (25:4519a)
- [10] William Jaco, David Letscher, and J. Hyam Rubinstein, *Algorithms for essential surfaces in 3-manifolds*, Topology and Geometry: Commemorating SISTAG, Contemporary Mathematics, no. 314, Amer. Math. Soc., Providence, RI, 2002, pp. 107–124. MR1941626 (2003m:57043)
- [11] William Jaco and Ulrich Oertel, *An algorithm to decide if a 3-manifold is a Haken manifold*, Topology **23** (1984), no. 2, 195–209. MR744850 (85j:57014)
- [12] William Jaco, Hyam Rubinstein, and Stephan Tillmann, *Minimal triangulations for an infinite family of lens spaces*, J. Topol. **2** (2009), no. 1, 157–180. MR2499441 (2010b:57016)
- [13] William Jaco and J. Hyam Rubinstein, *0-efficient triangulations of 3-manifolds*, J. Differential Geom. **65** (2003), no. 1, 61–168. MR2057531 (2005d:57034)
- [14] Sergei V. Matveev, *Computer recognition of three-manifolds*, Experiment. Math. **7** (1998), no. 2, 153–161. MR1677162 (2000b:57033)
- [15] P. McMullen, *The maximum numbers of faces of a convex polytope*, Mathematika **17** (1970), 179–184. MR0283691 (44:921)
- [16] Ulrich Oertel, *Homology branched surfaces: Thurston's norm on $H_2(M^3)$* , Low-Dimensional Topology and Kleinian Groups (Coventry/Durham, 1984), London Math. Soc. Lecture Note Ser., vol. 112, Cambridge Univ. Press, Cambridge, 1986, pp. 253–272. MR903869 (89e:57011)
- [17] Udo Pachner, *P.L. homeomorphic manifolds are equivalent by elementary shellings*, European J. Combin. **12** (1991), no. 2, 129–145. MR1095161 (92d:52040)
- [18] Jeffrey L. Tollefson, *Isotopy classes of incompressible surfaces in irreducible 3-manifolds*, Osaka J. Math. **32** (1995), no. 4, 1087–1111. MR1380743 (97f:57023)
- [19] ———, *Normal surface Q-theory*, Pacific J. Math. **183** (1998), no. 2, 359–374. MR1625962 (99c:57047)
- [20] C. Weber and H. Seifert, *Die beiden Dodekaederräume*, Math. Z. **37** (1933), no. 1, 237–253. MR1545392

SCHOOL OF MATHEMATICS AND PHYSICS, THE UNIVERSITY OF QUEENSLAND, BRISBANE QLD 4072, AUSTRALIA

E-mail address: bab@maths.uq.edu.au

DEPARTMENT OF MATHEMATICS AND STATISTICS, THE UNIVERSITY OF MELBOURNE, VIC 3010, AUSTRALIA

E-mail address: rubin@ms.unimelb.edu.au

SCHOOL OF MATHEMATICS AND PHYSICS, THE UNIVERSITY OF QUEENSLAND, BRISBANE QLD 4072, AUSTRALIA

E-mail address: tillmann@maths.uq.edu.au