The Complexity of the Normal Surface Solution Space*

Benjamin A. Burton†
School of Mathematics and Physics
The University of Queensland
Brisbane QLD 4072, Australia
bab@maths.uq.edu.au

ABSTRACT

Normal surface theory is a central tool in algorithmic three-dimensional topology, and the enumeration of vertex normal surfaces is the computational bottleneck in many important algorithms. However, it is not well understood how the number of such surfaces grows in relation to the size of the underlying triangulation. Here we address this problem in both theory and practice. In theory, we tighten the exponential upper bound substantially; furthermore, we construct pathological triangulations that prove an exponential bound to be unavoidable. In practice, we undertake a comprehensive analysis of millions of triangulations and find that in general the number of vertex normal surfaces is remarkably small, with strong evidence that our pathological triangulations may in fact be the worst case scenarios. This analysis is the first of its kind, and the striking behaviour that we observe has important implications for the feasibility of topological algorithms in three dimensions.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*; G.2.1 [**Discrete Mathematics**]: Combinatorics

General Terms

Performance, theory, experimentation

Keywords

Computational topology, 3-manifolds, normal surfaces, vertex enumeration, complexity

*A more detailed version of this paper will be available as *Extreme cases in normal surface enumeration* (in preparation).

†Supported by the Australian Research Council under the Discovery Projects funding scheme (project DP1094516).

1. INTRODUCTION

Geometric topology is an inherently algorithmic subject, asking fundamental questions such as the *homeomorphism problem* (find an algorithm to determine whether two given spaces are topologically equivalent) and the *identification problem* (find an algorithm that can determine the topological name and/or structure of a given space). Three-dimensional topology is of particular interest, since in lower dimensions such problems become trivial [27], and in higher dimensions they become unsolvable [25].

Throughout this paper we restrict our attention to *closed 3-manifolds*. In essence, a closed 3-manifold is a compact 3-dimensional topological space that locally looks like $\mathbb{R}^3$ at every point. Much recent progress has been made on algorithms in 3-manifold topology. For example:

- Rubinstein gave an algorithm in 1992 for recognising the simplest of all closed 3-manifolds, namely the 3-sphere [32, 33]; this algorithm has been refined several times since [8, 19, 34].
- In 1995, Jaco and Tollefson presented an algorithm for breaking a 3-manifold down into a connected sum decomposition (essentially a topological "prime decomposition") [21].
- Perelman's proof of the geometrisation conjecture in 2002 finally resolved the general homeomorphism problem for 3-manifolds, completing a programme initiated decades earlier by pioneers such as Haken [15] and Thurston [36]. The full homeomorphism algorithm is a fusion of diverse and complex components, including both the 3-sphere recognition and connected sum decomposition algorithms above.

A recurring theme in these algorithms (and many others) is that they rely upon *normal surface theory*, a tool that allows us to convert difficult topology problems into simpler linear programming problems. In particular, we can search for an interesting surface within a 3-manifold by (i) constructing a high-dimensional polytope, (ii) enumerating the "admissible" vertices of this polytope, and then (iii) testing each admissible vertex to see whether it encodes the interesting surface that we are searching for.[1]

The concept of an "interesting surface" depends on the application at hand. For instance, in the connected sum

[1]Some other algorithms (such as knot genus [16] and Heegaard genus [24]) replace step (ii) with the more difficult enumeration of a Hilbert basis for a polyhedral cone, yielding what are known as *fundamental surfaces*.

decomposition algorithm we search for embedded spheres within our 3-manifold; in other algorithms we might search for non-trivial embedded discs [14] or embedded incompressible surfaces [17]. However, in all of these applications the high-dimensional polytope and its admissible vertices remain the same. That is, the polytope vertex enumeration problem is a *common component* for all of these topological algorithms and many others besides.

Furthermore, this common vertex enumeration problem is in fact the computational bottleneck for many of these algorithms [8, 10]. It is therefore important to improve the efficiency and understand the complexity of this vertex enumeration problem, since any improvements or results will have a widespread impact on computational 3-manifold topology as a whole. This impact also extends beyond three dimensions—for instance, in *4-manifold topology*, to understand whether a given triangulation represents a 4-manifold we require all of the complex machinery of 3-sphere recognition as discussed above.

In general, polytope vertex enumeration is difficult. The general problem is known to be NP-hard [12, 22], and the range of available algorithms is matched by a range of pathological cases that exploit their weaknesses [2]. However, in our context we have two advantages:

- We are not dealing with an arbitrary polytope, but rather one that derives from the machinery of normal surface theory; this polytope is known as the *projective solution space*. Such polytopes have additional constraints on their dimensions and the equalities and inequalities that define them.

- We do not need to enumerate all vertices of the polytope, but only the *admissible vertices*. These are the vertices that satisfy an additional family of non-linear constraints, known as the *quadrilateral constraints*.

These contextual advantages can be exploited in vertex enumeration algorithms with great success; see [7, 8, 9, 37] for details. Nevertheless, the enumeration problem remains a difficult one. In particular, Agol et al. [1] show that determining knot genus—yet another problem that employs normal surface theory—is in fact NP-complete.

In this paper we concern ourselves with the *complexity* of the enumeration problem. More specifically, we focus on the *number of admissible vertices* of the projective solution space, which we denote by σ. This quantity is important for the following reasons:

- The admissible vertex count σ gives a lower bound for the time complexity of vertex enumeration. Moreover, for the quadrilateral-to-standard conversion algorithm (a key component of the current state-of-the-art enumeration algorithm), there is strong evidence to suggest that the running time is in fact a low-degree polynomial in σ [7].

- Each admissible vertex corresponds to a surface in our 3-manifold upon which we must run some subsequent test. For some problems (such as Hakenness testing [10, 17]) this test is extremely expensive, and so the number of admissible vertices becomes a critical factor in the overall time complexity.

The input for a typical normal surface algorithm is a *3-manifold triangulation*, formed from n tetrahedra by joining their $4n$ faces together in pairs. We call n the *size* of the triangulation; not only does n represent the complexity of the input, but both the dimension and the number of facets of the projective solution space are linear in n.

The growth of σ as a function of n is currently not well understood. The only general theoretical bound in the literature is $\sigma \leq 128^n$, proven by Hass et al. [16]; in the special case of a one-vertex triangulation this has been improved to $\sigma \in O(15^n)$ [9]. Very little is known about the growth of σ in practice, though initial observations suggest that σ is in fact far smaller [7]. For example, in the proof that the Weber-Seifert dodecahedral space is non-Haken (one of the first significant computer proofs to employ normal surface theory), a "typical" triangulation of size $n = 23$ is found to generate just $\sigma = 1751$ admissible vertices [10].

In this paper we shed more light on the growth of σ, including new theoretical bounds and comprehensive practical experimentation. Following a brief outline of normal surface theory in Section 2, we present the following results:

- In Section 3 we show that $\sigma \in O(\phi^{7n})$, where ϕ is the golden ratio $(1+\sqrt{5})/2$. This tightens the general theoretical bound on σ from 128^n to just over $O(29^n)$. We prove this by extending McMullen's upper bound theorem [31] to show that any convex polytope with k facets must have $O(\phi^k)$ vertices.

  We push this bound from the other direction in Section 4 by constructing an infinite family of 3-manifold triangulations for which $\sigma = 17^{n/4} + n/4$. This yields the first known family for which σ is exponential in n, and disproves an earlier conjecture of the author that $\sigma \in O(2^n)$. By extending this family to all $n > 5$ we show that any theoretical upper bound must grow at least as fast as $\Omega(17^{n/4}) \simeq \Omega(2.03^n)$.

- In Section 5 we build a comprehensive census of *all* 3-manifold triangulations of size $n \leq 9$, and measure σ for each of the ~ 150 million triangulations that ensue. We find a remarkably slow growth rate—for $n > 5$ the worst cases are precisely the infinite family above, suggesting that the lower limit of $\Omega(17^{n/4}) \simeq \Omega(2.03^n)$ may in fact be tight. In the average case the mean $\overline{\sigma}$ appears to grow even slower, with an apparent growth rate of less than ϕ^n and a final mean of just $\overline{\sigma} \simeq 78.49$ for $n = 9$.

  This analysis is the first of its kind, primarily because the complex algorithms and software required for such a comprehensive study did not exist until very recently [5, 7]. Previous censuses have focused on restricted classes of triangulations (such as minimal triangulations of irreducible or hyperbolic manifolds [5, 11, 26, 30]), and previous measurements of σ have been for isolated or ad-hoc collections of cases [7, 10, 28].

Throughout this paper we work with Haken's original formulation of normal surface theory [14, 15]. Tollefson defines an alternative formulation called *quadrilateral coordinates* [37], which is only applicable for some problems but where the polytope becomes much simpler. In quadrilateral coordinates an upper bound of $\sigma \leq 4^n$ can be obtained through an analysis of *zero sets* [9], but again the growth rate is found to be significantly slower in practice. We address quadrilateral coordinates in detail in the full version of this paper.

2. PRELIMINARIES

Throughout this paper we assume that we are working with a *3-manifold triangulation of size n*. By this we mean a collection of n tetrahedra, some of whose $4n$ faces are affinely identified (or "glued together") in pairs so that the resulting topological space is a 3-manifold (possibly with boundary). If all $4n$ faces are identified in $2n$ pairs then we obtain a closed 3-manifold; otherwise we obtain a *triangulation with boundary*, and the unidentified faces become *boundary faces*. Unless otherwise specified, all triangulations in this paper are of closed 3-manifolds.

There is no need for a 3-manifold triangulation to be rigidly embedded in some larger space—tetrahedra can be "bent" or "stretched". Moreover, we allow multiple vertices of the same tetrahedron to be identified as a result of our face gluings, and likewise with edges. This allows us to build triangulations using very few tetrahedra, which becomes useful for computation.

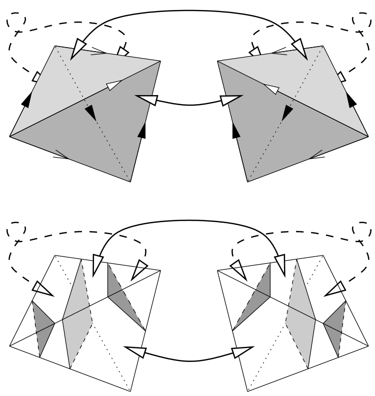

Figure 1: A 3-manifold triangulation and an embedded normal surface

To illustrate, the upper diagram of Figure 1 shows a triangulation of the product space $S^2 \times S^1$ using just $n = 2$ tetrahedra—the back two faces of each tetrahedron are identified with a twist, and the front two faces of the left tetrahedron are identified directly with the front two faces of the right tetrahedron. All eight vertices become identified together, and the 12 edges become identified in three distinct classes (represented in the diagram by three different types of arrowhead). We say that the resulting triangulation has *one vertex* and *three edges*.

Normal surfaces were introduced by Kneser [23], and further developed by Haken [14, 15] for use in algorithms. A *normal surface* is a 2-dimensional surface embedded within a 3-manifold triangulation that meets each tetrahedron in a (possibly empty) collection of *triangles* and/or *quadrilaterals*, as illustrated in Figure 2. For example, a normal surface within our $S^2 \times S^1$ triangulation is shown in the lower diagram of Figure 1; as a consequence of the tetrahedron gluings, the six triangles and quadrilaterals join together to form a 2-dimensional sphere.

There are four distinct *types* of triangle and three distinct *types* of quadrilateral within each tetrahedron (defined by which edges of the tetrahedron they meet). The *vector representation* of a normal surface is a collection of $7n$ integers counting the number of pieces of each type in each tetrahedron; from this vector in $\mathbb{R}^{7n}$ we can completely reconstruct

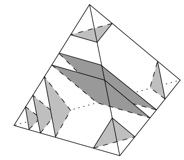

Figure 2: Normal triangles and quadrilaterals within a tetrahedron

the original surface. We treat surfaces and their vectors interchangeably (so, for instance, "adding" two surfaces means adding their two vectors and reconstructing a new surface from the result).

An early result of Haken is a set of necessary and sufficient conditions for a vector to represent a normal surface: (i) all coordinates must be non-negative; (ii) the vector must satisfy a set of linear homogeneous equations (the *matching equations*); and (iii) there can be at most one non-zero quadrilateral coordinate corresponding to each tetrahedron (the *quadrilateral constraints*). Vectors that satisfy all of these conditions are called *admissible*.

Jaco and Oertel [17] define the *projective solution space* to be the polytope in $\mathbb{R}^{7n}$ obtained as a cross-section of the cone defined by (i) and (ii) above. A *vertex normal surface* lies on an extremal ray of this cone and is not a multiple of some smaller surface. The vertex normal surfaces are in bijection with the admissible vertices of the projective solution space; we let σ denote the number of vertex normal surfaces, and we call σ the *admissible vertex count*.

The enumeration of vertex normal surfaces is a critical component—and often the computational bottleneck—for many important topological algorithms. This is because one can often prove that, if an interesting surface exists (such as an incompressible surface or an essential sphere), then one must appear as a vertex normal surface. See Hass et al. [16] for a more detailed introduction to normal surface theory and its role in computational topology.

3. THEORETICAL BOUNDS

As noted in the introduction, the best bound known to date for the admissible vertex count is $\sigma \leq 128^n$, proven by Hass et al. [16]. We begin by tightening this exponential bound as follows:

THEOREM 1. *Let $\phi = (1 + \sqrt{5})/2$. Then the admissible vertex count σ is bounded above by $O(\phi^{7n}) \simeq O(29.03^n)$.*

We prove this through a simple extension of McMullen's upper bound theorem [31]. McMullen gives a tight bound on the number of vertices for a convex polytope with k facets and d dimensions; we extend this here to a loose bound that covers all possible dimensions.

LEMMA 2. *Let $F_0 = 0$, $F_1 = 1$, $F_2 = 1$, ... represent the Fibonacci sequence, where $F_{i+2} = F_{i+1} + F_i$. Then for any $k \geq 3$, a convex polytope with precisely k facets has $\leq F_{k+1}$ vertices.*

PROOF. Suppose the polytope P is d-dimensional with precisely k facets. Then McMullen's theorem (taken in dual

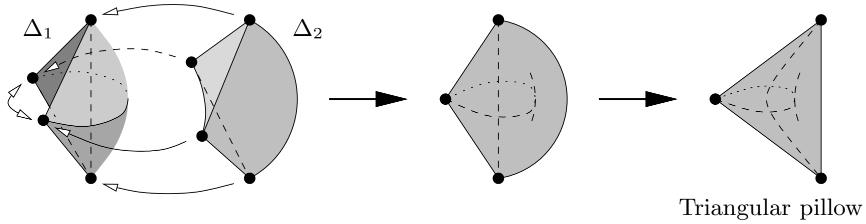

Figure 3: The two-tetrahedron triangular pillow at the centre of a 4-block

form) shows that P has at most

$$\binom{k-\lfloor\frac{d+1}{2}\rfloor}{k-d}+\binom{k-\lfloor\frac{d+2}{2}\rfloor}{k-d}$$
$$=\binom{k-\lfloor\frac{d+1}{2}\rfloor}{d-\lfloor\frac{d+1}{2}\rfloor}+\binom{k-\lfloor\frac{d+2}{2}\rfloor}{d-\lfloor\frac{d+2}{2}\rfloor} \qquad (1)$$

vertices.[2] For even d this can be rewritten as $\binom{k-a}{a}+\binom{(k-2)-b}{b}$ for suitable integers a, b, and for odd d it can be rewritten as $2\binom{(k-1)-a}{a}$ for a suitable integer a.

We now claim that $\binom{k-a}{a} \leq F_k$ for any k, a with $k \geq 1$. This is easily established for $k = 1, 2$, and the full claim follows from the inductive step $\binom{k-a}{a} = \binom{k-1-a}{a}+\binom{k-1-a}{a-1} = \binom{(k-1)-a}{a} + \binom{(k-2)-(a-1)}{a-1} \leq F_{k-1} + F_{k-2} = F_k$.

From here our lemma is straightforward. If d is even then the number of vertices of P is at most $\binom{k-a}{a} + \binom{(k-2)-b}{b} \leq F_k + F_{k-2} \leq F_k + F_{k-1} = F_{k+1}$, and if d is odd then the number of vertices is at most $2\binom{(k-1)-a}{a} \leq 2F_{k-1} \leq F_k + F_{k-1} = F_{k+1}$. □

Unlike McMullen's result, Lemma 2 is not tight. Nevertheless, it gives us a very good[3] asymptotic upper bound of $O(\phi^k)$, which is enough to prove our main theorem.

Proof of Theorem 1. The facets of the projective solution space in $\mathbb{R}^{7n}$ are defined by the $7n$ inequalities $x_1 \geq 0$, …, $x_{7n} \geq 0$, and so there are at most $7n$ facets in total. Lemma 2 then shows that the projective solution space has at most F_{7n+1} vertices, and so $\sigma \leq F_{7n+1}$. Using the standard formula $F_k = \lfloor \phi^k/\sqrt{5} + \frac{1}{2} \rfloor$ it follows that $\sigma \in O(\phi^{7n})$. □

It is interesting to note that Theorem 1 makes no use of admissibility—this suggests that, although the bound of ϕ^{7n} is a strong improvement on 128^n, this bound is still very loose. We confirm this through experimentation in Section 5. Although we only consider closed 3-manifolds in this paper, it should be noted that Theorem 1 and its proof apply equally well to triangulations with boundary, and also to the *ideal triangulations* of Thurston [35].

4. EXTREME CASES

Having tightened the upper bound from above, we now turn our attention to limiting the upper bound from below.

[2]This is the number of facets of the cyclic d-dimensional polytope with k vertices [13].

[3]Experimentation shows that this asymptotic upper bound of $\phi^k \simeq 1.618^k$ is close to optimal. If we maximise equation (1) over all d for each $k = 100, \ldots, 200$, the maximum grows at a rate of approximately 1.613^k.

We do this by building pathological triangulations for which $\sigma \in \Theta(17^{n/4}) \simeq \Theta(2.03^n)$. This growth rate shows that an exponential upper bound on σ is unavoidable, and furthermore disproves an earlier conjecture of the author that $\sigma \in O(2^n)$.

We begin by describing 4-blocks, which are small building blocks that appear repeatedly throughout our triangulations. Using these building blocks, we then construct the family of pathological triangulations $\mathcal{X}_1, \mathcal{X}_2, \ldots$.

Definition (4-block). A *4-block* is a triangulation with boundary, built from the four tetrahedra $\Delta_1, \Delta_2, \Delta_3, \Delta_4$ using the following construction.

We begin by folding together two faces of Δ_1, and then wrapping Δ_2 around the remaining two faces as illustrated in Figure 3. This forms a *triangular pillow* with three vertices, three boundary edges, two internal edges, and two boundary faces.

Next we fold together two faces of Δ_3 and two faces of Δ_4, as illustrated in the leftmost column of Figure 4. To finish, we join the pillow to both Δ_3 and Δ_4 as illustrated in the central column of Figure 4—the upper face $A_1B_1A_2$ of the pillow is glued to the lower face $A_3B_2A_3$ of Δ_3, and the lower face $A_1B_1A_2$ of the pillow is glued to the upper face $A_4B_3A_4$ of Δ_4.

The final result is shown in the rightmost column of Figure 4, with three boundary vertices and one internal vertex. The triangular pillow is buried in the middle of this structure, wrapped around the internal vertex; for simplicity the two edges inside the pillow are not shown.

Definition (Pathological triangulation $\mathcal{X}_k$). For each integer $k \geq 1$, the *pathological triangulation* $\mathcal{X}_k$ is constructed from $n = 4k$ tetrahedra in the following manner.

From these $4k$ tetrahedra we build k distinct 4-blocks, labelled $\mathcal{B}_1, \ldots, \mathcal{B}_k$. Within each 4-block $\mathcal{B}_i$ we label the three boundary vertices P_i, Q_i, R_i, where P_i sits between both boundary triangles as illustrated in Figure 5.

For each $i = 1, \ldots, k$ we join blocks $\mathcal{B}_i$ and $\mathcal{B}_{i+1}$ as follows (where $\mathcal{B}_{k+1}$ is taken to mean $\mathcal{B}_1$). Triangle $P_iP_iR_i$ is joined to triangle $Q_{i+1}P_{i+1}P_{i+1}$; note that this is "twisted", not a direct gluing, since it maps $P_i \leftrightarrow Q_{i+1}$ and $P_{i+1} \leftrightarrow R_i$. There are in fact two ways this gluing can be performed (one a reflection of the other); we resolve this ambiguity by orienting each block consistently, and then choosing the gluing that preserves orientation.

An effect of these gluings is to identify all of the P_i, Q_i and R_i to a single vertex, so that $\mathcal{X}_k$ has $k + 1$ vertices in total (counting also the k internal vertices from each original block).

It is not clear that each $\mathcal{X}_k$ is a 3-manifold triangulation (in particular, that $\mathcal{X}_k$ looks like $\mathbb{R}^3$ in the vicinity of each

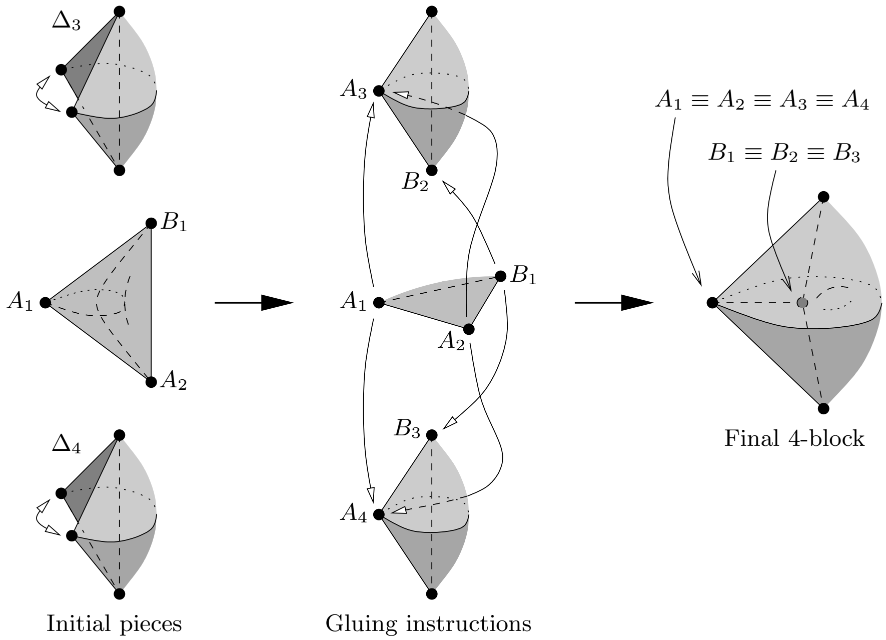

Figure 4: Building a 4-block from two tetrahedra and a triangular pillow

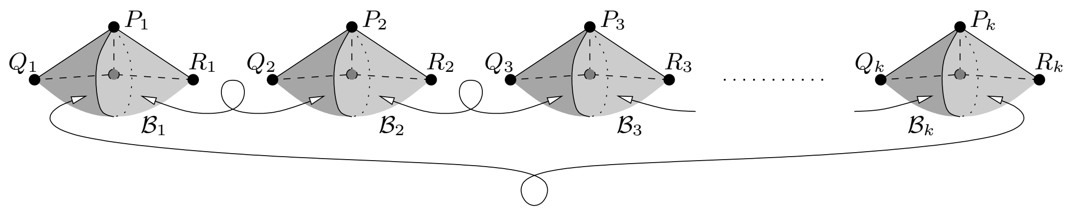

Figure 5: Building the pathological triangulation $\mathcal{X}_k$ from k distinct 4-blocks

vertex). The following sequence of results proves this by showing that every $\mathcal{X}_k$ is in fact a triangulation of the 3-sphere.

LEMMA 3. *A 4-block is a triangulation of the 3-ball (i.e., the solid 3-dimensional ball), with a boundary consisting of two triangles in the formation shown in Figure 6.*

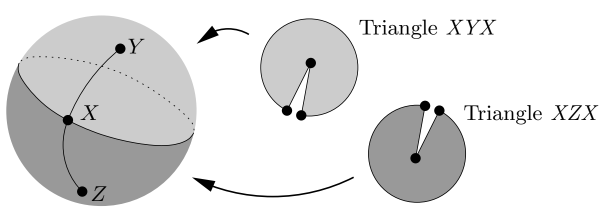

Figure 6: A 3-ball whose boundary consists of two triangles

PROOF. This is evident from the construction in Figure 4. It can also be verified computationally using the software package *Regina* [3], which implements 3-sphere and 3-ball recognition [4]. □

LEMMA 4. *Let $\mathcal{T}_1$ and $\mathcal{T}_2$ each be triangulations of the 3-ball with boundaries in the formation shown in Figure 6. If we identify one boundary triangle of $\mathcal{T}_1$ with one boundary triangle of $\mathcal{T}_2$ under any of the six possible identifications, the result is always another triangulation of the 3-ball with boundary in the formation shown in Figure 6.*

LEMMA 5. *Let $\mathcal{T}$ be a triangulation of the 3-ball with boundary in the formation shown in Figure 6. If we identify the two boundary triangles under any of the three possible orientation-preserving identifications, the result is always a closed 3-manifold triangulation of the 3-sphere.*

PROOF. Both of these results are essentially properties of 3-manifolds, not their underlying triangulations—if they hold for any selection of triangulations $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}$ then they must hold for all such selections. We verify these results using *Regina* by choosing 4-blocks for our triangulations and testing all six/three possible identifications. □

Since each $\mathcal{X}_k$ is built by joining together 4-blocks along boundary triangles in an orientation-preserving fashion, the following result follows immediately from Lemmata 3–5.

COROLLARY 6. *For each $k \geq 1$, $\mathcal{X}_k$ is a closed 3-manifold triangulation of the 3-sphere.*

We turn our attention now to counting the vertex normal surfaces for each triangulation $\mathcal{X}_k$. Recalling that $k = n/4$, the following result shows that for these pathological triangulations we have $\sigma \in \Theta(17^{n/4}) \simeq \Theta(2.03^n)$.

Lemma 7. *For each $k \geq 1$, $\mathcal{X}_k$ has precisely $\sigma = 17^k + k$ vertex normal surfaces.*

Proof. Consider a single 4-block with boundary vertices labelled P, Q, R as before, and let S denote the internal vertex. Define α, β and γ to be small loops on the 4-block boundary surrounding P, Q and R respectively, as illustrated in Figure 7.

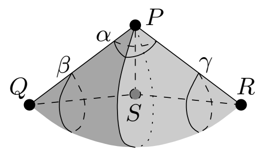

Figure 7: The curves α, β, γ on the boundary of a 4-block

Using the software package *Regina*, we can construct the projective solution space for this 4-block. There are 17 admissible vertices in total, corresponding to 17 vertex normal surfaces: one with empty boundary, and 16 whose boundary consists of some combination of α, β and γ. These surfaces are summarised in Table 1, and we label them $\mathbf{a}, \mathbf{b}, \ldots, \mathbf{q}$ as shown.

It is important to note that $\mathbf{a}, \mathbf{b}, \ldots, \mathbf{q}$ are all *compatible*; that is, no combination of their vectors can ever violate the quadrilateral constraints.[4] This is an unusual but extremely helpful state of affairs, since we can effectively ignore the quadrilateral constraints from here onwards.

Now consider the entire set of 4-blocks $\mathcal{B}_1, \ldots, \mathcal{B}_k$. Let $\mathbf{a}_i, \mathbf{b}_i, \ldots, \mathbf{q}_i$ denote the corresponding surfaces in $\mathcal{B}_i$, and let $\alpha_i, \beta_i, \gamma_i$ denote the corresponding boundary curves. Any normal surface in $\mathcal{X}_k$ is a union of normal surfaces in the individual blocks $\mathcal{B}_1, \ldots, \mathcal{B}_k$, and hence can be expressed as

$$(\lambda_{1,1}\, \mathbf{a}_1 + \ldots + \lambda_{1,17}\, \mathbf{q}_1) + \ldots + (\lambda_{k,1}\, \mathbf{a}_k + \ldots + \lambda_{k,17}\, \mathbf{q}_k)$$

for some family of constants $\lambda_{1,1}, \ldots, \lambda_{k,17} \geq 0$. In this form, it can be shown[5] that the matching equations for $\mathcal{X}_k$ reduce to the following statement:

> There is some non-negative $\mu \in \mathbb{R}$ such that, for every i, the sum $\lambda_{i,1}\, \mathbf{a}_i + \ldots + \lambda_{i,17}\, \mathbf{q}_i$ has boundary $\mu\alpha_i + \mu\beta_i + \mu\gamma_i$.

In other words, the portion of the overall surface within each 4-block $\mathcal{B}_i$ must have boundary $\mu\alpha_i + \mu\beta_i + \mu\gamma_i$, where μ is independent of i.

Return now to a single 4-block with admissible vertices $\mathbf{a}, \ldots, \mathbf{q}$, and let $\lambda_1\, \mathbf{a} + \ldots + \lambda_{17}\, \mathbf{q}$ be some point in the projective solution space for this 4-block. We can ensure that the corresponding surface has boundary of the form $\mu\alpha + \mu\beta + \mu\gamma$ by imposing the linear constraints depicted in Table 2 (where each line of these constraints corresponds to a section of the *Boundary* column in Table 1).

This has the effect of intersecting the original projective solution space for the 4-block with two new hyperplanes. A standard application of the filtered double description method [9] shows that the resulting polytope has 18 admissible vertices, described by the following 18 normal surfaces: the original $\mathbf{a}$ with no boundary, and 17 new surfaces[6] all with boundary $\alpha + \beta + \gamma$. Within each block $\mathcal{B}_i$, we label these 17 new surfaces $\mathbf{v}_{i,1}, \ldots, \mathbf{v}_{i,17}$.

Given the formulation of the matching equations above, it follows that the normal surfaces in $\mathcal{X}_k$ are described completely by the linear combinations

$$\rho_{1,1}\, \mathbf{v}_{1,1} + \ldots + \rho_{k,17}\, \mathbf{v}_{k,17} + \eta_1 \mathbf{a}_1 + \ldots + \eta_k \mathbf{a}_k,$$

where each $\rho_{i,j}, \eta_i \geq 0$ and where $\sum_j \rho_{1,j} = \sum_j \rho_{2,j} = \ldots = \sum_j \rho_{k,j}$. The full projective solution space for $\mathcal{X}_k$ therefore has $17^k + k$ admissible vertices, corresponding to the k surfaces $\mathbf{a}_1, \ldots, \mathbf{a}_k$ and the 17^k combinations $\mathbf{v}_{1,j_1} + \mathbf{v}_{2,j_2} + \ldots + \mathbf{v}_{k,j_k}$ for $j_1, j_2, \ldots, j_k \in \{1, \ldots, 17\}$. $\square$

The pathological triangulations $\mathcal{X}_1, \mathcal{X}_2, \ldots$ cover all sizes of the form $n = 4k$. We can generalise this construction to include $n = 4k + 1$, $4k + 2$ and $4k + 3$ by replacing one of our 4-blocks with a single "exceptional" block. The general constructions and analyses are detailed in the full version of this paper, and the results are summarised in the following theorem.

Theorem 8. *For every positive $n \neq 1, 2, 3, 5$, there exists a closed 3-manifold triangulation of size n whose admissible vertex count is as follows:*

$$\begin{array}{lll} n = 4k & (k \geq 1) & \Longrightarrow \quad \sigma = 17^k + k \\ n = 4k+1 & (k \geq 2) & \Longrightarrow \quad \sigma = 581 \cdot 17^{k-2} + k + 1 \\ n = 4k+2 & (k \geq 1) & \Longrightarrow \quad \sigma = 69 \cdot 17^{k-1} + k \\ n = 4k+3 & (k \geq 1) & \Longrightarrow \quad \sigma = 141 \cdot 17^{k-1} + k + 2 \end{array} \qquad (2)$$

Lemma 7 proves this result for the first case $n = 4k$. For an extra measure of verification, equation (2) has been confirmed numerically for all $n \leq 14$ by building the relevant triangulations and using *Regina* to enumerate all vertex normal surfaces.

The main result of this section is the following limit on any upper bound for σ, which follows immediately from Theorem 8. Moreover, as we discover in the following section, there is reason to believe that this may in fact give the tightest possible asymptotic bound.

Corollary 9. *Any upper bound for the admissible vertex count σ must grow at a rate of at least $\Omega(17^{n/4}) \simeq \Omega(2.03^n)$.*

5. PRACTICAL GROWTH

We turn now to a comprehensive study of the admissible vertex count σ for real 3-manifold triangulations. The basis of this study is a complete census of *all* closed 3-manifold triangulations of size $n \leq 9$. This is a significant undertaking, and such a census has never been compiled before; the paper [5] details some of the sophisticated algorithms involved.

The result is a collection of 149 676 922 triangulations, each counted once up to *isomorphism* (a relabelling of tetrahedra and their vertices). It is worth noting that within this large collection of triangulations there is a much smaller

[4]This is because, within each tetrahedron, we observe that two of the three quadrilateral types never appear *anywhere* amongst the surfaces $\mathbf{a}, \mathbf{b}, \ldots, \mathbf{q}$.

[5]The argument uses the facts that curves $\alpha_i, \beta_i, \gamma_i$ surround vertices P_i, Q_i, R_i respectively, and that all of these vertices are identified together in the overall triangulation $\mathcal{X}_k$.

[6]These are the six surfaces $(\mathbf{c}+\mathbf{g},\ \mathbf{d}+\mathbf{f},\ \text{or}\ \mathbf{j}) + (\mathbf{b}\ \text{or}\ \mathbf{l})$, the five surfaces $\mathbf{c} + \mathbf{d} + (\mathbf{b},\ \mathbf{e},\ \mathbf{l},\ \mathbf{m},\ \text{or}\ \mathbf{n})$, and the six surfaces $\mathbf{c} + \mathbf{i}$, $\mathbf{c} + \mathbf{p}$, $\mathbf{d} + \mathbf{h}$, $\mathbf{d} + \mathbf{o}$, $\mathbf{k}$ and $\mathbf{q}$.

Label	Boundary	Description
a	—	Small sphere around internal vertex S
b	α	Small disc around boundary vertex P
c	β	Small disc around boundary vertex Q
d	γ	Small disc around boundary vertex R
e	α	Tube from P to S, closed around S
f	β	Tube from Q to S, closed around S
g	γ	Tube from R to S, closed around S
h	$\alpha+\beta$	Tube from P to Q via S, open at both ends
i	$\alpha+\gamma$	Tube from P to R via S, open at both ends
j	$\beta+\gamma$	Tube from Q to R via S, open at both ends
k	$\alpha+\beta+\gamma$	Forked tube joining all of P, Q, R via S, open at all three ends
l	α	Surface **b** with large "balloon" disc attached inside the pillow
m	α	Surface **b** with punctured torus attached inside the pillow
n	α	Surface **e** with punctured torus attached inside the pillow
o	$\alpha+\beta$	Surface **h** with punctured torus attached inside the pillow
p	$\alpha+\gamma$	Surface **i** with punctured torus attached inside the pillow
q	$\alpha+\beta+\gamma$	Surface **k** with punctured torus attached inside the pillow

Table 1: The 17 vertex normal surfaces within a 4-block

$$\begin{array}{rl} & \lambda_2 + \lambda_5 + \lambda_8 + \lambda_9 + \lambda_{11} + \lambda_{12} + \lambda_{13} + \lambda_{14} + \lambda_{15} + \lambda_{16} + \lambda_{17} \\ = & \lambda_3 + \lambda_6 + \lambda_8 + \lambda_{10} + \lambda_{11} + \lambda_{15} + \lambda_{17} \\ = & \lambda_4 + \lambda_7 + \lambda_9 + \lambda_{10} + \lambda_{11} + \lambda_{16} + \lambda_{17} \end{array}$$

Table 2: Ensuring a boundary of the form $\mu\alpha + \mu\beta + \mu\gamma$ within a single 4-block

number of distinct 3-manifolds, as indicated by the 3-manifold census data of Martelli and Petronio [26] and the author [5].

For each of these ~ 150 million triangulations we enumerate all vertex normal surfaces using the algorithms described in [7, 9]. The resulting admissible vertex counts σ are summarised in Table 3. All computations were performed using the software package *Regina* [3, 4].

Size (n)	Number of triangulations	Admissible vertex count (σ) Mean	Std dev	Min	Max
1	4	2.00	0.71	1	3
2	17	3.94	1.39	2	7
3	81	5.49	1.97	2	11
4	577	8.80	3.38	2	18
5	5 184	13.34	5.49	4	36
6	57 753	20.76	9.21	4	70
7	722 765	32.17	15.29	4	144
8	9 787 509	50.20	25.52	4	291
9	139 103 032	78.49	42.51	4	584

Table 3: Summary of admissible vertex counts for all triangulations ($n \leq 9$)

The figures that we see are remarkably small. For $n = 9$ tetrahedra, although Theorem 1 places the theoretical bound at $\simeq O(29^n)$, we have just 584 vertex normal surfaces in the worst case. The mean admissible vertex count for $n = 9$ is much smaller again, evaluated at just 78.49. The full distribution of all admissible vertex counts for $n = 9$ is shown in the upper graph of Figure 8.

Indeed, our pathological triangulations $\mathcal{X}_1, \mathcal{X}_2$ represent the worst cases for $n = 4, 8$ respectively, giving the maximum observed values of $\sigma = 17^1 + 1 = 18$ and $\sigma = 17^2 + 2 = 291$. More generally, the pathological triangulations of Theorem 8 give the maximum cases in our census wherever they are defined (i.e., $n \neq 1, 2, 3, 5$). This leads us to the following general conjecture:

Conjecture 1. *For every positive $n \neq 1, 2, 3, 5$, equation (2) gives a tight upper bound on the admissible vertex count σ. As a consequence, we have $\sigma \in O(17^{n/4})$.*

The growth rate of σ for $n = 1, \ldots, 9$ is illustrated in the lower graph of Figure 8 (note that the vertical axis is plotted on a log scale). The growth rate of the maximum σ is roughly $17^{n/4} \simeq 2.03^n$ as suggested above; the growth rate of the average $\overline{\sigma}$ is in the range 1.5^n to 1.6^n. This is just below the Fibonacci growth rate of $\phi^n \simeq 1.62^n$. Indeed, if we let $\overline{\sigma}_n$ denote the mean admissible vertex count amongst all triangulations of size n, we find that $\overline{\sigma}_n < \overline{\sigma}_{n-1} + \overline{\sigma}_{n-2}$ throughout our census. This leads us to our next general conjecture:

Conjecture 2. *For every $n \geq 3$, the mean admissible vertex count $\overline{\sigma}_n$ satisfies the relation $\overline{\sigma}_n < \overline{\sigma}_{n-1} + \overline{\sigma}_{n-2}$. As a consequence, $\overline{\sigma}_n$ is bounded above by $O(\phi^n)$ where $\phi = (1 + \sqrt{5})/2$.*

In particular, our census analysis gives us the following computational result:

Theorem 10. *Conjectures 1 and 2 are true for $n \leq 9$.*

6. CONCLUSIONS

We have pushed the theoretical bounds on the admissible vertex count σ from both directions, and we have shown through an exhaustive study of ~ 150 million triangulations that σ is surprisingly small in practice. We close with a brief discussion of the implications of this study.

Most importantly, it suggests that topological algorithms that employ normal surfaces might not be as infeasible as theory suggests. Hints of this have already been seen with the quadrilateral-to-standard conversion algorithm for normal surfaces [7], which (against theoretical expectations) appears to have a running time polynomial in its output size.

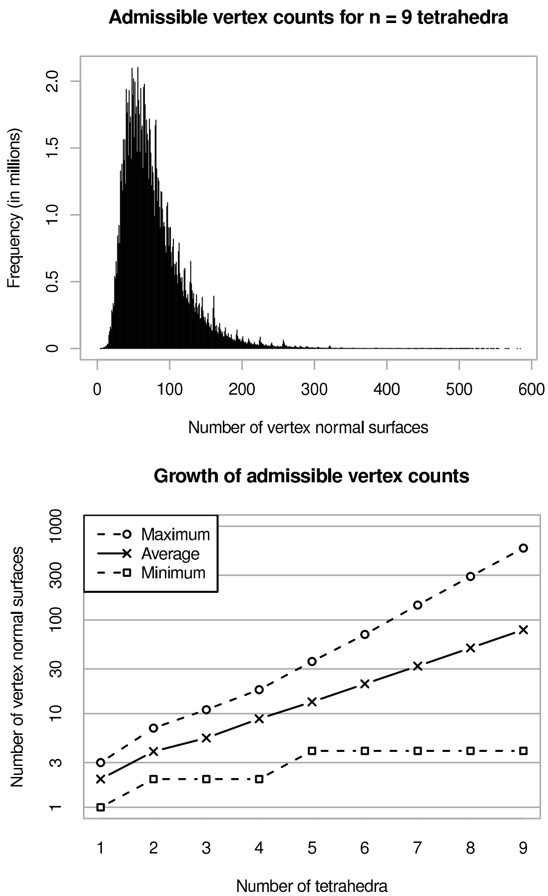

Figure 8: Aggregate results for admissible vertex counts

In many fields, a census for size $n \leq 9$ might not seem large enough for drawing conclusions and conjectures. However, there is evidence elsewhere to suggest that 3-manifold triangulations are flexible enough for important patterns to establish themselves for very low n. For example, the papers [6, 29] discuss several combinatorial patterns for $n \leq 6$; these patterns have later been found to generalise well for larger n [5, 26], and some are now proven in general [18, 20].

Finally, it is clear from this practical study that the theoretical bounds on σ still have much room for improvement. One possible direction is to incorporate the quadrilateral constraints directly into McMullen's theorem. This is difficult because the quadrilateral constraints break convexity, but the outcome may be significantly closer to the $O(17^{n/4})$ that we see in practice.

7. ACKNOWLEDGMENTS

The author is grateful to both the University of Victoria (Canada) and the Victorian Partnership for Advanced Computing (Australia) for the use of their excellent computing resources.

8. REFERENCES

[1] I. Agol, J. Hass, and W. Thurston. 3-manifold knot genus is NP-complete. In *STOC '02: Proceedings of the Thiry-Fourth Annual ACM Symposium on Theory of Computing*, pages 761–766. ACM Press, 2002.

[2] D. Avis, D. Bremner, and R. Seidel. How good are convex hull algorithms? *Comput. Geom.*, 7(5-6):265–301, 1997.

[3] B. A. Burton. Regina: Normal surface and 3-manifold topology software. `http://regina.sourceforge.net/`, 1999–2009.

[4] B. A. Burton. Introducing Regina, the 3-manifold topology software. *Experiment. Math.*, 13(3):267–272, 2004.

[5] B. A. Burton. Enumeration of non-orientable 3-manifolds using face-pairing graphs and union-find. *Discrete Comput. Geom.*, 38(3):527–571, 2007.

[6] B. A. Burton. Structures of small closed non-orientable 3-manifold triangulations. *J. Knot Theory Ramifications*, 16(5):545–574, 2007.

[7] B. A. Burton. Converting between quadrilateral and standard solution sets in normal surface theory. *Algebr. Geom. Topol.*, 9(4):2121–2174, 2009.

[8] B. A. Burton. Quadrilateral-octagon coordinates for almost normal surfaces. To appear in Experiment. Math., `arXiv:0904.3041`, Apr. 2009.

[9] B. A. Burton. Optimizing the double description method for normal surface enumeration. *Math. Comp.*, 79(269):453–484, 2010.

[10] B. A. Burton, J. H. Rubinstein, and S. Tillmann. The Weber-Seifert dodecahedral space is non-Haken. Preprint, `arXiv:0909.4625`, Sept. 2009.

[11] P. J. Callahan, M. V. Hildebrand, and J. R. Weeks. A census of cusped hyperbolic 3-manifolds. *Math. Comp.*, 68(225):321–332, 1999.

[12] M. E. Dyer. The complexity of vertex enumeration methods. *Math. Oper. Res.*, 8(3):381–402, 1983.

[13] B. Grünbaum. *Convex Polytopes*. Number 221 in Graduate Texts in Mathematics. Springer, New York, 2nd edition, 2003.

[14] W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961.

[15] W. Haken. Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I. *Math. Z.*, 80:89–120, 1962.

[16] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. Assoc. Comput. Mach.*, 46(2):185–211, 1999.

[17] W. Jaco and U. Oertel. An algorithm to decide if a 3-manifold is a Haken manifold. *Topology*, 23(2):195–209, 1984.

[18] W. Jaco, H. Rubinstein, and S. Tillmann. Minimal triangulations for an infinite family of lens spaces. *J. Topol.*, 2(1):157–180, 2009.

[19] W. Jaco and J. H. Rubinstein. 0-efficient triangulations of 3-manifolds. *J. Differential Geom.*, 65(1):61–168, 2003.

[20] W. Jaco, J. H. Rubinstein, and S. Tillmann. Coverings and minimal triangulations of 3-manifolds. To appear in Algebr. Geom. Topol., `arXiv:0903.0112`, Feb. 2009.

[21] W. Jaco and J. L. Tollefson. Algorithms for the complete decomposition of a closed 3-manifold. *Illinois J. Math.*, 39(3):358–406, 1995.

[22] L. Khachiyan, E. Boros, K. Borys, K. Elbassioni, and

V. Gurvich. Generating all vertices of a polyhedron is hard. *Discrete Comput. Geom.*, 39(1-3):174–190, 2008.

[23] H. Kneser. Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten. *Jahresbericht der Deut. Math. Verein.*, 38:248–260, 1929.

[24] T. Li. An algorithm to determine the Heegaard genus of a 3-manifold. Preprint, `arXiv:1002.1958`, Feb. 2010.

[25] A. A. Markov. Insolubility of the problem of homeomorphy. In *Proc. Internat. Congress Math. 1958*, pages 300–306. Cambridge Univ. Press, New York, 1960.

[26] B. Martelli and C. Petronio. Three-manifolds having complexity at most 9. *Experiment. Math.*, 10(2):207–236, 2001.

[27] W. S. Massey. *A Basic Course in Algebraic Topology.* Number 127 in Graduate Texts in Mathematics. Springer-Verlag, New York, 1991.

[28] S. Matsumoto and R. Rannard. The regular projective solution space of the figure-eight knot complement. *Experiment. Math.*, 9(2):221–234, 2000.

[29] S. V. Matveev. Tables of 3-manifolds up to complexity 6. *Max-Planck-Institut für Mathematik Preprint Series*, (67), 1998. Available from `http://www.mpim-bonn.mpg.de/html/preprints/preprints.html`.

[30] S. V. Matveev. Recognition and tabulation of three-dimensional manifolds. *Dokl. Akad. Nauk*, 400(1):26–28, 2005.

[31] P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[32] J. H. Rubinstein. An algorithm to recognize the 3-sphere. In *Proceedings of the International Congress of Mathematicians (Zürich, 1994)*, volume 1, pages 601–611. Birkhäuser, 1995.

[33] J. H. Rubinstein. Polyhedral minimal surfaces, Heegaard splittings and decision problems for 3-dimensional manifolds. In *Geometric Topology (Athens, GA, 1993)*, volume 2 of *AMS/IP Stud. Adv. Math.*, pages 1–20. Amer. Math. Soc., 1997.

[34] A. Thompson. Thin position and the recognition problem for S^3. *Math. Res. Lett.*, 1(5):613–630, 1994.

[35] W. P. Thurston. The geometry and topology of 3-manifolds. Lecture notes, Princeton University, 1978.

[36] W. P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.

[37] J. L. Tollefson. Normal surface Q-theory. *Pacific J. Math.*, 183(2):359–374, 1998.