Computational topology and normal surfaces: Theoretical and experimental complexity bounds*

Benjamin A. Burton† João Paixão‡ Jonathan Spreer§

Abstract

In three-dimensional computational topology, the theory of normal surfaces is a tool of great theoretical and practical significance. Although this theory typically leads to exponential time algorithms, very little is known about how these algorithms perform in "typical" scenarios, or how far the best known theoretical bounds are from the real worst-case scenarios. Here we study the combinatorial and algebraic complexity of normal surfaces from both the theoretical and experimental viewpoints. Theoretically, we obtain new exponential lower bounds on the worst-case complexities in a variety of settings that are important for practical computation. Experimentally, we study the worst-case and average-case complexities over a comprehensive body of roughly three billion input triangulations. Many of our lower bounds are the first known exponential lower bounds in these settings, and experimental evidence suggests that many of our theoretical lower bounds on worst-case growth rates may indeed be asymptotically tight.

1 Introduction

In three-dimensional computational topology, many important problems are solved by exponential-time algorithms: key examples include Haken's algorithm for recognising the unknot [14], or breaking down a triangulated 3-manifold into its prime decomposition [18, 19]. This is in contrast to two dimensions in which many problems are solved in polynomial time, or higher dimensions in which important topological problems can become undecidable [11, 24].

A common feature of many such three-dimensional algorithms, and the source of both their solvability and their exponential running times, is their use of *normal surfaces*. In essence, normal surfaces are embedded 2-dimensional surfaces that intersect the surrounding 3-dimensional triangulation in a simple fashion. Most importantly, they describe topological features using combinatorial data, and are thereby well-suited for algorithmic enumeration and analysis.

Amongst the most important normal surfaces are the *vertex normal surfaces*. These correspond to the vertices of a high-dimensional polytope (called the *projective solution space*), and together they generate the space of all possible normal surfaces within the input triangulation. Many topological algorithms begin by enumerating all vertex normal surfaces in the input triangulation, and for many problems (such as unknot recognition and prime decomposition) this enumeration is in fact the main bottleneck for the entire algorithm.

*The first author is grateful to the Australian Research Council for their support under the Discovery Projects funding scheme (projects DP1094516 and DP110101104). This work was done while the second author was visiting The University of Queensland and he is thankful for the funding by FAPERJ for this visit. Computational resources used in this work were provided by the Queensland Cyber Infrastructure Foundation.

†School of Mathematics and Physics, The University of Queensland, Brisbane, Australia. bab@maths.uq.edu.au

‡Department of Mathematics, Pontifícia Universidade Católica, Rio de Janeiro, Brazil, jpaixao@mat.puc-rio.br

§School of Mathematics and Physics, The University of Queensland, Brisbane, Australia. j.spreer@uq.edu.au

One remarkable feature of many algorithms in three-dimensional computational topology is that, although they have extremely large theoretical worst-case complexity bounds, they appear to be *much* easier to solve in practice than these bounds suggest. For example:

- In 1980, Thurston asked if the Weber-Seifert dodecahedral space is Haken (the precise meaning of this is not important here) [1]. This long-standing question became a symbolic benchmark for computational topology, and was only resolved by computer proof after 30 years [10]. At the heart of the proof was an enumeration of all vertex normal surfaces in an $n = 23$-tetrahedron triangulation: despite a prohibitive $O(16^n \times \text{poly}(n))$-time enumeration algorithm (the best available at the time) and a best known bound of $O(3.303^n)$ vertex normal surfaces [5], the enumeration ran in just $5\frac{1}{2}$ hours with only 1751 vertex normal surfaces in total.

- The problem of unknot recognition is of particular interest. Modern derivatives of Haken's original algorithm [19] have an exponential time complexity [16], but there is a growing discussion as to whether a faster algorithm might exist [12, 15]. Certainly unknot recognition lies in **NP** [16], and also **co-NP** if we assume the generalised Riemann hypothesis [21]; moreover, recent algorithmic developments based on linear programming now exhibit an experimental polynomial-time behaviour [6]. Deciding whether unknot recognition has a worst-case polynomial-time solution is now a major open problem in computational topology.

This severe gap between theory and practice is still poorly understood. There appear to be two causes: (i) the best theoretical complexity bounds are far from tight; (ii) "pathological" inputs that exhibit high-complexity behaviour are rare, with "typical" inputs often far easier to work with.

Proving such claims mathematically remains extremely elusive. Obtaining tight complexity bounds requires a deep interaction between topology, normal surfaces and polytope theory, and it is difficult to avoid making very loose estimates in at least one of these areas. Understanding "typical" behaviour (such as average- or generic-case complexity) is hampered by our very limited understanding of random 3-manifold triangulations: even the simple task of generating a random 3-manifold triangulation with n tetrahedra has no known sub-exponential-time solution [27]. In this setting, experimental work plays a crucial role in understanding the realistic performance of algorithms, as well as the innate difficulty of the problems that they aim to solve.

In this paper we focus our attention on the problem of enumerating all vertex normal surfaces within a given n-tetrahedron input triangulation: as mentioned earlier, this is a central component—and often the main bottleneck—of many algorithms in computational 3-manifold topology. Enumeration algorithms are still evolving [6, 8], and they are often hand-tailored to a particular topological problem of interest. For this reason we do not focus on the complexity of any specific algorithm, but instead we study two aspects of normal surface theory that affect and constrain all of these algorithms:

- *Combinatorial complexity:* We study the *total number* of vertex normal surfaces within the input triangulation $\mathcal{T}$, which we denote by $\sigma(\mathcal{T})$. This is our main quantity of interest. It yields an immediate *lower bound* for the time complexity of any enumeration algorithm, since it determines the output size.[1] Moreover, $\sigma(\mathcal{T})$ also factors into *upper bounds*, since modern enumeration algorithms are designed to run faster in situations where $\sigma(\mathcal{T})$ is small [8].[2]

- *Algebraic complexity:* As detailed in Section 2, each normal surface is described by a non-negative integer vector in $\mathbb{R}^{7n}$ (or in some settings, $\mathbb{R}^{3n}$). We investigate the *maximum coordinate* of this vector of any vertex normal surface within the input triangulation $\mathcal{T}$, which

[1] Specifically, since each vertex normal surface can be described in $O(n)$ space [16], the output size is $O(\sigma(\mathcal{T}) \times n)$.

[2] The tree traversal enumeration algorithm (the current state of the art) has running time $O(4^n \sigma(\mathcal{T}) \times \text{poly}(n))$.

we denote by $\kappa(\mathcal{T})$. This quantity is important for the *implementation* of enumeration algorithms, since it affects whether we can work with fast native machine integer types or whether we must fall back to significantly more expensive arbitrary-precision integer arithmetic [8]. Moreover, $\kappa(\mathcal{T})$ features in algorithms that extend or even avoid the enumeration problem:

- Some algorithms, such as recognising small Seifert fibred spaces [28], require the complete enumeration of not just vertex normal surfaces but a much larger "lattice" of normal surfaces whose size is a function of $\kappa(\mathcal{T})$.
- Some algorithms, such as determining the crosscap number of a knot [9], avoid vertex enumeration entirely by solving an integer program instead; here the bounds on $\kappa(\mathcal{T})$ feature as coefficients in the integer program, and directly affect whether the program can be solved using off-the-shelf integer programming software.

In summary, by focusing our attention on the quantities $\sigma(\mathcal{T})$ and $\kappa(\mathcal{T})$, we learn not only about the behaviour of current enumeration algorithms, but also about the intrinsic limits and behaviour of the problem that they seek to solve.

We approach these combinatorial and algebraic quantities $\sigma(\mathcal{T})$ and $\kappa(\mathcal{T})$ through both theory and experiment. Theoretically, we construct infinite "pathological" families of triangulations in Section 3 that establish *exponential lower bounds* on the worst-case scenario for both $\sigma(\mathcal{T})$ and $\kappa(\mathcal{T})$. Experimentally, we examine both the *worst case* and *average case* behaviour of these quantities in Section 4, using a comprehensive census of billions of input triangulations.

Such results are highly important for practitioners in three-dimensional computational topology, particularly given the exponential nature of many key algorithms. Despite this, just one preliminary study of this type appears in the literature [3]. This scarcity of results has two causes:

(i) *The lack of large, comprehensive censuses of both "typical" and "atypical" triangulations.*

There are many censuses of 3-manifold triangulations in the literature, but these typically focus on well-structured triangulations with special properties (such as minimal triangulations, or irreducible manifolds). Such triangulations are often easy to work with [18], and offer little insight into an algorithm's worst-case (or even average-case) behaviour.

It is only recently that large, comprehensive bodies of census data have been developed to study *all* triangulations of a given input size [3, 4]. By using such censuses for our experimental data, we ensure that we identify pathological cases, and also gain a clear understanding of how common or rare they are.

(ii) *The intense computation required to study normal surfaces with such large bodies of data.*

Normal surface enumeration algorithms have enjoyed significant advances in recent years, and modern algorithms now run many orders of magnitude faster than their earlier counterparts [6, 8]. The experimental work in this paper required several years of combined CPU time, and without recent algorithmic advances [2, 4, 6] this work would not have been possible.

The preliminary study in [3] examines only the combinatorial complexity $\sigma(\mathcal{T})$, and works with a data set of roughly 150 million triangulations of closed 3-manifolds. The study in this paper is significantly richer, both in scope and detail:

- We examine the algebraic complexity $\kappa(\mathcal{T})$ in addition to the combinatorial complexity $\sigma(\mathcal{T})$;
- We work with a comprehensive data set of almost *three billion* triangulations, spanning both *closed* manifolds (which are important for algorithms such as prime decomposition) and *bounded* manifolds (which are important for knot algorithms);

- We also examine these quantities in "optimised" settings that arise in practical computation—in particular, *one vertex triangulations* (a common optimisation used in many topological algorithms), and the restricted problem of enumerating only *vertex normal discs* (which is important for unknot recognition, or testing surfaces for incompressibility).

Our pathological families yield the first known explicit exponential lower bounds on worst-case complexity for the computationally important settings of bounded triangulations, closed 1-vertex triangulations, and normal discs. In many settings our pathological families match the experimental worst-case growth rates precisely, and we conjecture that the resulting exponential bounds are in fact exact.

Of related note is a result of Hass et al. [17], who establish an exponential lower bound on the worst-case complexity of a triangulated disc spanned by the unknot in $\mathbb{R}^3$ (this has particular relevance for the complexity of unknot recognition). Their result operates under stricter geometric constraints, and it is not yet known how it translates to the more flexible setting of normal surfaces.

We emphasise again that our experimental data sets use exhaustive censuses of *all* possible input triangulations below a given size. This paper introduces the first such census of bounded 3-manifold triangulations in the literature, totalling over 20 billion triangulations of size $n \leq 9$.

We use exhaustive censuses because there is no known efficient algorithm for randomly sampling large triangulations [13, 27], and although there are other methods for generating random 3-manifolds [13, 23], nothing is known about the bias of the resulting sample of triangulations. As a result, although our *census* is large, the triangulations it contains are all relatively small. Nevertheless, there are strong reasons to believe that our experimental results are indicative of behaviour for larger inputs; we discuss this further in Section 5.

2 Preliminaries

By a *triangulation* $\mathcal{T}$, we mean a collection of n abstract tetrahedra $\Delta_i = i(0123)$, $1 \leq i \leq n$, some or all of whose faces are affinely identified or "glued together" in pairs; here (0123) refers to the four vertices of tetrahedron Δ_i. As a consequence of these face gluings, many tetrahedron edges may become identified together; we refer to the result as a single *edge of the triangulation*, and likewise with vertices. The gluings must be in a way that no edge is identified with itself in reverse as a result. Moreover, each tetrahedron face must be identified with one and only one partner (we call these *internal faces*), or with nothing at all (we call these *boundary faces*). The set of boundary faces is called the *boundary of the triangulation* and denoted by $\partial\mathcal{T}$. If $\partial\mathcal{T} = \emptyset$ then $\mathcal{T}$ is called a *closed triangulation*, otherwise it is said to be *bounded*. Not all triangulations (closed or bounded) represent 3-manifolds. However, unless otherwise specified, this is true for all triangulations presented in this paper.

Throughout this article, the gluings of the triangles are given by a bijection of their vertices $i(abc) \mapsto j(def)$ where the symbol $i(abc)$ denotes triangle (abc) from tetrahedron i and the order of the vertices determines the gluing. A triangulation as defined above is sometimes referred to as *generalised triangulation*; these are more general and flexible than simplicial complexes. An important case is a *1-vertex triangulation*, in which all tetrahedron vertices become identified together.

The *face pairing graph* of a triangulation $\mathcal{T}$ is the multigraph whose nodes represent tetrahedra, and whose arcs represent pairs of tetrahedron faces that are glued together. A face pairing graph may contain loops (if two faces of the same tetrahedron are glued together), and/or multiple edges (if two tetrahedra are joined together along more than one face). See Figure 3 for examples.

A *properly embedded* surface in $\mathcal{T}$ is a surface $\mathbf{s} \subseteq \mathcal{T}$ with no self-intersections, and whose boundary lies entirely within $\partial\mathcal{T}$. A *normal surface* in $\mathcal{T}$ is a properly embedded surface that

Figure 1: Normal triangles and quadrilaterals within a tetrahedron.

meets each tetrahedron Δ of $\mathcal{T}$ in a disjoint collection of triangles and quadrilaterals, each running between distinct edges of Δ, as illustrated in Figure 1. There are four *triangle types* and three *quadrilateral types* in Δ according to which edges they meet. Within each tetrahedron there may be several triangles or quadrilaterals of any given type; collectively these are referred to as *normal pieces*. The intersection of a normal piece of a tetrahedron with one of its faces is called *normal arc*; each face has three *arc types* according to which two edges of the face an arc meets.

Counting the number of pieces of each type for a normal surface $\mathbf{s}$ gives rise to a 7-tuple per tetrahedron of $\mathcal{T}$ and hence a $7n$-tuple of non-negative integers describing $\mathbf{s}$ as a point in $\mathbb{R}^{7n}_{\geq 0}$, called its *normal coordinates*. Such a point must satisfy a set of linear homogeneous *matching equations* (one for each arc type of each internal face). These equations are necessary but not sufficient: the normal coordinates must also satisfy a set of combinatorial constraints called the *quadrilateral constraints*, which we discuss further in the appendix.

The solution set to the matching equations in $\mathbb{R}^{7n}_{\geq 0}$ is a polyhedral cone (the cross-section polytope of this cone is also known as the *projective solution space*). A *vertex normal surface* is one whose normal coordinates lie on an extremal ray of this polyhedral cone and, in addition, its normal coordinates are minimal for all integer points on this ray. Thus, there are only finitely many such vertex normal surfaces; every normal surface can then be expressed as a positive rational linear combination of these surfaces just like every point in a polyhedral cone is a positive rational linear combination of points on its extremal rays. This is why, when enumerating normal surfaces in a triangulation, we typically just consider the finite set of vertex normal surfaces.

3 Theoretical lower bounds

Here we establish lower bounds for the worst-case values of $\sigma(\mathcal{T})$ and $\kappa(\mathcal{T})$, i.e., the maximum possible $\sigma(\mathcal{T})$ or $\kappa(\mathcal{T})$ for an n-tetrahedron triangulation $\mathcal{T}$. Recall that $\sigma(\mathcal{T})$ measures the combinatorial complexity, i.e., the number of vertex normal surfaces within $\mathcal{T}$, and that $\kappa(\mathcal{T})$ measures the algebraic complexity, i.e., the maximum coordinate of any vertex normal surface of $\mathcal{T}$.

Few such results are known: there are no explicit lower bounds on the worst-case $\kappa(\mathcal{T})$ in the literature, and the only explicit lower bound on the worst-case $\sigma(\mathcal{T})$ is given by a family of closed triangulations with $\sigma(\mathcal{T}) \in \Theta(17^{n/4}) \simeq \Theta(2.03^n)$ [3]. In this section, we give new exponential lower bounds for $\sigma(\mathcal{T})$ and $\kappa(\mathcal{T})$ in a variety of settings that hold particular relevance for key algorithms in computational geometry and topology. We sketch the main constructions and results here; see the appendix for detailed proofs.

3.1 Closed triangulations with many normal surfaces

Important 3-manifold algorithms such as prime decomposition often begin by converting the input triangulation to a 1-vertex triangulation, whereupon the subsequent processing becomes significantly easier [18]. The $\Theta(2.03^n)$ family of [3] is not of this type (each triangulation has $n+1$ vertices), which raises the question of how such bounds behave in a 1-vertex setting:

Theorem 1. *There is a family $\mathcal{A}_n$, $n \geq 1$, of closed 1-vertex triangulations with n tetrahedra and $\sigma(\mathcal{A}_n) = 2^n$ vertex normal surfaces.*

We call these triangulations *binomial triangulations*, because more precisely they have $\binom{n}{k}$ vertex normal surfaces of genus k for each $k = 0, \ldots, n$ (whereby $\sigma(\mathcal{A}_n) = \sum \binom{n}{k} = 2^n$). We construct each $\mathcal{A}_n$ from n tetrahedra $\Delta_1, \ldots, \Delta_n$ in the following manner.

We begin by folding together two faces of the same tetrahedron $\Delta_i = i(0123)$ by the gluing $i(012) \mapsto i(013)$ for $i = 1, ..., n$. Then we identify tetrahedra Δ_i and Δ_{i+1} by $i(123) \mapsto (i+1)(230)$ (where $\Delta_{n+1} = \Delta_1$). These gluings identify all the vertices to a single vertex, therefore $\mathcal{A}_n$ is a 1-vertex triangulation. See Figure 3 for a picture of the face pairing graph of $\mathcal{A}_n$ for $n = 6$. It can be shown that each $\mathcal{A}_n$ is a closed 1-vertex triangulation of the 3-sphere.

To see why $\sigma(\mathcal{A}_n) = 2^n$, we observe that in each tetrahedron there are two normal "subsurfaces" which are compatible with any other normal surface of the triangulation. One of these adds genus to the overall surface and the other does not. We show that the vertex normal surfaces are precisely combinations of these subsurfaces, whereby the binomial coefficients and 2^n growth rate easily follow.

Remark. Experimentation suggests that this family $\mathcal{A}_n$ might in fact yield a tight upper bound for closed 1-vertex triangulations; see Section 4 for details.

3.2 Bounded triangulations with many normal surfaces

The number of vertex normal surfaces in a bounded triangulation has a direct impact on algorithms such as unknot recognition and incompressibility testing [19]. Here we give the first explicit exponential lower bound on the worst-case growth rate of this quantity. The proof is based on a general construction principle (Lemma 2) which, for an arbitrary bounded triangulation $\mathcal{G}_0$ satisfying certain weak constraints, uses a recursive squaring argument to obtain a family of triangulations $\{\mathcal{G}_k\}$ with $\Omega(\beta^n)$ vertex normal surfaces, where the exponential base β is derived from $\mathcal{G}_0$. By choosing a suitable starting triangulation $\mathcal{G}_0$, we obtain the explicit base $\beta \simeq 2.3715$ (Corollary 3).

For both unknot recognition and incompressibility testing, we can improve the underlying algorithms by only considering vertex normal *discs* (vertex normal surfaces that are topologically trivial). In Theorem 4 we show that this restricted quantity is also worst-case exponential: we build a family of triangulations with $\Theta(2^n)$ vertex normal discs.

Lemma 2. *Suppose $\mathcal{G}_0$ is a bounded triangulation with n_0 tetrahedra, f_0 is a boundary face of $\mathcal{G}_0$ such that not all vertices of f_0 are identified in $\mathcal{G}_0$, and c_0 is one of the three normal arc types on f_0. If there are α_0 vertex normal surfaces in $\mathcal{G}_0$ that meet f_0 in at least one arc of type c_0 but in no other normal arc types, then $\mathcal{G}_0$ can be extended to a family of triangulations $\{\mathcal{G}_k\}$ in which the number of vertex normal surfaces grows at a rate of $\Omega(\beta^n)$, where $\beta = \alpha_0^{1/(n_0+1)}$.*

In the proof (for a detailed proof see the appendix), we recursively construct $\mathcal{G}_k$ by joining two copies of $\mathcal{G}_{k-1}$ to an additional tetrahedron Δ_k along their faces f_{k-1} (see Figure 2). Suppose there are α_{k-1} vertex normal surfaces in $\mathcal{G}_{k-1}$ that meet f_{k-1} in only arcs of type c_{k-1}. For each pair of such surfaces in the two copies of $\mathcal{G}_{k-1}$, we can combine these surfaces in a way that extends through Δ_k to meet one of the free faces f_k of Δ_k in just one chosen normal arc type c_k, and this extension yields a vertex normal surface of $\mathcal{G}_k$. There are α_{k-1}^2 such pairings, and therefore $\alpha_k \geq \alpha_{k-1}^2$ such vertex normal surfaces of $\mathcal{G}_k$. This recurrence yields the final growth rate of $\Omega(\beta^n)$ where $\beta = \alpha_0^{1/(n_0+1)}$. We need the assumption that not all vertices of the boundary face f_0 are identified in $\mathcal{G}_0$ to show that each triangulation $\mathcal{G}_k$ represents a bounded 3-manifold.

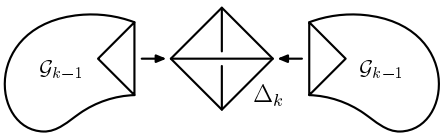

Figure 2: Attaching two copies of $\mathcal{G}_{k-1}$ to the tetrahedron Δ.

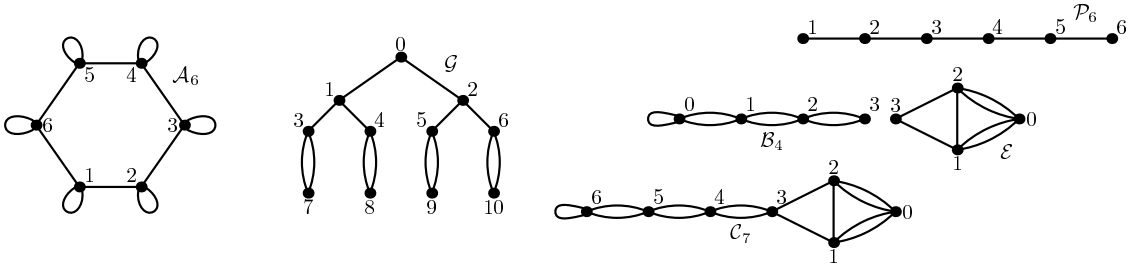

Figure 3: Face pairing graphs of the binomial family $\mathcal{A}_6$, triangulation $\mathcal{G}$, the path family $\mathcal{P}_6$, the bounded family $\mathcal{B}_4$, triangulation $\mathcal{E}$, and the closed family $\mathcal{C}_7$.

Corollary 3. *There is a triangulation $\mathcal{G} = \mathcal{G}_0$ that is the starting point of a family of bounded triangulations $\{\mathcal{G}_k\}$, $k \geq 0$, with $n_k = 12 \cdot 2^k - 1$ tetrahedra and $\sigma(\mathcal{G}_k) \geq 2.3715^{n_k}$.*

We prove this using Lemma 2 by choosing a starting triangulation $\mathcal{G} = \mathcal{G}_0$ with $n_0 = 11$ tetrahedra and a choice of face f_0 and arc type c_0 with $\alpha_0 = 31\,643$ corresponding vertex normal surfaces. This yields a growth rate of $\Omega(\beta^n)$ with $\beta = 31\,643^{1/12} \simeq 2.3715$. The face pairing graph of $\mathcal{G}$ is shown in Figure 3; for a detailed construction see the appendix.

For our final result, we construct the *path triangulation* $\mathcal{P}_n$ from n tetrahedra $\Delta_i = i(0123)$, $i = 1, ..., n$, by joining tetrahedra Δ_i and Δ_{i+1} by the map $i(012) \mapsto (i+1)(013)$. It can be shown that each $\mathcal{P}_n$ is a bounded triangulation whose underlying 3-manifold is the 3-ball (for details see the appendix).

Theorem 4. *For each $n \geq 1$, $\mathcal{P}_n$ has $2^{n+1} + \frac{(n+1)(n+2)}{2} \in \Theta(2^n)$ vertex normal discs.*

We prove this by obtaining explicit recurrences for the number of vertex normal surfaces with different normal arcs based on the matching equations. We have essentially two choices for each normal arc, and either a triangle or quadrilateral can be added to the previous triangulation giving the $\Theta(2^n)$ growth rate. Again, see the appendix for details.

3.3 Lower bounds for the size of normal coordinates

Here we give exponential lower bounds on the worst-case algebraic complexity $\kappa(\mathcal{T})$. Our bounds follow a Fibonacci growth rate of $\Omega([(1+\sqrt{5})/2]^n) \simeq \Omega(1.618)^n$. Understanding the worst-case $\kappa(\mathcal{T})$ is important for improving the time and space complexity of normal surface enumeration algorithms due to a better handling of the integer arithmetic involved (see Section 1).

To obtain such lower bounds, we first construct a family of bounded triangulations, each containing a vertex normal disc with coordinates growing exponentially in the number of tetrahedra. We then close the bounded family using a constant number of additional tetrahedra so that the vertex normal surface with exponential coordinates is preserved. In this way, we are able to construct two families of triangulations, bounded and closed, with Fibonacci type growth rates for $\kappa(\mathcal{T})$.

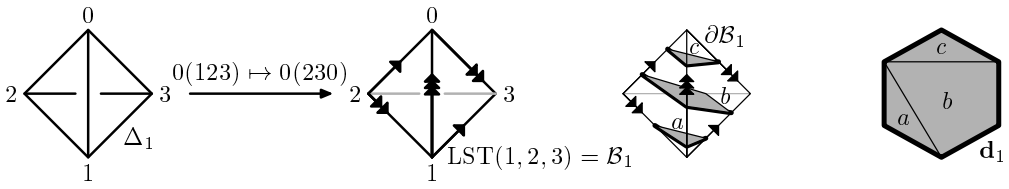

Figure 4: Left: construction of the 1-tetrahedron triangulation of $\mathrm{LST}(1,2,3)$. Right: the meridian disc $\mathbf{d}_1$ which is a 6-gon with 1 normal quad, 2 normal triangles and a maximum of 3 edge intersections.

Figure 5: Intersection of $\mathbf{s}$ (left) and $\mathbf{t}$ (right) with $\partial\mathcal{E}$

The key objects of the construction are so-called *layered solid tori* [18]. These are parameterised triangulations of the solid torus: the layered solid torus denoted $\mathrm{LST}(a,b,a+b)$ has as its boundary a triangulation of the torus with exactly three boundary edges, such that the *meridian disc* (the unique disc of the solid torus meeting the boundary in a non-contractible closed curve) intersects the boundary edges in a, b and $a+b$ points.

Layered solid tori are very common tools when constructing triangulations of a given type of 3-manifold (see [20, 22, 26] for more about constructing 3-manifolds). The most prominent example of a layered solid torus, the one tetrahedron triangulation of $\mathrm{LST}(1,2,3)$, is shown in Figure 4.

Theorem 5. *There is a family $\mathcal{B}_n$ of bounded 1-vertex triangulations with n tetrahedra, where $\mathcal{B}_n$ contains a vertex normal disc $\mathbf{d}_n$ with maximum coordinate F_{n+1}, where F_k denotes the k-th Fibonacci number.*

The family $\mathcal{B}_n$ consists of layered solid tori of type $\mathrm{LST}(\mathrm{F}_{n+1}, \mathrm{F}_{n+2}, \mathrm{F}_{n+3})$ and each vertex normal surface $\mathbf{d}_n$ is the corresponding meridian disc. See the appendix for details of the proof.

The key idea for the construction of a closed family $\mathcal{C}_n$ of n-tetrahedron triangulations containing a vertex normal surface with exponentially growing coordinates is to find a small m-tetrahedron triangulation $\mathcal{E}$ with the same boundary as the family of layered solid tori $\mathcal{B}_n$ acting like a type of plug. By this we mean that $\mathcal{E}$ contains a normal surface intersecting $\partial\mathcal{E}$ in the same way as $\mathbf{d}_n$ intersects $\partial\mathcal{B}_n$ for all $n \geq 1$. This gives rise to a vertex normal surface in $\mathcal{C}_n$ with maximum coordinate greater than or equal to F_{n+1}. Since m is constant, this gives the same asymptotic lower bound on $\kappa(\mathcal{T})$ for closed triangulations as for bounded triangulations.

Remark (The triangulation $\mathcal{E}$)**.** There is a 4-tetrahedron bounded triangulation $\mathcal{E}$ with $\partial\mathcal{E} = \partial\mathcal{B}_m$ for all m which contains two vertex normal surfaces $\mathbf{s}$ and $\mathbf{t}$ that can be combined into a normal surface intersecting $\partial\mathcal{E}$ in the same pattern as $\mathbf{d}_m$. The face pairing graph of $\mathcal{E}$ is shown in Figure 3, and the intersection of $\mathbf{s}$ and $\mathbf{t}$ with its boundary $\partial\mathcal{E}$ is shown in Figure 5. A detailed description of $\mathcal{E}$ can be found in the appendix.

Theorem 6. *There is a family $\mathcal{C}_n$ of closed 1-vertex triangulations with n tetrahedra, $n \geq 5$, each containing a vertex normal surface with maximum coordinate at least F_{n-3} if $n \equiv 2 \bmod 3$ or at least $2\,\mathrm{F}_{n-3}$ otherwise.*

Input size n	Closed	Closed and 1-vertex	Bounded
1	4	3	3
2	17	12	17
3	81	63	156
4	577	433	2 308
5	5 184	3 961	45 046
6	57 753	43 584	995 920
7	722 765	538 409	25 225 447
8	9 787 509	7 148 483	695 134 018
9	139 103 032	99 450 500	19 933 661 871
10	2 046 869 999	1 430 396 979	
Total	$n \leq 10$: 2 196 546 921	$n \leq 10$: 1 537 582 427	$n \leq 8$: 721 402 915 $n \leq 9$: 20 655 064 786

Table 1: The number of 3-manifold triangulations in the census

As outlined above, we construct $\mathcal{C}_n$ by gluing $\mathcal{B}_{n-4}$ and $\mathcal{E}$ along their boundary tori. If $n \equiv 2 \bmod 3$, the meridian disc $\mathbf{d}_{n-4}$ glued to a combination of $\mathbf{s}$ and $\mathbf{t}$ yields a vertex normal projective plane or, if $n \equiv 0, 1 \bmod 3$, twice $\mathbf{d}_{n-4}$ with a combination of $\mathbf{s}$ and $\mathbf{t}$ yields a vertex normal sphere; the maximum coordinates are then as stated. See the appendix for details.

We note that the vertex normal sphere from above (in the case $n \equiv 0, 1 \bmod 3$) is the only non-vertex linking normal sphere in $\mathcal{C}_n$. Detecting these normal surface types is one of the key tasks in important 3-manifold problems such as prime decomposition. Hence, the family $\mathcal{C}_n$ is an example for a case where, in order to prove the existence of such a normal sphere, dealing with exponentially large normal coordinates cannot be avoided. This is a hint towards the conjecture that these problems are intrinsically hard to solve using normal surface enumeration methods.

4 Experimental behaviour

We turn now to an experimental study of the combinatorial and algebraic complexities of vertex normal surfaces. Our experimental data consists of *all closed 3-manifold triangulations* with $n \leq 10$ tetrahedra, and *all bounded 3-manifold triangulations* with $n \leq 8$ tetrahedra (each appearing precisely once up to relabelling). As shown in Table 1, this yields almost 3 billion triangulations in total ($2\,196\,546\,921 + 721\,402\,915 = 2\,917\,949\,836$). We also extract the ~ 1.5 billion *1-vertex* triangulations from the closed census for additional study.

Generating exhaustive censuses of all possible inputs requires sophisticated algorithms and significant computational resources. The $n \leq 10$-tetrahedron census of all closed triangulations first appeared in [4] (which also describes some of the underlying algorithms). The $n \leq 8$-tetrahedron census of all bounded triangulations is new to this paper. Moreover, we have constructed this bounded census for $n = 9$, with over 20 billion triangulations; however, we only use $n \leq 8$ for our experiments because the subsequent analysis of normal surfaces for $n = 9$ remains out of our computational reach (for $n = 8$ this analysis already consumed years of CPU time).

Table 2 summarises our experimental results, and gives worst-case and average-case measurements (labelled *Max* and *Avg* respectively) for the quantities $\sigma(\mathcal{T})$ and $\kappa(\mathcal{T})$ in our various settings. Each measurement is taken over the relevant census of triangulations from Table 1. The *Closed* and *Bounded* columns refer to all closed or bounded triangulations respectively; in the *Closed 1-vertex* column we restrict our attention to 1-vertex triangulations of closed manifolds, and in the *Bounded, discs* column we only count vertex normal discs (not all vertex normal surfaces). We have also measured the algebraic complexity $\kappa(\mathcal{T})$ in the 1-vertex and discs-only settings, but we

Input size n	Combinatorial complexity $\sigma(\mathcal{T})$								Algebraic complexity $\kappa(\mathcal{T})$			
	Closed		Closed 1-vertex		Bounded		Bounded, discs		Closed		Bounded	
	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg	Max	Avg
1	3	2.0	2	1.7	7	5.0	7	4.0	1	1.0	1	1.0
2	7	3.9	4	3.3	14	8.2	14	5.2	2	1.2	2	1.3
3	11	5.5	8	4.9	35	14.0	27	7.0	3	1.5	3	1.7
4	18	8.8	16	7.8	85	31.3	69	11.6	4	1.8	6	2.5
5	36	13.3	32	12.0	236	69.5	176	20.2	7	2.1	12	3.4
6	70	20.8	64	18.6	688	152.6	440	34.8	10	2.3	20	4.3
7	144	32.2	128	28.8	1943	376.6	1109	61.7	16	2.6	36	5.8
8	291	50.2	256	44.7	5725	947.4	2768	112.4	26	2.9	65	7.5
9	584	78.5	512	69.4					42	3.2		
10	1175	123.2	1024	108.2					68	3.6		
Growth	2.03^n	1.56^n	2^n	1.56^n	2.73^n	2.23^n	2.45^n	1.69^n	1.62^n	1.14^n	1.81^n	1.34^n

Table 2: Experimental worst-case and average-case results

omit the details due to space constraints; see the appendix for the details.

The final row of Table 2 gives a "best estimate" of the exponential growth rate of each quantity with respect to n (we just list the base of the exponential, ignoring any coefficients or polynomial factors). Most growth rates are estimated by linear regression[3], though for cases where the worst cases matches a known family of triangulations (see below) we give the corresponding known rate.

We can make some broad observations from Table 2:

- The average-case scenarios grow at a significantly slower rate than the worst-case scenarios, sometimes astonishingly so. This is consistent with past observations in which "typical" triangulations exhibit significantly smaller complexity properties than expected (see Section 1).
- For closed manifolds, 1-vertex triangulations only give a very slight improvement: the worst case drops from the $\Theta(17^{n/4}) \simeq \Theta(2.03^n)$ family described in [3] down to the $\Theta(2^n)$ family of Theorem 1. The closeness of these results is surprising, since the theoretical bounds on $\sigma(\mathcal{T})$ for 1-vertex triangulations are much smaller than the general bounds (see below), and algorithms for working with them are often much simpler [18].
- Bounded triangulations exhibit higher complexity properties than their closed counterparts. For the combinatorial complexity $\sigma(\mathcal{T})$ this discrepancy is very pronounced—even the average case for bounded triangulations is well above the worst case for closed triangulations. This is again consistent with past experiences in working with normal surface algorithms [10]. Restricting our attention to normal discs (e.g., for unknot recognition) does alleviate this problem somewhat.

Table 3 compares the experimental behaviour of the worst-case $\sigma(\mathcal{T})$ against its best known theoretical lower and upper bounds (here "lower bounds" refers to families of pathological triangulations with the highest known growth rate of $\sigma(\mathcal{T})$, such as those constructed in Section 3).

Regarding $\sigma(\mathcal{T})$: the lower bound of $\Theta(17^{n/4}) \simeq \Theta(2.03^n)$ is known from [3], and the remaining three lower bounds are new to this paper (Theorem 1, Corollary 3 and Theorem 4). The first two upper bounds are taken from [5]; the final two $O(64^n)$ bounds are well known but do not appear in the literature (for a proof we refer to the appendix). Regarding $\kappa(\mathcal{T})$: all four lower bounds of

[3]Specifically, we take a weighted linear regression of $\log \sigma$ or $\log \kappa$ as a function of n. The weights are taken to be $1, \ldots, n$, in order to limit the influence of anomalous small cases.

	Combinatorial complexity $\sigma(\mathcal{T})$			Algebraic complexity $\kappa(\mathcal{T})$		
	Lower bound	Experimental growth	Upper bound	Lower bound	Experimental growth	Upper bound
Closed	$\Omega(2.03^n)$	$\longleftarrow$	$O(14.556^n)$	$\Omega(1.62^n)$	$\longleftarrow$	$O(3.17^n)$
Closed 1-vertex	$\Omega(2^n)$	$\longleftarrow$	$O(4.852^n)$	$\Omega(1.62^n)$	$\longleftarrow$	$O(3.17^n)$
Bounded	$\Omega(2.37^n)$	$\simeq 2.73^n$	$O(64^n)$	$\Omega(1.62^n)$	$\simeq 1.82^n$	$O(31.63^n)$
Bounded, discs only	$\Omega(2^n)$	$\simeq 2.45^n$	$O(64^n)$	$\Omega(1.62^n)$	$\longleftarrow$	$O(31.63^n)$

Table 3: Summary of worst-case theoretical and experimental results

$\Theta([(1+\sqrt{5})/2]^n) \simeq \Theta(1.62^n)$ are new to this paper (Theorems 5 and 6), and all four upper bounds are taken from [8].

Here we see that the experimental growth rates are much closer to the lower bounds than the upper bounds; in particular, an arrow ($\longleftarrow$) indicates that the experimental worst-case growth rate is *identical* to the best lower bound (up to a constant factor). This invites the following conjectures:

Conjecture 7. *For closed triangulations, the maximum number of vertex normal surfaces $\sigma(\mathcal{T})$ for any given n grows at an asymptotic rate of $\Theta(17^{n/4}) \simeq \Theta(2.03^n)$, and for closed 1-vertex triangulations this reduces to $\Theta(2^n)$.*

For closed triangulations as well as closed 1-vertex triangulations, the maximum coordinate of any vertex normal surface $\kappa(\mathcal{T})$ for any given n grows at an asymptotic rate of $\Theta([(1+\sqrt{5})/2]^n) \simeq \Theta(1.62^n)$. For bounded triangulations, the maximum coordinate of any vertex normal disc for any given n likewise grows at an asymptotic rate of $\Theta([(1+\sqrt{5})/2]^n) \simeq \Theta(1.62^n)$.

For closed 1-vertex triangulations, our experimental data gives an even stronger result:

Theorem 8. *For closed 1-vertex triangulations, the maximum number of vertex normal surfaces $\sigma(\mathcal{T})$ for any given $n \leq 10$ is precisely 2^n, and is attained by the binomial triangulations $\mathcal{A}_n$ as described in Section 3.1.*

Conjecture 9. *Theorem 8 is true for all positive integers n.*

5 Discussion

As noted in the introduction, although we consider close to 3 billion distinct triangulations, they are all relatively small with less than or equal to 10 tetrahedra. Despite this, there are reasons to believe that our experimental results might generalise. Because we allow more flexible triangulations (not just simplicial complexes), this census contains a rich diversity of 3-manifolds (including 5114 distinct closed $\mathbb{P}^2$-irreducible manifolds [4, 25]). Moreover, several of the patterns that we see in Table 2 (in particular Theorem 8) are established early on and are remarkably consistent.

As an exception, in the case of arbitrarily bounded manifolds, a greater number of tetrahedra seems to allow a better choice for the starting triangulation $\mathcal{G}_0$ in Lemma 2 in order to obtain a higher exponential base.

As seen in Table 3, there is still a long way to go before lower bounds and upper bounds on worst-case complexities converge. Not only does this paper produce the first explicit lower bounds in several computationally important settings, but it also gives strong experimental evidence that these lower bounds are close to (or even exactly) tight. This suggests that it is now the upper bounds that require significant improvement, inviting new directions of research with a rich interplay between topology, polytopes and complexity theory.

References

[1] Joan S. Birman, *Problem list: Nonsufficiently large 3-manifolds*, Notices Amer. Math. Soc. **27** (1980), no. 4, 349.

[2] Benjamin A. Burton, *Converting between quadrilateral and standard solution sets in normal surface theory*, Algebr. Geom. Topol. **9** (2009), no. 4, 2121–2174.

[3] ———, *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.

[4] ———, *Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations*, ISSAC 2011: Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation, ACM, 2011, pp. 59–66.

[5] ———, *Maximal admissible faces and asymptotic bounds for the normal surface solution space*, J. Combin. Theory Ser. A **118** (2011), no. 4, 1410–1435.

[6] ———, *Computational topology with Regina: Algorithms, heuristics and implementations*, To appear in Geometry & Topology Down Under, Amer. Math. Soc., `arXiv:1208.2504`, 2012.

[7] Benjamin A. Burton, Ryan Budney, William Pettersson, et al., *Regina: Software for 3-manifold topology and normal surface theory*, `http://regina.sourceforge.net/`, 1999–2012.

[8] Benjamin A. Burton and Melih Ozlen, *A tree traversal algorithm for decision problems in knot theory and 3-manifold topology*, To appear in Algorithmica, `arXiv:1010.6200`, 2010.

[9] ———, *Computing the crosscap number of a knot using integer programming and normal surfaces*, To appear in ACM Trans. Math. Software, `arXiv:1107.2382`, 2011.

[10] Benjamin A. Burton, J. Hyam Rubinstein, and Stephan Tillmann, *The Weber-Seifert dodecahedral space is non-Haken*, Trans. Amer. Math. Soc. **364** (2012), no. 2, 911–932.

[11] Maurice Chiodo, *Finding non-trivial elements and splittings in groups*, J. Algebra **331** (2011), 271–284.

[12] Nathan M. Dunfield and Anil N. Hirani, *The least spanning area of a knot and the optimal bounding chain problem*, SCG '11: Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry, ACM, 2011, pp. 135–144.

[13] Nathan M. Dunfield and William P. Thurston, *Finite covers of random 3-manifolds*, Invent. Math. **166** (2006), no. 3, 457–521.

[14] Wolfgang Haken, *Theorie der Normalflächen*, Acta Math. **105** (1961), 245–375.

[15] Joel Hass, *New results on the complexity of recognizing the 3-sphere*, To appear in Oberwolfach Rep., 2012.

[16] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger, *The computational complexity of knot and link problems*, J. Assoc. Comput. Mach. **46** (1999), no. 2, 185–211.

[17] Joel Hass, Jack Snoeyink, and William P. Thurston, *The size of spanning disks for polygonal curves*, Discrete Comput. Geom. **29** (2003), no. 1, 1–17.

[18] William Jaco and J. Hyam Rubinstein, *0-efficient triangulations of 3-manifolds*, J. Differential Geom. **65** (2003), no. 1, 61–168.

[19] William Jaco and Jeffrey L. Tollefson, *Algorithms for the complete decomposition of a closed* 3*-manifold*, Illinois J. Math. **39** (1995), no. 3, 358–406.

[20] Robion Kirby, *A calculus for framed links in* S^3, Invent. Math. **45** (1978), no. 1, 35–56.

[21] Greg Kuperberg, *Knottedness is in NP, modulo GRH*, Preprint, `arXiv:1112.0845`, November 2011.

[22] W. B. Raymond Lickorish, *An introduction to knot theory*, Graduate Texts in Mathematics, no. 175, Springer, New York, 1997.

[23] Joseph Maher, *Random Heegaard splittings*, J. Topol. **3** (2010), no. 4, 997–1025.

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

[25] Sergei Matveev, *Algorithmic topology and classification of 3-manifolds*, Algorithms and Computation in Mathematics, no. 9, Springer, Berlin, 2003.

[26] Peter Orlik, *Seifert manifolds*, Lecture Notes in Mathematics, no. 291, Springer-Verlag, Berlin, 1972.

[27] Vic Reiner and John Sullivan, *Triangulations: Minutes of the open problem session*, To appear in Oberwolfach Rep., 2012.

[28] J. Hyam Rubinstein, *An algorithm to recognise small Seifert fiber spaces*, Turkish J. Math. **28** (2004), no. 1, 75–87.

Appendix: Additional proofs

Additional Preliminaries

In order to fully understand the detailed proofs given in this appendix we need some more terminology concerning triangulations and normal surface theory.

From Section 2 we know that normal surfaces within a triangulation can be described by their normal coordinates which, due to the matching equations and the non-negativity constraints, lie within a high-dimensional polyhedral cone. However, not all of the integer points in this cone give rise to a normal surface. Those normal coordinates which in fact do belong to a normal surface are called *admissible*.[4] Given an admissible vector in this cone, the corresponding normal surface can be reconstructed uniquely (up to certain forms of isotopy).

The reason why not all normal coordinates are admissible is because there is no way to insert two quadrilaterals of different types into the same tetrahedron without having a non-empty intersection (cf. Figure 1). Since normal surfaces are defined to be embedded, it follows that for every tetrahedron there can be at most one type of normal quadrilateral (though there may be many quadrilaterals of this one type). Translated to normal coordinates, this means that for every tetrahedron of a triangulation at most one of the three coordinates accounting for the quadrilaterals may be non-zero. This condition is called the *quadrilateral constraints*, and it can be shown that the admissible points—that is, the normal coordinates that do correspond to a normal surface—are precisely those integer points in the solution cone that satisfy the quadrilateral constraints.

Inside the cone it is straightforward to define the *sum* of two normal surfaces as the sum of their normal coordinates: if two normal coordinates satisfy the matching equations and non-negativity constraints then their sum will also satisfy them. However, not all of these sums represent normal surfaces, since they need not satisfy the quadrilateral constraints. Because of this we call two normal surfaces *compatible* if the sum of their normal coordinates is admissible.

Due to the gluings of a triangulation, many vertices of the tetrahedra may become identified together and, as a consequence, a small neighbourhood of a vertex in a triangulation looks like a union of cusps of tetrahedra which are themselves smaller tetrahedra. Hence, the outer boundary (or frontier) of this neighbourhood is a collection of triangles which all are normal pieces. Therefore this outer boundary is a normal surface, which we call a *vertex linking normal surface* or just a *vertex link*. In a closed triangulation (which represents a closed manifold), all vertex links must be spheres (otherwise the neighbourhood of some vertex is not locally $\mathbb{R}^3$). In a bounded triangulation (which represents a manifold with boundary), all vertex links must be spheres or discs.

The proof of Theorem 1

In this section, we give a full proof of a more detailed version of Theorem 1 from Section 3.1:

Theorem 10. *For each $n \geq 1$, there are $\binom{n}{k}$ vertex normal surfaces in $\mathcal{A}_n$ of genus k. Therefore there are $\sigma(\mathcal{A}_n) = 2^n$ vertex normal surfaces.*

Proof. Consider a single tetrahedron Δ_i with two faces identified as in Section 3.1. Define α, β and γ to be loops surrounding vertices 0, 1 and 2 respectively. The projective solution space has four vertex normal surfaces **a**, **b**, **c** and **d**. All of these surfaces have a non-empty boundary which consists of a combination of α, β and γ and all of them are compatible. Therefore we do not have to consider the quadrilateral constraints, which simplifies the rest of the argument.

[4] Some authors define "admissible" to include all non-negative multiples of such points (so admissible vectors need not have integer coordinates) [2, 5]. In this paper we restrict the notion of admissibility to integer vectors only.

Let $\mathbf{a}_i$, $\mathbf{b}_i$, $\mathbf{c}_i$, $\mathbf{d}_i$ denote the surfaces and α_i, β_i, γ_i the boundary curves in the corresponding Δ_i. Any normal surface in $\mathcal{A}_n$ can be written as

$$\sum_{i=1}^{n} \tau_{a_i}\mathbf{a}_i + \tau_{b_i}\mathbf{b}_i + \tau_{c_i}\mathbf{c}_i + \tau_{d_i}\mathbf{d}_i$$

for constants $\tau_{a_i}, \tau_{b_i}, \tau_{c_i}, \tau_{d_i} \geq 0$ for $i = 1, ..., n$. Following the argument in [3], there are only two ways in which a vertex normal surface can meet each Δ_i and satisfy the matching equations, namely $\mathbf{u}_i = \mathbf{b}_i + \mathbf{c}_i + \mathbf{d}_i$ and $\mathbf{v}_i = \mathbf{a}_i + \mathbf{d}_i$, both with boundary $\alpha_i + \beta_i + \gamma_i$. It follows that every normal surface in $\mathcal{A}_n$ can be described by

$$\sum_{i=1}^{n} p_i\mathbf{u}_i + q_i\mathbf{v}_i$$

where each $p_i, q_i \geq 0$ and where $\epsilon := p_1 + q_1 = p_2 + q_2 = ... = p_n + q_n$ is an integer.

It follows from the above that each normal surface meets the boundary of Δ_i, $i = 1, \ldots, n$, in $\epsilon(\alpha_i + \beta_i + \gamma_i)$. Each normal surface with $\epsilon > 1$ is the sum of two other normal surfaces in $\mathcal{A}_n$ and thus is not a vertex normal surface. On the other hand, if $\epsilon = 1$ it can be shown that the only rational linear combinations $p_i + q_i = 1$, $i = 1, \ldots, n$, resulting in a normal surface of $\mathcal{A}_n$ must be $p_i = 1$ and $q_i = 0$ or $p_i = 0$ and $q_i = 1$.

Therefore, the projective solution space has 2^n admissible vertices, corresponding to the 2^n combinations of either $p_i = 1$ and $q_i = 0$ or $p_i = 0$ and $q_i = 1$ at each Δ_i. The binomial coefficients easily follow from a simple Euler characteristic argument since each surface $\mathbf{u}_i$ is a disc and each $\mathbf{v}_i$ is a punctured torus. □

Proof of Lemma 2

Here we give a full proof of Lemma 2 from Section 3.2, which we restate below.

Lemma 2. *Suppose $\mathcal{G}_0$ is a bounded triangulation with n_0 tetrahedra, f_0 is a boundary face of $\mathcal{G}_0$ such that not all vertices of f_0 are identified in $\mathcal{G}_0$, and c_0 is one of the three normal arc types on f_0. If there are α_0 vertex normal surfaces in $\mathcal{G}_0$ that meet f_0 in at least one arc of type c_0 but in no other normal arc types, then $\mathcal{G}_0$ can be extended to a family of triangulations $\{\mathcal{G}_k\}$ in which the number of vertex normal surfaces grows at a rate of $\Omega(\beta^n)$, where $\beta = \alpha_0^{1/(n_0+1)}$.*

Proof. The proof consists of constructing such a family of triangulations. To do so, we take two copies of $\mathcal{G}_0$ and an additional tetrahedron Δ_1 and join both copies along their faces f_0 to Δ_1 (see Figure 6 where $k = 1$), yielding a new triangulation $\mathcal{G}_1$. Any pair of the α_0 vertex normal surfaces meeting f_0 in a copy of normal arcs of type c_0 can be combined such that this linear combination has a unique extension throughout Δ_1 meeting a third face f_1 of Δ_1 in only one of the three normal arcs c_1. An argument exploiting the uniqueness of this extension shows that any such pairing yields a new vertex normal surface and there are $\alpha_1 = \alpha_0^2$ such pairings. The assumption that not all vertices of the boundary face f_0 are identified in $\mathcal{G}_0$ is necessary to show that the resulting triangulation still represents a bounded manifold (i.e., all vertex links are discs or spheres).

In a next step, two copies of $\mathcal{G}_1$ are combined together in an analogous way to create $\mathcal{G}_2$, and so on. After k steps we have a full binary tree of depth k where the "leaves" of the tree are all copies of $\mathcal{G}_0$. In each step the number of tetrahedra is doubled plus one extra "root tetrahedron" is added. On the other hand, we at least square the number of vertex normal surfaces. These observations put together are sufficient to yield the desired result.

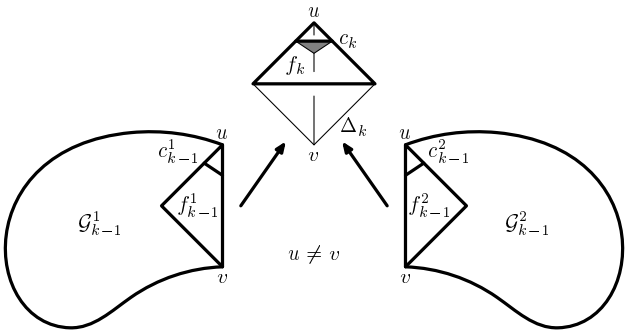

Figure 6: Attaching two copies of $\mathcal{G}_{k-1}$ to the root tetrahedron Δ_k.

More precisely: let $\mathcal{G}_k$ be the triangulation after the k-th step with n_k tetrahedra, with a designated boundary face f_k with a fixed normal arc type c_k of the root tetrahedron Δ_k which will be glued to the new root tetrahedron Δ_{k+1}, and with at least α_k vertex normal surfaces meeting f_k only in copies of c_k (see Figure 6).

Then we have $n_{k+1} = 2n_k + 1$ and $\alpha_{k+1} \geq (\alpha_k)^2$. Now let β be a real number such that $\alpha_0 \geq \beta^{n_0+1}$ and let us assume that $\alpha_k \geq \beta^{n_k+1}$. Since $\alpha_{k+1} \geq (\alpha_k)^2$ it then follows that

$$\begin{aligned} \alpha_{k+1} &\geq (\beta^{n_k+1})^2 \\ &= \beta^{2n_k+2} \\ &= \beta^{n_{k+1}+1} \end{aligned}$$

and thus $\mathcal{G}_{k+1}$ has at least $\beta^{n_{k+1}+1}$ vertex normal surfaces. To finish the induction argument we now simply choose $\beta = \alpha_0^{1/(n_0+1)}$.

It remains to show that (i) $\mathcal{G}_k$ is a triangulation of a bounded 3-manifold and (ii) any of the α_k normal surfaces is in fact a vertex normal surface.

To show that $\mathcal{G}_k$ represents a bounded 3-manifold, we first have to take a closer look at how the two copies of $\mathcal{G}_{k-1}$ (denoted by $\mathcal{G}^1_{k-1}$ and $\mathcal{G}^2_{k-1}$) have to be attached to Δ_k: the two boundary triangles of type f_{k-1} (denoted by f^1_{k-1} and f^2_{k-1}) are joined to Δ_k along an edge e. The gluing has to be in a way that the endpoints of e are not previously identified in *both* f^1_{k-1} and f^2_{k-1}, and that the normal arcs of type c_{k-1} of both triangles (denoted by c^1_{k-1} and c^2_{k-1}) are next to the same vertex of Δ_k which thus has to be an endpoint of e (see Figure 6). Note that this is always possible if not all vertices of f^1_{k-1} are identified together and likewise with f^2_{k-1}.

Since $\mathcal{G}_0 = \mathcal{G}$ is a bounded triangulation, in order to prove that $\mathcal{G}_k$ is a bounded triangulation it suffices to assume that $\mathcal{G}_{k-1}$ is a bounded manifold and then show that all vertex links of vertices of the root tetrahedron Δ_k are still triangulations of the disc after attaching $\mathcal{G}^1_{k-1}$ and $\mathcal{G}^2_{k-1}$ to Δ_k.

If none of the vertices of f^i_{k-1} are identified in $\mathcal{G}^i_{k-1}$, $i = 1, 2$, none of the vertices of Δ_k will be identified. Thus, all vertex links are either just copies of the vertex links in $\mathcal{G}^1_{k-1}$ and $\mathcal{G}^2_{k-1}$ with an additional triangle attached or two vertex linking discs attached to each other along an additional triangle (see Figure 7 on the left).

If two of the vertices of f^1_{k-1} and f^2_{k-1} are identified, and if these are the vertices opposite the arcs c^1_{k-1} and c^2_{k-2}, then one of these vertices in each face must be disjoint from e in Δ_k. In this case, three vertices of Δ_k are identified to the same vertex and the vertex link consists of the two disjoint vertex links in $\mathcal{G}^1_{k-1}$ and $\mathcal{G}^2_{k-1}$ joined together along a triangle (the normal piece near the endpoint of e opposite of c^1_{k-1} and c^2_{k-2}), and two additional disjoint triangles which are added at opposite ends of the new disc. The result is still a disc (see Figure 7 on the right). The fourth

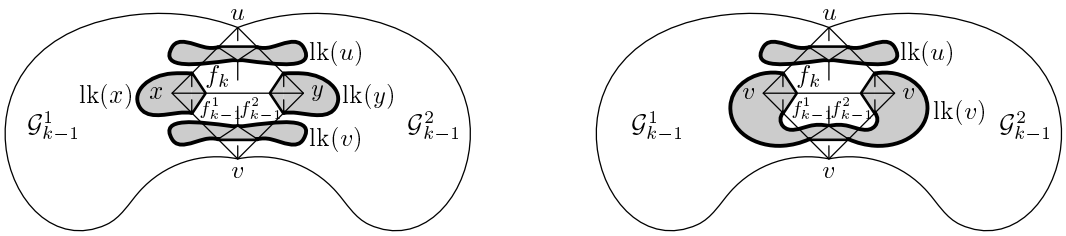

Figure 7: Left: no two vertices of f^1_{k-1} or f^2_{k-1} are identified. Right: two vertices of f^1_{k-1} and f^2_{k-1} are identified. $\mathrm{lk}(v)$ denotes the link of vertex v.

vertex of Δ_k simply connects two discs as described in the previous situation. It follows that $\mathcal{G}_k$ must represent a bounded 3-manifold, and the new face f_k likewise does not have all three vertices identified.

If two of the vertices of f^1_{k-1} and f^2_{k-1} are identified, and if these include the vertices inside arcs c^1_{k-1} and c^2_{k-2}, then we must be careful how we attach the two copies of $\mathcal{G}_{k-1}$ to Δ_k: in one copy $\mathcal{G}^1_{k-1}$ we ensure that the two identified vertices join to the common edge e, and in the other copy $\mathcal{G}^2_{k-1}$ we only allow one of these vertices to join to e. By a similar argument, this ensures that the new vertex links in $\mathcal{G}_k$ are discs, and that the new boundary face f_k again does not have all three vertices identified.

To prove that all α_k normal surfaces constructed in the k-th step are in fact vertex normal surfaces, consider two normal surfaces $\mathbf{u}_1$ and $\mathbf{u}_2$ meeting f_{k-1} in exactly x_1 and x_2 normal arcs of type c_{k-1}. Let m be the smallest common multiple of x_1 and x_2, then $\frac{m}{x_1}\,\mathbf{u}_1 + \frac{m}{x_2}\,\mathbf{u}_2$ can be extended to a normal surface in $\mathcal{G}_k$ by inserting m matching triangles into Δ_k (cf. Figure 6). We will denote this normal surface by $\mathbf{u}$. Now assume that $\mathbf{u}$ can be written as $\mathbf{u} = \lambda\mathbf{s} + \mu\mathbf{t}$ for two normal surfaces $\mathbf{s}$ and $\mathbf{t}$ of $\mathcal{G}_k$ and two rational numbers λ and μ. By construction, the restriction of this linear combination to $\mathcal{G}^1_{k-1}$ or $\mathcal{G}^2_{k-1}$ is a multiple of a vertex normal surface. Hence, inside $\mathcal{G}^1_{k-1}$ or $\mathcal{G}^2_{k-1}$ both $\mathbf{s}$ and $\mathbf{t}$ are rational multiples of $\mathbf{u}$. Moreover, since $\mathbf{u}$ inside Δ_k only consists of m triangles of a given type, $\mathbf{s}$ and $\mathbf{t}$ have to be multiples of $\mathbf{u}$ inside Δ_k as well and by the matching equations for $\mathbf{s}$ and $\mathbf{t}$ at f^1_{k-1} and f^2_{k-1} it follows that $\mathbf{s}$ and $\mathbf{t}$ are multiples of $\mathbf{u}$ throughout the entire triangulation $\mathcal{G}_k$. Thus, $\mathbf{u}$ is a vertex normal surface of $\mathcal{G}_k$. □

Construction of triangulation $\mathcal{G}$

In Section 3.2 we refer to an 11-tetrahedra triangulation $\mathcal{G}$ as a starting point for the family $\{\mathcal{G}_k\}$ of bounded triangulations in Corollary 3. Here we present this 11-tetrahedron triangulation in full.

The triangulation $\mathcal{G}$, consisting of $n = 11$ tetrahedra $\Delta_i = i(0123)$, $i = 0, \ldots, 10$, is given by the gluings in Table 4, where missing entries denote boundary faces. Face f_0 is given by $f_0 = 0(012)$ and the normal arc type c_0 is the one around vertex 0 of f_0, hence $c_0 = 0$. Using `Regina` [6, 7] for vertex normal surface enumeration yields $\sigma(\mathcal{G}) = 61\,526$ and $\alpha_0 = 31\,643$. The face pairing graph of $\mathcal{G}$ is shown in Figure 3.

Proof of Theorem 4

Here we give a full proof of Theorem 4 from Section 3.2, which we restate below.

Theorem 4. *For each $n \geq 1$, $\mathcal{P}_n$ has $2^{n+1} + \frac{(n+1)(n+2)}{2} \in \Theta(2^n)$ vertex normal discs.*

Δ_i	$i(012)$	$i(013)$	$i(023)$	$i(123)$
0		1(012)	2(021)	
1	0(013)	3(012)	4(021)	
2	0(032)	5(012)	6(021)	
3	1(013)	7(231)	7(023)	
4	1(032)	8(231)	8(023)	
5	2(013)	9(231)	9(023)	
6	2(032)	10(231)	10(023)	
7			3(023)	3(301)
8			4(023)	4(301)
9			5(023)	5(301)
10			6(023)	6(301)

Table 4: The triangulation $\mathcal{G}$

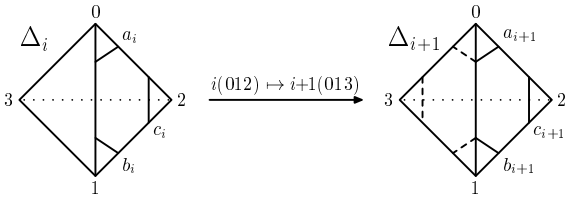

Figure 8: Notation for the normal arcs of the path family.

Proof. First of all, let $\mathcal{P}_i$ be the triangulation after the i-th step with i tetrahedra with a designated boundary face f_i of Δ_i, and with fixed normal arcs a_i, b_i and c_i on this face which will be glued to tetrahedron Δ_{i+1}, as shown in Figure 8. Also let α_i, β_i and γ_i be the number of vertex normal surfaces meeting face f_i in a_i, b_i and c_i respectively. Finally, let δ_i be the number of vertex normal surfaces not meeting face f_i.

First, let us consider the case of a single tetrahedron Δ_1. There are seven admissible vertices in the projective solution space of Δ_1 corresponding to seven vertex normal discs: four single triangles around each vertex and three single quadrilaterals. There are two surfaces meeting at each type of normal arc in face f_1 and one surface without any normal arcs in this face. Therefore $\alpha_1 = 2$, $\beta_1 = 2$, $\gamma_1 = 2$ and $\delta_1 = 1$.

Now, let us construct explicit formulae for each of the sequences α_i, β_i, γ_i and δ_i. In tetrahedron $\Delta_{i+1} = (i+1)(0123)$, normal surfaces with normal arcs a_{i+1} either have a triangle around vertex 0 or a quadrilateral separating edge 03 from edge 12, where the notation is as indicated in Figure 8. The triangle meets face f_i in normal arc a_i and the quadrilateral meets f_i in normal arc b_i. Therefore the number of normal surfaces α_{i+1} is exactly $\alpha_i + \beta_i$. A similar argument shows that $\beta_{i+1} = \beta_i + \alpha_i$, and since $\alpha_1 = \beta_1 = 2$, it follows that $\alpha_i = \beta_i = 2^i$.

Normal surfaces with normal arcs c_{i+1} on the other hand either have a triangle around vertex 2 or a quadrilateral separating edge 01 from edge 23 (see Figure 8). The triangle does not meet f_i and the quadrilateral meets f_i in normal arc c_i. Therefore the number of normal surfaces $\gamma_{i+1} = \gamma_i + 1$. Therefore since $\gamma_1 = 2$ we have $\gamma_i = i + 1$.

Now consider δ_{i+1}, the number of normal surfaces not meeting face f_{i+1}. These surfaces either have a triangle around vertex 3 or do not meet face f_i (see Figure 8). Since there are exactly δ_i normal surfaces disjoint to f_i and the triangle around vertex 3 meets face f_i in a normal arc of type

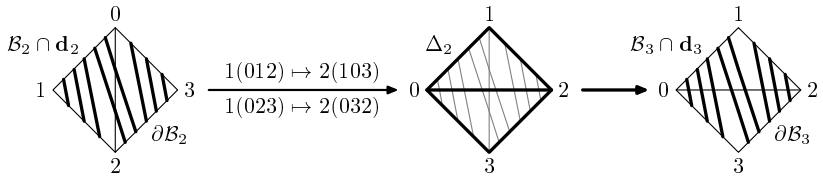

Figure 9: Building $\mathcal{B}_n$ in the case $n = 2$. Left: boundary of $\mathbf{d}_2$ in $\partial\mathcal{B}_2$. Center: the 3rd tetrahedron is glued to $\mathcal{B}_2$ Right: boundary of $\mathbf{d}_3$ in $\partial\mathcal{B}_3$.

c_i, we have $\delta_{i+1} = \gamma_i + \delta_i$. It follows that $\delta_i = \frac{i(i+1)}{2}$, since $\delta_1 = 1$ and $\gamma_i = i + 1$.

We note at this point that all vertex normal surfaces in $\mathcal{P}_n$ have been accounted for, since any normal surface that meets the face f_{i+1} in more than one arc can be expressed as a sum of the surfaces described above. Therefore the total number of vertex normal surfaces in $\mathcal{P}_n$ is

$$\alpha_n + \beta_n + \gamma_n + \delta_n = 2^n + 2^n + (n+1) + \frac{n(n+1)}{2} = 2^{n+1} + \frac{(n+1)(n+2)}{2}.$$

It remains to show that each of the above vertex normal surfaces is a disc. At the $(i+1)$-st step, there is only one triangle or one quadrilateral in Δ_{i+1} glued to one normal arc in f_i. Therefore, we either add one vertex, two edges and one face in the case of the triangle or two vertices, three edges and one face in the case of the quadrilateral to the surfaces in $\mathcal{P}_i$. In both cases the Euler characteristic does not change in the $(i+1)$-st step. Since in tetrahedron Δ_1 all seven surfaces are discs, it follows that every surface in $\mathcal{P}_n$ is a disc. □

Proof of Theorem 5

In the following, we give a full proof of Theorem 5 from Section 3.3, which we restate below.

Theorem 5. *There is a family $\mathcal{B}_n$ of bounded 1-vertex triangulations with n tetrahedra, where $\mathcal{B}_n$ contains a vertex normal disc $\mathbf{d}_n$ with maximum coordinate F_{n+1}, where F_k denotes the k-th Fibonacci number.*

Proof. The family of triangulations $\mathcal{B}_n$ consists of layered solid tori of type $\mathrm{LST}(\mathrm{F}_{n+1}, \mathrm{F}_{n+2}, \mathrm{F}_{n+3})$ and each vertex normal surface $\mathbf{d}_n$ will be the corresponding meridian disc with algebraic complexity of at least F_{n+1}.

In order to prove this we will start with the base case $\mathcal{B}_1 = \mathrm{LST}(1, 2, 3)$. Figure 4 shows a one tetrahedron triangulation of $\mathrm{LST}(1, 2, 3)$ with the meridian disc having maximum coordinate $\mathrm{F}_2 = 1$. Furthermore, by using `Regina` [6, 7] we can check that the meridian disc $\mathcal{B}_n$ is a vertex normal surface of $\mathcal{B}_n$.

Now, let us assume that $\mathcal{B}_n$ is a n-tetrahedra triangulation of $\mathrm{LST}(\mathrm{F}_{n+1}, \mathrm{F}_{n+2}, \mathrm{F}_{n+3})$ containing the meridian disc $\mathbf{d}_n$ as a vertex normal surface with maximum coordinate F_{n+1}, and intersecting the boundary edges in F_{n+1}, F_{n+2} and F_{n+3} points (see Figure 9 on the left).

To construct $\mathcal{B}_{n+1}$ and $\mathbf{d}_{n+1}$ we now glue the $(n+1)$-st tetrahedron to $\mathrm{LST}(\mathrm{F}_{n+1}, \mathrm{F}_{n+2}, \mathrm{F}_{n+3})$ such that the boundary edge with F_{n+1} intersecting points becomes an internal edge (see Figure 9 in the center). Note that $\mathcal{B}_{n+1}$ still triangulates a solid torus. There are F_{n+2} parallel normal arcs both in the upper left and in the lower right corner (in each case F_{n+1} arcs coming from triangles and F_n arcs coming from quadrilaterals of $\mathbf{d}_n$). For each of them we will insert a normal triangle into the $(n+1)$-st tetrahedron of $\mathcal{B}_{n+1}$, intersecting the boundary of $\mathcal{B}_{n+1}$ in two arcs and the newly inserted boundary edge e in one point. For each of the remaining F_{n+1} pairs of arcs in the middle

of the boundary of $\mathcal{B}_n$ we insert F_{n+1} parallel normal quadrilaterals into the $(n+1)$-st tetrahedron of $\mathcal{B}_{n+1}$, each intersecting the boundary of $\mathcal{B}_{n+1}$ in two normal arcs and edge e in one point.

Altogether, we extend the meridian disc of $\mathcal{B}_n$ to a new normal surface by adding $2\,\mathrm{F}_{n+2}$ triangles and F_{n+1} quadrilaterals, hence F_{n+4} normal pieces each intersecting e in one point. The new normal surface has two types of normal triangles with each F_{n+2} copies, and thus algebraic complexity of at least F_{n+2}. Starting from $\mathbf{d}_n$, each of these extensions applied one after another turns the previous disc into a new disc. As a consequence, the normal surface constructed by attaching all F_{n+4} normal pieces to $\mathbf{d}_n$ is still a disc and will be denoted by $\mathbf{d}_{n+1}$. The boundary curve of $\mathbf{d}_{n+1}$ is still simply closed and non-contractible in the boundary of $\mathcal{B}_{n+1}$, hence, it is the meridian disc of $\mathcal{B}_{n+1}$ intersecting the boundary edges in F_{n+2}, F_{n+3} and F_{n+4} points. By construction, $\mathcal{B}_{n+1}$ is a layered solid torus of type $\mathrm{LST}(\mathrm{F}_{n+2}, \mathrm{F}_{n+3}, \mathrm{F}_{n+4})$ with $n+1$ tetrahedra (see Figure 9 on the right).

It remains to show that $\mathbf{d}_{n+1}$ is a vertex normal surface. This is equivalent to the statement that for any rational combination

$$\mathbf{d}_{n+1} = \lambda \mathbf{s} + \mu \mathbf{t} \tag{1}$$

with $\mathbf{s}$ and $\mathbf{t}$ normal surfaces in $\mathcal{B}_{n+1}$ and $\lambda, \mu \in \mathbb{Q}$, the normal surfaces $\mathbf{s}$ and $\mathbf{t}$ must be rational multiples of $\mathbf{d}_{n+1}$.

Equation 1 is valid for every subset of normal coordinates in $\mathcal{B}_{n+1}$. In particular, it holds if we restrict $\mathbf{s}$, $\mathbf{t}$ and $\mathbf{d}_{n+1}$ to the normal coordinates associated to the first n tetrahedra of $\mathcal{B}_{n+1}$. However, this restriction applied to $\mathbf{d}_{n+1}$ yields $\mathbf{d}_n$ which is a vertex normal surface by assumption. It follows that $\mathbf{s} = \theta \mathbf{d}_n$ and $\mathbf{t} = \psi \mathbf{d}_n$ in $\mathcal{B}_n$ for some rational numbers θ and ψ.

Extending $\mathcal{B}_n$ to $\mathcal{B}_{n+1}$, we glue the $(n+1)$-st tetrahedron onto the two boundary triangles of $\mathcal{B}_n$, thus adding seven new variables and six new matching constraints to the projective solution space of $\mathcal{B}_n$ in order to enumerate the vertex normal surfaces of $\mathcal{B}_{n+1}$. Following the labelling of Figure 9 on the right, we will denote the seven normal coordinates of the $(n+1)$-st tetrahedron by

$$(t_0, t_1, t_2, t_3 \,|\, q_{01}, q_{02}, q_{03})$$

where t_X denotes the triangle type isolating vertex X from the rest of the tetrahedron and q_{XY} denotes the quadrilateral type isolating edge XY. Taking the normal coordinates from the boundary curve of $\mathbf{d}_n$ into account, this results in the following equations

$$\begin{aligned} 0 &= t_2 + q_{03} \\ 0 &= t_3 + q_{03} \\ \mathrm{F}_{n+2} &= t_0 + q_{01} \\ \mathrm{F}_{n+2} &= t_1 + q_{01} \\ \mathrm{F}_{n+1} &= t_2 + q_{02} \\ \mathrm{F}_{n+1} &= t_3 + q_{02}. \end{aligned}$$

It follows immediately that $t_2 = t_3 = q_{03} = 0$, and $q_{02} = \mathrm{F}_{n+1}$, $q_{01} = 0$ by the quadrilateral constraints and thus $t_0 = t_1 = \mathrm{F}_{n+2}$. Hence, the matching equations assure that $\mathbf{d}_n$ has a unique extension and the same holds for the restricted versions of $\mathbf{s}$ and $\mathbf{t}$ for some rational multiple of the matching equations. It follows that $\mathbf{s}$ and $\mathbf{t}$ are rational multiples of $\mathbf{d}_{n+1}$ and $\mathbf{d}_{n+1}$ is a vertex normal surface. □

The triangulation $\mathcal{E}$

In Section 3.3 we refer to a 4-tetrahedron triangulation $\mathcal{E}$ acting as a kind of "plug" for the bounded family $\mathcal{B}_n$ in order to obtain the closed family $\mathcal{C}_n$ with large algebraic complexity. Here we present

Figure 10: Left: boundary pattern of $\mathbf{s}$, two disjoint circles. Right: boundary pattern of $\mathbf{t}$

this 4-tetrahedron triangulation including its vertex normal surfaces $\mathbf{s}$ and $\mathbf{t}$ in detail.

The triangulation $\mathcal{E}$ is given by the 4 tetrahedra $\Delta_i = i(0123)$, $i = 0, \ldots, 3$, and the following table of gluings:

Δ_i	$i(012)$	$i(013)$	$i(023)$	$i(123)$
0	2(231)	1(230)	2(023)	1(123)
1	3(012)	2(102)	0(301)	0(123)
2	1(103)	3(230)	0(023)	0(201)
3	1(012)	∂_1	2(301)	∂_2

where $\partial_1 = 3(013)$ and $\partial_2 = 3(123)$ denote the two boundary faces of $\mathcal{E}$ (see Figure 10). The face pairing graph of $\mathcal{E}$ is shown in Figure 3.

Enumeration of all vertex normal surfaces of $\mathcal{E}$ using `Regina` [6, 7] yields 13 surfaces, 12 of them bounded and one of them closed. However, in the following we will only take a closer look at two of the bounded surfaces which will be denoted by $\mathbf{s}$ and $\mathbf{t}$. Surface $\mathbf{s}$ is a cylinder with two boundary components (see Figure 10 on the left), $\mathbf{t}$ is a Möbius strip with one boundary component (see Figure 10 on the right). In particular, both surfaces have Euler characteristic 0.

The surfaces are given by the following standard normal coordinates (cf. proof of Theorem 5 for more about the notation of standard normal coordinates):

$$\begin{array}{rcl} \mathbf{s} & = & (0,1,1,0\,|\,0,0,0),\quad (0,1,0,0\,|\,0,1,0),\quad (0,1,0,0\,|\,0,0,1),\quad (0,2,0,2\,|\,0,0,0) \\ \mathbf{t} & = & (0,0,0,0\,|\,1,0,0),\quad (1,1,0,0\,|\,0,0,0),\quad (1,1,0,0\,|\,0,0,0),\quad (0,1,0,1\,|\,0,0,1), \end{array}$$

and boundary patterns

$$\begin{array}{rcl} \partial\mathbf{s} & = & (0,2,2),\quad (2,0,2) \\ \partial\mathbf{t} & = & (0,2,1),\quad (1,0,2), \end{array}$$

where the first triple counts the normal arcs of boundary face ∂_1 and the second triple the ones of boundary face ∂_2. The triples are ordered starting with the normal arcs around the smallest vertex label.

The surfaces $\mathbf{s}$ and $\mathbf{t}$ are compatible, meaning that their normal coordinates can be added to each other yielding a new normal surface. Since the Euler characteristic is additive under this summing operation it follows that all linear combinations of $\mathbf{s}$ and $\mathbf{t}$ will have Euler characteristic 0. Moreover, any linear combination of $\mathbf{s}$ and $\mathbf{t}$ is connected if and only if it is orientable with two boundary components or non-orientable with one boundary component. There cannot be any interior closed connected components since the only closed vertex normal surface is not compatible with $\mathbf{t}$ and a scalar multiple of $\mathbf{s}$ cannot contain closed connected components. Hence, any connected linear combination of $\mathbf{s}$ and $\mathbf{t}$ is a cylinder or a Möbius strip where the latter is true if and only if the linear combination has an odd number of copies of $\mathbf{t}$.

Keeping this in mind, simple addition of the boundary patterns of $\mathbf{s}$ and $\mathbf{t}$ shows that

$$\mathrm{F}_{n-4}\ \mathbf{t} + \frac{1}{2}\,\mathrm{F}_{n-5}\ \mathbf{s},$$

$n \geq 5$, yields the boundary pattern of the meridian disc $\mathbf{d}_{n-4}$ from the layered solid torus family $\mathcal{B}_{n-4}$ from Theorem 5 above. Moreover, since the boundary patterns of $\mathbf{s}$ and $\mathbf{t}$ are not multiples of each other, this is the only linear combination of $\mathbf{s}$ and $\mathbf{t}$ with this property. Now, since $\mathbf{s}$ contains odd coordinates, this boundary pattern bounds a surface if and only if F_{n-5} is even, or equivalently $n \equiv 2 \bmod 3$. In this case F_{n-4} is odd and since the boundary pattern is connected it bounds a Möbius strip which will be denoted by M_n. If, on the other hand $n \equiv 0, 1 \bmod 3$, the boundary pattern

$$2\,\mathrm{F}_{n-4}\ \mathbf{t} + \mathrm{F}_{n-5}\ \mathbf{s}$$

has two connected components and since $2\,\mathrm{F}_{n-4}$ is always even, it bounds a cylinder Cyl_n.

Proof of Theorem 6

In the following, we give a full proof of Theorem 5 from Section 3.3, which we restate below.

Theorem 6. *There is a family $\mathcal{C}_n$ of closed 1-vertex triangulations with n tetrahedra, $n \geq 5$, each containing a vertex normal surface with maximum coordinate at least F_{n-3} if $n \equiv 2 \bmod 3$ or at least $2\,\mathrm{F}_{n-3}$ otherwise.*

Proof. We will construct the family $\mathcal{C}_n$ by gluing $\mathcal{B}_{n-4}$ and $\mathcal{E}$ along their boundary components. Then, if $n \equiv 2 \bmod 3$, the meridian disc $\mathbf{d}_{n-4}$ glued with M_n yields a vertex normal projective plane or, if $n \equiv 0, 1 \bmod 3$, twice $\mathbf{d}_{n-4}$ glued with Cyl_n yields a vertex normal sphere of $\mathcal{C}_n$ and the maximum normal coordinates are as stated. We will call these surfaces $\mathcal{S}_n$.

It remains to show that $\mathcal{S}_n$ is a vertex normal surface of $\mathcal{C}_n$ for all $n \geq 5$. To see this, recall that $\mathbf{d}_{n-4}$ is a vertex normal surface of $\mathcal{B}_{n-4}$ for all $n \geq 5$. Hence, it suffices to show that $\mathbf{d}_n$ ($2\mathbf{d}_n$) has a unique extension in $\mathcal{C}_n$ yielding $\mathcal{S}_n$ if $n \equiv 2 \bmod 3$ ($n \equiv 0, 1 \bmod 3$) and that $\mathcal{S}_n$ has the smallest integer normal coordinates amongst all rational multiples of $\mathcal{S}_n$.

From the section above we know that there is at least one linear combination of vertex normal surfaces in $\mathcal{E}$ realising a valid extension of $\mathbf{d}_n$ or $2\mathbf{d}_n$ yielding $\mathcal{S}_n$. Using the classification of all vertex normal surfaces of $\mathcal{E}$ we can see that due to conflicting quadrilateral constraints and matching equations there are only four vertex normal surfaces in $\mathcal{E}$ which may occur in such a linear combination. Due to further compatibility constraints such a linear combination can be shown to be a linear combination of $\mathbf{s}$ and $\mathbf{t}$ and the only closed vertex normal surfaces of $\mathcal{E}$. Following the observations made in the section above it follows that there is no other linear combination of vertex normal surfaces and hence the extension of $\mathbf{d}_n$ ($2\mathbf{d}_n$) in $\mathcal{E}$ is unique.

Altogether it follows that $\mathcal{S}_n$ is a vertex normal surface. □

An upper bound for $\sigma(\mathcal{T})$

The following statement (mentioned in Section 4) is well known but does not appear in the literature, and so we give the simple proof here.

Lemma 11. *Let $\mathcal{T}$ be a triangulation with n tetrahedra. Then $\sigma(\mathcal{T}) \leq 64^n$.*

Proof. It is known that each vertex normal surface is uniquely defined by its zero set, i. e., the set of normal coordinates which are zero (see [5]). Let $(t_0, t_1, t_2, t_3 \,|\, q_{01}, q_{02}, q_{03})$ be the set of normal coordinates of a tetrahedron of $\mathcal{T}$. Then each of the t_i can be either zero or non-zero which leaves us with $2^4 = 16$ choices, and at most one of the q_{0j} can be non-zero due to the quadrilateral constraints which leaves us with 4 additional choices. All together we have $4 \cdot 2^4 = 64$ distinct zero sets per tetrahedron which yields the result. □