Maximal admissible faces and asymptotic bounds for the normal surface solution space

Benjamin A. Burton

Author's self-archived version
Available from `http://www.maths.uq.edu.au/~bab/papers/`

Abstract

The enumeration of normal surfaces is a key bottleneck in computational three-dimensional topology. The underlying procedure is the enumeration of admissible vertices of a high-dimensional polytope, where admissibility is a powerful but non-linear and non-convex constraint. The main results of this paper are significant improvements upon the best known asymptotic bounds on the number of admissible vertices, using polytopes in both the standard normal surface coordinate system and the streamlined quadrilateral coordinate system.

To achieve these results we examine the layout of admissible points within these polytopes. We show that these points correspond to well-behaved substructures of the face lattice, and we study properties of the corresponding "admissible faces". Key lemmata include upper bounds on the number of maximal admissible faces of each dimension, and a bijection between the maximal admissible faces in the two coordinate systems mentioned above.

AMS Classification Primary 52B05; Secondary 57N10, 57Q35

Keywords 3-manifolds, normal surfaces, polytopes, face lattice, complexity

1 Introduction

Computational topology in three dimensions is a diverse and expanding field, with algorithms drawing on a range of ideas from geometry, combinatorics, algebra, analysis, and operations research. A key tool in this field is *normal surface theory*, which allows us to convert difficult topological decision and decomposition problems into more tractable enumeration and optimisation problems over convex polytopes and polyhedra.

In this paper we develop new asymptotic bounds on the complexity of problems in normal surface theory, which in turn impacts upon a wide range of topological algorithms. The techniques that we use are based on ideas from polytope theory, and the bulk of this paper focuses on the combinatorics of the various polytopes and polyhedra that arise in the study of normal surfaces.

Normal surface theory was introduced by Kneser [21], and further developed by Haken [12, 13] and Jaco and Oertel [16] for use in algorithms. The core machinery of normal surface theory is now central to many important algorithms in three-dimensional topology, including unknot recognition [12], 3-sphere recognition [7, 17, 25, 27], connected sum decomposition [17, 18], and testing for embedded incompressible surfaces [8, 16].

The core ideas behind normal surface theory are as follows. Suppose we are searching for an "interesting" surface embedded within a 3-manifold (such as a disc bounded by the unknot, or a sphere that splits apart a connected sum). We construct a high-dimensional convex polytope called the *projective solution space*, and we define the *admissible* points within this polytope to be those that satisfy an additional set of non-linear and non-convex constraints. The importance of this polytope is that every admissible and rational point within it corresponds to an embedded surface within our 3-manifold, and moreover all embedded "normal" surfaces within our 3-manifold are represented in this way.

We then prove that, if any interesting surfaces exist, at least one must be represented by a *vertex* of the projective solution space. Our algorithm is now straightforward: we construct this polytope, enumerate its admissible vertices, reconstruct the corresponding surfaces, and test whether any of these surfaces is "interesting".

The development of this machinery was a breakthrough in computational topology. However, the algorithms that it produces are often extremely slow. The main bottleneck lies in enumerating the admissible vertices of the projective solution space—polytope vertex enumeration is NP-hard in general [10, 20], and there is no evidence to suggest that our particular polytope is simple enough or special enough to circumvent this.[1]

Nevertheless, there *is* strong evidence to suggest that these procedures can be made significantly faster than current theoretical bounds imply. For instance, detailed experimentation with the quadrilateral-to-standard conversion procedure—a key step in the current state-of-the-art enumeration algorithm—suggests that this conversion runs in small polynomial time, even though the best theoretical bound remains exponential [3]. Comprehensive experimentation with the projective solution space [4] suggests that the number of admissible vertices, though exponential, grows at a rate below $O(1.62^n)$ in the average case and around $O(2.03^n)$ in the worst case, compared to the best theoretical bound of approximately $O(29.03^n)$ (which we improve upon in this paper). Here the "input size" n is the number of tetrahedra in the underlying 3-manifold triangulation.

The key to this improved performance is our *admissibility* constraint. Admissibility is a powerful constraint that eliminates almost all of the complexity of the projective solution space (we see this vividly illustrated in Section 3). However, as a non-linear and non-convex constraint it is difficult to weave admissibility into complexity arguments, particularly if we wish to draw on the significant body of work from the theory of convex polytopes.

The ultimate aim of this paper is to bound the number of admissible vertices of the projective solution space. This is a critical quantity for the running times of normal surface algorithms. First, however well we exploit admissibility in our vertex enumeration algorithms, running times *must* be at least as large as the output size—that is, the number of admissible vertices. Moreover, for some topological algorithms, the procedure that we perform on each admissible vertex is significantly slower than the enumeration of these vertices (see Hakenness testing for an example [8]). In these cases, the number of admissible vertices becomes a central factor in the overall running time.

Enumeration algorithms typically work in one of two coordinate systems: *standard coordinates* of dimension $7n$, and *quadrilateral coordinates* of dimension $3n$. The strongest bounds known to date are as follows:

- In standard coordinates, the first bound on the number of admissible vertices of the projective solution space was 128^n, due to Hass et al. [14]. The author has recently refined this bound to $O(\phi^{7n}) \simeq O(29.03^n)$, where ϕ is the golden ratio [4].[2]
- In quadrilateral coordinates, the best general bound is 4^n (this bound does not appear in the literature but is well known, and we outline the simple proof in Section 2.1).

[1] In fact, Agol et al. have proven that the *knot genus* problem is NP-complete [1]. The knot genus algorithm uses normal surface theory, but in a more complex way than we describe here.

[2] The paper [4] also places a *lower* bound on the worst case complexity of $\Omega(17^{n/4}) \simeq \Omega(2.03^n)$.

- In the case where the input is a one-vertex triangulation, the author sketches a bound of approximately $O(15^n/\sqrt{n})$ admissible vertices in standard coordinates [6]. This case is important for practical computation, as we discuss further in Section 2.

The main results of this paper are as follows. In standard coordinates, we tighten the general bound from approximately $O(29.03^n)$ to $O(14.556^n)$ (Theorem 6.3). In quadrilateral coordinates, we tighten the general bound from 4^n to approximately $O(3.303^n)$ (Theorem 5.4). For the one-vertex case in standard coordinates, we strengthen $O(15^n/\sqrt{n})$ to approximately $O(4.852^n)$ (Theorem 6.4).

We achieve these results by studying not just the admissible vertices, but the broader region formed by *all* admissible points within the projective solution space. Although this region is not convex, we show that it corresponds to a well-behaved structure within the face lattice of the surrounding polytope. By working through maximal elements of this structure—that is, *maximal admissible faces* of the polytope—we are able to draw on strong results from polytope theory such as McMullen's upper bound theorem [24], yet still enjoy the significant reduction in complexity that admissibility provides.

To contrast this paper from earlier work: The bound of $O(29.03^n)$ in [4] is a straightforward consequence of McMullen's theorem, applied once to the entire projective solution space without using admissibility at all. In this paper, the key innovations are the decomposition of the admissible region into maximal admissible faces, and the combinatorial analysis of these maximal admissible faces. These new techniques allow us to apply McMullen's theorem repeatedly in a careful and targeted fashion, ultimately yielding the stronger bounds outlined above.

Throughout this paper, we restrict our attention to closed and connected 3-manifolds. In addition to the main results listed above, we also prove several key lemmata that may be useful in future work. These include an upper bound of 3^{n-1-d} maximal admissible faces of dimension d in quadrilateral coordinates (Lemma 5.2), a bijection between maximal admissible faces in quadrilateral coordinates and standard coordinates (Lemma 6.1), and a tight upper bound of $n+1$ vertices for any triangulation with $n > 2$ tetrahedra (Lemma 6.2).

The layout of this paper is as follows. Section 2 begins with an overview of relevant results from normal surface theory and polytope theory. In Section 3 we study the structure of admissible points in detail, focusing in particular on admissible faces and maximal admissible faces of the projective solution space.

We turn our attention to asymptotic bounds in Section 4, focusing on properties of the bounds obtained by McMullen's theorem. In Section 5 we prove our main results in quadrilateral coordinates, and in Section 6 we transport these results to standard coordinates with the help of the aforementioned bijection. Section 7 finishes with a discussion of our techniques, including experimental comparisons and possibilities for further improvement.

2 Preliminaries

In this section we recount key definitions and results from the two core areas of normal surface theory and polytope theory. Section 2.1 covers 3-manifold triangulations and normal surfaces, and Section 2.2 discusses convex polytopes and polyhedra.

In this brief summary we only give the details necessary for this paper. For a more thorough overview of these topics, the reader is referred to Hass et al. [14] for the theory of normal surfaces and its role in computational topology, and to Grünbaum [11] or Ziegler [30] for the theory of convex polytopes.

Assumptions. The following assumptions and conventions run throughout this paper:

- We always assume that we are working with a closed 3-manifold triangulation $\mathcal{T}$ constructed from precisely n tetrahedra (see Section 2.1 for details), and we always assume that this triangulation is connected;
- The words "polytope" and "polyhedron" refer exclusively to *convex* polytopes and polyhedra;
- For convenience, we allow arbitrary integers a, b in the binomial coefficients $\binom{a}{b}$ but we define $\binom{a}{b} = 0$ unless $0 \le b \le a$.

2.1 Triangulations and normal surfaces

A *closed 3-manifold* is a compact topological space that locally "looks" like $\mathbb{R}^3$ at every point.[3] A *closed 3-manifold triangulation* is a collection of n tetrahedra whose 2-dimensional faces are affinely identified (or "glued together") in pairs so that the resulting topological space is a closed 3-manifold.

We do not require these tetrahedra to be rigidly embedded in some larger space—in other words, tetrahedra can be "bent" or "stretched". In particular, we allow identifications between two faces of the same tetrahedron; likewise, we may find that multiple edges or vertices of the same tetrahedron become identified together as a result of our face gluings. Some authors refer to such triangulations as *semi-simplicial triangulations* or *pseudo-triangulations*. This more flexible definition allows us to represent complex topological spaces using relatively few tetrahedra, which is extremely useful for computation.

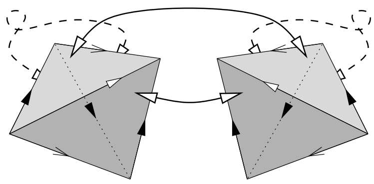

Figure 1: An example of a closed 3-manifold triangulation

Tetrahedron vertices that become identified together are collectively referred to as a single *vertex of the triangulation*; similarly for edges and 2-dimensional faces. Figure 1 illustrates a triangulation formed from $n = 2$ tetrahedra: the two front faces of the left tetrahedron are identified directly with the two front faces of the right tetrahedron, and in each tetrahedron the two back faces are identified together with a twist.[4] This triangulation has only one vertex (since all eight tetrahedron vertices become identified together), and it has precisely three edges (indicated by the three different types of arrowhead).

One-vertex triangulations are of particular interest to computational topologists, since they often simplify to very few tetrahedra, and since some algorithms become significantly simpler and/or faster in a one-vertex setting. Several authors have shown that one-vertex triangulations exist for a wide range of 3-manifolds with a variety of procedures to construct them; see [17, 22, 23] for details. We devote particular attention to one-vertex triangulations in Theorem 6.4 of this paper.

As indicated earlier, for the remainder of this paper we assume that we are working with a closed (and connected) 3-manifold triangulation $\mathcal{T}$ constructed from n tetrahedra. A

[3] More precisely, a closed 3-manifold is a compact and separable metric space in which every point has an open neighbourhood homeomorphic to $\mathbb{R}^3$ [15].

[4] The underlying 3-manifold described by this triangulation is the product space $S^2 \times S^1$.

normal surface within $\mathcal{T}$ is a closed 2-dimensional surface embedded within $\mathcal{T}$ that intersects each tetrahedron of $\mathcal{T}$ in a collection of zero or more *normal discs*. A normal disc is either an embedded *triangle* (meeting three distinct edges of the tetrahedron) or an embedded *quadrilateral* (meeting four distinct edges), as illustrated in Figure 2.

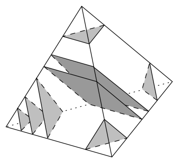

Figure 2: Normal triangles and quadrilaterals within a tetrahedron

Like the tetrahedra themselves, triangles and quadrilaterals need not be rigidly embedded (i.e., they can be "bent"). However, they must intersect the edges of the tetrahedron transversely, and they cannot meet the vertices of the tetrahedron at all. Figure 3 illustrates a normal surface within the example triangulation given earlier.[5] Normal surfaces may be disconnected or empty.

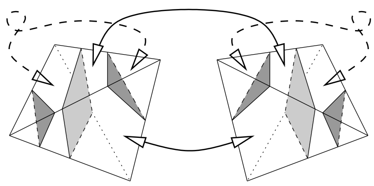

Figure 3: A normal surface within a closed 3-manifold triangulation

Within each tetrahedron there are four *types* of triangle and three *types* of quadrilateral, defined by which edges of the tetrahedron they intersect (Figure 2 includes discs of all four triangle types but only one of the three quadrilateral types). We can represent a normal surface by the integer vector

$$(\ t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4},\ q_{1,1}, q_{1,2}, q_{1,3}\ ;\ t_{2,1}, t_{2,2}, t_{2,3}, t_{2,4},\ q_{2,1}, q_{2,2}, q_{2,3}\ ;\ \ldots, q_{n,3}\ ) \in \mathbb{Z}^{7n},$$

where each $t_{i,j}$ or $q_{i,j}$ is the number of triangles or quadrilaterals respectively of the jth type within the ith tetrahedron.

A key theorem of Haken [12] states that an arbitrary integer vector in $\mathbb{R}^{7n}$ represents a normal surface if and only if:

(i) all coordinates of the vector are non-negative;

(ii) the vector satisfies the *standard matching equations*, which are $6n$ linear homogeneous equations in $\mathbb{R}^{7n}$ that depend on $\mathcal{T}$;

(iii) the vector satisfies the *quadrilateral constraints*, which state that for each i, at most one of the three quadrilateral coordinates $q_{i,1}, q_{i,2}, q_{i,3}$ is non-zero.

Any vector in $\mathbb{R}^{7n}$ that satisfies all three of these constraints is called *admissible* (note that we extend this definition to apply to non-integer vectors). The quadrilateral constraints are the most problematic of these three conditions, since they are non-linear constraints with a non-convex solution set.

[5] This surface is an embedded essential 2-dimensional sphere.

We refer to the region of $\mathbb{R}^{7n}$ that satisfies the non-negativity constraints and the standard matching equations as the *standard solution cone*, which we denote $\mathscr{S}^\vee$; this is a pointed polyhedral cone in $\mathbb{R}^{7n}$ with apex at the origin. We also consider the cross-section of this cone with the *projective hyperplane* $\sum t_{i,j} + \sum q_{i,j} = 1$, which we call the *standard projective solution space* and denote $\mathscr{S}$; this is a bounded polytope in $\mathbb{R}^{7n}$. The admissible vertices of the standard projective solution space—that is, the vertices that also satisfy the quadrilateral constraints—are called the *standard solution set*.

Tollefson [29] defines a smaller vector representation in $\mathbb{R}^{3n}$, obtained by considering only the quadrilateral coordinates $q_{i,j}$ and ignoring the triangular coordinates $t_{i,j}$. This smaller coordinate system is more efficient for computation, but its use is restricted to a smaller range of topological algorithms. Tollefson proves a theorem similar to Haken's, in that an arbitrary integer vector in $\mathbb{R}^{3n}$ represents a normal surface if and only if:

(i) all coordinates of the vector are non-negative;

(ii) the vector satisfies the *quadrilateral matching equations*, which is a smaller family of linear homogeneous equations in $\mathbb{R}^{3n}$ that again depend on $\mathcal{T}$;

(iii) the vector satisfies the *quadrilateral constraints* as defined above.

Again, any vector in $\mathbb{R}^{3n}$ that satisfies all three of these constraints is called *admissible*. The region of $\mathbb{R}^{3n}$ that satisfies the non-negativity constraints and the quadrilateral matching equations is the *quadrilateral solution cone*, denoted $\mathscr{Q}^\vee$, which is a pointed polyhedral cone in $\mathbb{R}^{3n}$ with apex at the origin. The cross-section with the *projective hyperplane* $\sum q_{i,j} = 1$ is likewise called the *quadrilateral projective solution space* and denoted $\mathscr{Q}$; this is a bounded polytope in $\mathbb{R}^{3n}$. The admissible vertices of the quadrilateral projective solution space are called the *quadrilateral solution set*.

In general, when we work in $\mathbb{R}^{7n}$ we say we are working in *standard coordinates*, and when we work in $\mathbb{R}^{3n}$ we say we are working in *quadrilateral coordinates*. See [3] for a detailed discussion of the relationship between these coordinate systems as well as fast algorithms for converting between them.

Enumerating the standard and quadrilateral solution sets is a common feature of high-level algorithms in 3-manifold topology. Moreover, this enumeration is often the computational bottleneck, and so it is important to have fast enumeration algorithms as well as good complexity bounds on the size of each solution set. The latter problem is the main focus of this paper.

As noted in the introduction, the only upper bound to date on the size of the quadrilateral solution set is the well-known but unpublished[6] bound of 4^n. The proof is simple. For any vector $\mathbf{x} \in \mathscr{Q}$, the *zero set* of $\mathbf{x}$ is defined as $\{k \mid x_k = 0\}$; in other words, the set of indices at which $\mathbf{x}$ has zero coordinates. It is shown in [6] that any vertex of $\mathscr{Q}$ can be completely reconstructed from its zero set. The quadrilateral constraints allow for at most four different zero / non-zero patterns amongst the three quadrilateral coordinates for each tetrahedron, restricting us to at most 4^n distinct zero sets in total, and therefore at most 4^n admissible vertices of $\mathscr{Q}$.

Two admissible vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{7n}$ or $\mathbf{u}, \mathbf{v} \in \mathbb{R}^{3n}$ are said to be *compatible* if the quadrilateral constraints are satisfied by both of them together. That is, for each i, at most one of the three quadrilateral coordinates $q_{i,1}, q_{i,2}, q_{i,3}$ can be non-zero in *either* $\mathbf{u}$ or $\mathbf{v}$.

Some particular vectors in standard and quadrilateral coordinates are worthy of note:

- For each vertex V of the triangulation $\mathcal{T}$, the *vertex link* of V is the vector in $\mathbb{R}^{7n}$ describing a small embedded normal sphere surrounding V. This normal surface consists

[6] Although the bound of $\leq 4^n$ does not appear in the literature, an asymptotic bound of $O(4^n/\sqrt{n})$ is sketched in [6] for the special case of a one-vertex triangulation.

of triangles only, and so the corresponding vector is zero on all quadrilateral coordinates. If $\mathcal{T}$ contains v distinct vertices then there are v corresponding vertex links, all of which are admissible and linearly independent.

- For each $i = 1, \ldots, n$, the *tetrahedral solution* $\tau^{(i)} \in \mathbb{R}^{3n}$ is the vector with $q_{i,1} = q_{i,2} = q_{i,3} = 1$ and all other quadrilateral coordinates equal to zero. The tetrahedral solutions were introduced by Kang and Rubinstein [19] as part of a "canonical basis" for normal surface theory. They satisfy the quadrilateral matching equations (so $\tau^{(i)} \in \mathscr{Q}^\vee$), but they do not satisfy the quadrilateral constraints (so $\tau^{(i)}$ is not admissible).

There is a natural relationship between standard and quadrilateral coordinates. We define the *quadrilateral projection map* $\pi\colon \mathbb{R}^{7n} \to \mathbb{R}^{3n}$ as the map that deletes all $4n$ triangular coordinates $t_{i,j}$ and retains all $3n$ quadrilateral coordinates $q_{i,j}$. This map is linear, and it maps the admissible points of $\mathscr{S}^\vee$ *onto* the admissible points of $\mathscr{Q}^\vee$. This map is not one-to-one, but the kernel is precisely the subspace of $\mathbb{R}^{7n}$ generated by the (linearly independent) vertex links. The relevant results are proven by Tollefson for integer vectors in [29]; see [3] for extensions into $\mathbb{R}^{7n}$ and $\mathbb{R}^{3n}$.

For points within the solution cones, the quadrilateral projection map preserves admissibility and inadmissibility, and it preserves compatibility and incompatibility. That is, $\mathbf{v} \in \mathscr{S}^\vee$ is admissible if and only if $\pi(\mathbf{v}) \in \mathscr{Q}^\vee$ is admissible, and admissible vectors $\mathbf{u}, \mathbf{v} \in \mathscr{S}^\vee$ are compatible if and only if $\pi(\mathbf{u}), \pi(\mathbf{v}) \in \mathscr{Q}^\vee$ are compatible.

We finish this overview of normal surface theory with an important dimensional result. This theorem is due Tillmann [28], and extends earlier work of Kang and Rubinstein for non-closed manifolds [19].

Theorem 2.1 (Tillmann, 2008)**.** *The solution space to the quadrilateral matching equations in $\mathbb{R}^{3n}$ has dimension precisely $2n$.*

2.2 Polytopes and polyhedra

We follow Ziegler [30] for our terminology: *polytopes* are always bounded (like the projective solution spaces $\mathscr{S}$ and $\mathscr{Q}$), and *polyhedra* may be bounded or unbounded (like the solution cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$). The reader is referred to [30] for background material on standard concepts such as faces, facets and supporting hyperplanes.

In this paper we work with the *face lattice* of a polytope or polyhedron P, which encodes all of the combinatorial information about the facial structure of P. Specifically, the face lattice is the poset consisting of all faces of P ordered by the subface relation, and is denoted by $L(P)$. See Figure 4 for an illustration in the case where P is the 3-dimensional cube.

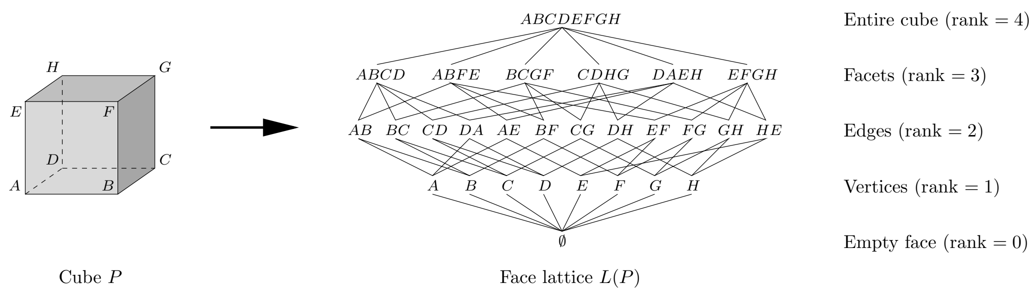

Figure 4: The face lattice of a cube

We recount some key properties of the face lattice. Any two faces $F, G \in L(P)$ have a unique greatest lower bound in $L(P)$, called the *meet* $F \wedge G$ (this corresponds to the intersection $F \cap G$), and also a unique least upper bound in $L(P)$, called the *join* $F \vee G$. There is a unique minimal element of $L(P)$ (corresponding to the empty face) and a unique maximal element of $L(P)$ (corresponding to P itself). Moreover, $L(P)$ is a *graded lattice*: it is equipped with a *rank function* $r\colon L(P) \to \mathbb{N}$ defined by $r(F) = \dim F + 1$, so that whenever G *covers* F in the poset (that is, $F < G$ and there is no X for which $F < X < G$), we have $r(G) = r(F) + 1$. Once again we refer to Ziegler [30] for details.

For any polytope F, we define the *cone over* F to be $F^\vee = \{\lambda\mathbf{x} \,|\, \mathbf{x} \in F,\ \lambda \geq 0\}$. As a special case, for the empty face $\emptyset$ we define $\emptyset^\vee = \{\mathbf{0}\}$. It is clear that the solution cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ are indeed the cones over the projective solution spaces $\mathscr{S}$ and $\mathscr{Q}$, as the notation suggests. The facial structures of polytopes and their cones are tightly related, as described by the following well-known result:

Lemma 2.2. *Let P be a d-dimensional polytope whose affine hull does not contain the origin. Then $P^\vee$ is a $(d+1)$-dimensional polyhedron, and the cone map $F \mapsto F^\vee$ is a bijection from the faces of P to the non-empty faces of $P^\vee$. This bijection maps i-faces of P to $(i+1)$-faces of $\mathscr{S}^\vee$ for all i. Both the bijection and its inverse preserve subfaces; in other words, $F^\vee \subseteq G^\vee$ if and only if $F \subseteq G$.*

A celebrated milestone in polytope complexity theory was McMullen's *upper bound theorem*, proven in 1970 [24]. In essence, this result places an upper bound on the number of i-faces of a d-dimensional k-vertex polytope, for any $i \leq d < k$. This upper bound is tight, and equality is achieved in the case of *cyclic polytopes* (and more generally, *neighbourly simplicial polytopes*). Taken in dual form, McMullen's theorem bounds the number of i-faces of a d-dimensional polytope with k facets. In this paper we use the dual form for the case $i = 0$, which reduces to the following result:

Theorem 2.3 (McMullen, 1970)**.** *For any integers $2 \leq d < k$, a d-dimensional polytope with precisely k facets can have at most*

$$\binom{k - \lfloor\frac{d+1}{2}\rfloor}{k-d} + \binom{k - \lfloor\frac{d+2}{2}\rfloor}{k-d} \tag{2.1}$$

vertices.[7]

3 Admissibility and the face lattice

In this section we explore the facial structures of the bounded polytopes $\mathscr{S}$ and $\mathscr{Q}$ (the standard and quadrilateral projective solution spaces) and the tightly-related polyhedral cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ (the standard and quadrilateral solution cones). In particular we focus on *admissible faces*, which are faces along which the quadrilateral constraints are always satisfied.

We begin by showing that the admissible faces together contain all admissible points (that is, all of the "interesting" points from the viewpoint of normal surface theory). Following this, we study the layout of admissible faces within the larger face lattice of each solution space, and we examine the relationships between admissible faces and pairs of compatible points. We finish the section by categorising *maximal* admissible faces in a variety of ways.

Definition 3.1 (Admissible face)**.** Let F be a face of the standard projective solution space $\mathscr{S}$. Then F is an *admissible face* of $\mathscr{S}$ if every point in F satisfies the quadrilateral constraints. We say that F is a *maximal admissible face* if F is not a subface of some other admissible face of $\mathscr{S}$. The same definitions apply if we replace $\mathscr{S}$ with $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$.

[7] The expression (2.1) is the number of facets of the cyclic d-dimensional polytope with k vertices; see a standard reference such as Grünbaum [11] for details.

There are always admissible points in $\mathscr{S}$ (for instance, scaled multiples of the vertex links in the underlying triangulation). Likewise, there are always admissible points in the cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ (the origin, for example). However, it might be the case that the quadrilateral solution space $\mathscr{Q}$ has no admissible points at all, in which case the empty face becomes the unique maximal admissible face of $\mathscr{Q}$.

In general, faces of a polytope are simpler to deal with than arbitrary sets of points—they have convenient representations (such as intersections with supporting hyperplanes) and useful combinatorial properties (which we discuss shortly). Our first result is to show that, in each solution space, the admissible faces together hold all of the admissible points. Jaco and Oertel make a similar remark in [16], at the point where they introduce the projective solution space.

Lemma 3.2. *Every admissible point within the standard projective solution space $\mathscr{S}$ belongs to some admissible face of $\mathscr{S}$. The same is true if we replace $\mathscr{S}$ with $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$.*

Proof. We work with $\mathscr{S}$ only; the arguments for $\mathscr{Q}$, $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ are identical. Let $\mathbf{p} \in \mathscr{S}$ be any admissible point, and let F be the minimal-dimensional face of $\mathscr{S}$ containing $\mathbf{p}$ (we can construct F by taking the intersection of all faces containing $\mathbf{p}$).

We claim that F is an admissible face. If not, let $\mathbf{q} \in F$ be some inadmissible point in F. Because $\mathbf{p}$ is admissible but $\mathbf{q}$ is not, there must be some coordinate position i for which $p_i = 0$ and $q_i > 0$.

Consider now the hyperplane $H = \{\mathbf{x} \in \mathbb{R}^{7n} \mid x_i = 0\}$. It is clear that H is a supporting hyperplane for $\mathscr{S}$ and that $\mathbf{p} \in H$ but $\mathbf{q} \notin H$. It follows that $F \cap H$ is a strict subface of F containing our original point $\mathbf{p}$, contradicting the minimality of F. $\square$

Because polyhedra have finitely many faces, every admissible face must belong to some maximal admissible face. This gives us the following immediate corollary:

Corollary 3.3. *The set of all admissible points in $\mathscr{S}$ is precisely the union of all maximal admissible faces of $\mathscr{S}$. The same is true if we replace $\mathscr{S}$ with $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$.*

Remarks. It should be noted that this union of maximal admissible faces is generally not convex. This means that we cannot (easily) apply the theory of convex polytopes to the "admissible region" within $\mathscr{S}$, which causes difficulties both for theoretical analysis (as in this paper) and for practical algorithms (see [6] for a detailed discussion). The maximal admissible faces are the largest admissible regions that *can* be described as convex polytopes, and our strategy in Sections 5 and 6 of this paper is to work with each maximal admissible face one at a time.

It should also be noted that there may be faces of $\mathscr{S}$ that are *not* admissible faces, but which contain admissible points. In particular, $\mathscr{S}$ itself is such a face. We also see this in lower dimensions; for instance, $\mathscr{S}$ might have a non-admissible edge whose endpoints are both admissible vertices.

We turn our attention now to the face lattices of the various solution spaces, and the structures formed by the admissible faces within them.

Definition 3.4 (Admissible face semilattice)**.** Let P represent one of the solution spaces $\mathscr{S}$, $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$. The *admissible face semilattice* of P, denoted $L_A(P)$, is the poset consisting of all admissible faces of P, ordered again by the subface relation.

The use of the word "semilattice" will be justified shortly. In the meantime, it is clear that the admissible face semilattice $L_A(P)$ is a substructure of the face lattice $L(P)$. Figure 5 illustrates this for the quadrilateral projective solution space, showing both $L(\mathscr{Q})$ and $L_A(\mathscr{Q})$

for a three-tetrahedron triangulation[8] of the product space $\mathbb{R}P^2 \times S^1$. The full face lattice is shown in grey, and the admissible face semilattice is highlighted in black. The admissible face semilattice contains one maximal admissible edge, two maximal admissible vertices, and no other maximal admissible faces at all.

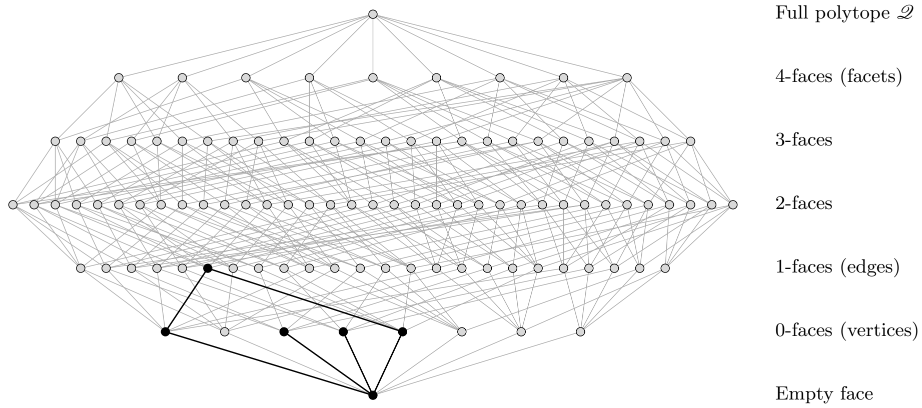

Figure 5: The face lattice and admissible face semilattice for an example triangulation

One striking observation from Figure 5 is how few admissible faces there are in comparison to the size of the full face lattice. This is a pervasive phenomenon in normal surface theory, and it highlights the importance of incorporating admissibility into enumeration algorithms and complexity bounds.

The admissible face semilattice retains several key properties of the face lattice, which we outline in the following lemma. For this result we use interval notation: in a poset S with elements $x \leq y$, the notation $[x, y]$ denotes the *interval* $\{w \in S \mid x \leq w \leq y\}$.

Lemma 3.5. *The admissible face semilattice $L_A(\mathscr{S})$ is the union of all intervals $[\emptyset, F]$ in the face lattice $L(\mathscr{S})$, where F ranges over all maximal admissible faces of $\mathscr{S}$.*

Every pair of faces $F, G \in L_A(\mathscr{S})$ has a meet (i.e., a unique greatest lower bound), and $L_A(\mathscr{S})$ has a unique minimal element (the empty face). The rank function of the face lattice $r\colon L(\mathscr{S}) \to \mathbb{N}$ maintains its covering property when restricted to $L_A(\mathscr{S})$; that is, whenever G covers F in the poset $L_A(\mathscr{S})$, we have $r(G) = r(F) + 1$.

All of these results remain true if we replace $\mathscr{S}$ with $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$.

Proof. The fact that $L_A(\mathscr{S})$ is the union of intervals $[\emptyset, F]$ for all maximal admissible faces F follows immediately from Corollary 3.3. The remaining observations follow from the properties of the face lattice $L(\mathscr{S})$ and the observation that, for any face $F \in L_A(\mathscr{S})$, all subfaces of F are also in $L_A(\mathscr{S})$. The arguments are identical for $\mathscr{Q}$, $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$. □

The poset $L_A(\mathscr{S})$ is generally not a lattice, since joins $F \vee G$ need not exist. Because meets exist however, $L_A(\mathscr{S})$ is a *meet-semilattice* (and likewise for $\mathscr{Q}$, $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$); see [26] for details.

Throughout this section we work in all four solution spaces $\mathscr{S}$, $\mathscr{Q}$, $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$. However, the cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ are precisely the cones over the projective solution spaces $\mathscr{S}$ and $\mathscr{Q}$, and so their facial structures are tightly related. The following result formalises this relationship, allowing us to transport results between different spaces where necessary.

[8]The precise triangulation is described by the dehydration string `dafbcccxaqh`, using the notation of Callahan, Hildebrand and Weeks [9].

Lemma 3.6. *Consider the cone map $F \mapsto F^\vee$ from faces of $\mathscr{S}$ into the cone $\mathscr{S}^\vee$. This cone map satisfies all of the properties described in Lemma 2.2; in particular, $F \mapsto F^\vee$ is a bijection between the faces of $\mathscr{S}$ and the non-empty faces of $\mathscr{S}^\vee$.*

Moreover, this bijection and its inverse both preserve admissibility. In other words, $F^\vee$ is an admissible face of $\mathscr{S}^\vee$ if and only if F is an admissible face of $\mathscr{S}$. This means that the cone map is also a bijection between the admissible faces of $\mathscr{S}$ and the non-empty admissible faces of $\mathscr{S}^\vee$, and a bijection between the maximal admissible faces of $\mathscr{S}$ and the maximal admissible faces of $\mathscr{S}^\vee$.

All of these results remain true if we replace $\mathscr{S}$ and $\mathscr{S}^\vee$ with $\mathscr{Q}$ and $\mathscr{Q}^\vee$ respectively.

Proof. We are able to use Lemma 2.2 because $\mathscr{S}$ lies entirely within the projective hyperplane $\sum x_i = 1$, and so the origin lies outside the affine hull of $\mathscr{S}$. It is simple to show that the bijection $F \mapsto F^\vee$ and its inverse preserve admissibility: any inadmissible point in F is also an inadmissible point in $F^\vee$, and if $\mathbf{x}$ is an inadmissible point in $F^\vee$ then $\mathbf{x}/\sum x_i$ is an inadmissible point in F. The remaining claims follow immediately from Lemma 2.2. $\square$

One consequence of Lemma 2.2 is that the face lattice of $\mathscr{S}^\vee$ is "almost isomorphic" to the face lattice of $\mathscr{S}$; the only difference is that $L(\mathscr{S}^\vee)$ contains one new element (the empty face) that is dominated by all others. What Lemma 3.6 shows is that the same relationship exists between the admissible face semilattices.

From here we turn our attention to admissible faces and compatible pairs of points. Throughout the remainder of this section we explore the relationships between these two concepts, culminating in Corollary 3.12 which categorises maximal admissible faces in terms of pairwise compatible points and vertices.

Lemma 3.7. *Let F be an admissible face of $\mathscr{S}$, $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$. Then any two points in F are compatible.*

Proof. Suppose that F contains two incompatible points $\mathbf{x}, \mathbf{y}$. Because $\mathbf{x}$ and $\mathbf{y}$ are admissible but incompatible, their sum $\mathbf{x}+\mathbf{y}$ must have non-zero entries in the coordinate positions for two distinct quadrilateral types within the same tetrahedron. Therefore the midpoint $\mathbf{z} = (\mathbf{x}+\mathbf{y})/2$ is inadmissible, contradicting the admissibility of the face F. $\square$

From this result we obtain a simple but useful bound on the complexity of admissible faces within our solution spaces. Note that by a "facet" of some i-face F, we mean an $(i-1)$-dimensional subface of F.

Corollary 3.8. *Every admissible face of $\mathscr{Q}$ or $\mathscr{Q}^\vee$ has at most n facets, and every admissible face of $\mathscr{S}$ or $\mathscr{S}^\vee$ has at most $5n$ facets.*

Proof. Let F be an admissible face of $\mathscr{Q}^\vee$. Because any two points in F are compatible (Lemma 3.7), it follows that for each tetrahedron of the underlying triangulation, two of the three corresponding quadrilateral coordinates are *simultaneously* zero for all points in F. In other words, F lies within $2n$ distinct hyperplanes of the form $x_i = 0$ (and possibly more).

Recall that $\mathscr{Q}^\vee$ is the intersection of $\mathbb{R}^{3n}$ with the hyperplanes defined by the matching equations and the $3n$ half-spaces defined by the inequalities $x_i \geq 0$. Because F is the intersection of $\mathscr{Q}^\vee$ with a supporting hyperplane, the argument above shows that F is precisely the intersection of $\mathbb{R}^{3n}$ with some number of hyperplanes and at most $3n-2n=n$ half-spaces of the form $x_i \geq 0$.

It is a standard result of polytope theory [30] that the number of half-spaces in any representation of a polytope is at least the number of facets, whereupon the number of facets of F can be at most n.

The corresponding result in $\mathscr{Q}$ is immediate from Lemma 3.6, and the corresponding arguments in $\mathscr{S}^\vee, \mathscr{S} \subseteq \mathbb{R}^{7n}$ show that F has at most $7n-2n=5n$ facets instead. $\square$

In Lemma 3.7 we showed that every admissible face must be filled with pairwise compatible points. In the following result we turn this around, showing that any set of pairwise compatible points must belong to some maximal admissible face.

Lemma 3.9. *Let $X \subseteq \mathscr{S}$ be any set of admissible points in which every two points are compatible. Then there is some maximal admissible face F of $\mathscr{S}$ for which $X \subseteq F$. The same is true if we replace $\mathscr{S}$ with $\mathscr{Q}$, $\mathscr{S}^\vee$ or $\mathscr{Q}^\vee$.*

Proof. We consider the case $X \subseteq \mathscr{S}$; again the arguments for $\mathscr{Q}$, $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ are identical. As in the proof of Corollary 3.8, the pairwise compatibility constraint shows that, for each tetrahedron of the underlying triangulation, two of the three corresponding quadrilateral coordinates are simultaneously zero for all points in X. As a consequence, X lies within all $2n$ corresponding hyperplanes of the form $x_i = 0$.

Let G be the intersection of $\mathscr{S}$ with these $2n$ hyperplanes. It follows that every point in G is admissible, and that $X \subseteq G \subseteq \mathscr{S}$. Moreover, because each hyperplane $x_i = 0$ is a supporting hyperplane for $\mathscr{S}$, it follows that G is a face of $\mathscr{S}$ (and therefore an admissible face). By finiteness of the face lattice, the admissible face G must in turn belong to some maximal admissible face F containing all of the points in X. □

Note that the set X might be contained in several distinct maximal admissible faces. However, there is always a unique admissible face of minimal dimension containing X (specifically, the intersection of *all* admissible faces containing X).

We come now to our categorisation of maximal admissible faces. Lemma 3.10 gives necessary and sufficient conditions for a face to be a maximal admissible face, and Corollary 3.12 extends these to necessary and sufficient conditions for an arbitrary set of points.

Lemma 3.10. *Let F be any admissible face of the projective solution space $\mathscr{S}$. Then the following conditions are equivalent:*

(i) F is a maximal admissible face of $\mathscr{S}$;

(ii) there is no admissible point in $\mathscr{S}$ that is not in F but that is compatible with every point in F;

(iii) there is no admissible vertex of $\mathscr{S}$ that is not in F but that is compatible with every vertex of F.

The same is true if we replace $\mathscr{S}$ with $\mathscr{Q}$. In the solution cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$, conditions (i) and (ii) are equivalent but we cannot use (iii).

Proof. We first prove (i) ⇔ (ii) for all four solution spaces. As usual we work in $\mathscr{S}$ only, since the arguments in the other solution spaces are identical.

For (i) ⇒ (ii), suppose that F is a maximal admissible face and there is some admissible point $\mathbf{x} \in \mathscr{S} \backslash F$ compatible with every point in F. Then by Lemma 3.9 there is some admissible face containing $F \cup \{\mathbf{x}\}$, contradicting the maximality of F.

For (ii) ⇒ (i), suppose that F is not a maximal admissible face. This means that there is some larger admissible face $G \supset F$, and from Lemma 3.7 it follows that there is some point $\mathbf{x} \in G \backslash F$ that is admissible and compatible with every point in F.

To prove (ii) ⇔ (iii) we require the additional fact that $\mathscr{S}$ (or $\mathscr{Q}$) is a polytope, which means that every face is the convex hull of its vertices. This is why condition (iii) fails in the cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$, where the only vertex is the origin.

For (i) ⇒ (iii), suppose that F is a maximal admissible face with vertex set V, and suppose there is some admissible vertex $\mathbf{u}$ of $\mathscr{S}$ not in F but compatible with every $\mathbf{v} \in V$. By Lemma 3.9 there is some admissible face G containing $V \cup \{\mathbf{u}\}$, and by convexity of faces it follows that $G \supseteq \mathrm{conv}(V) = F$. Because $\mathbf{u} \notin F$ we have $G \neq F$, contradicting the maximality of F.

For (iii) $\Rightarrow$ (i), suppose that F is not a maximal admissible face; again there must be some larger admissible face $G \supset F$. Because faces are convex hulls of their vertices, G must contain some admissible vertex $\mathbf{v}$ not in F, and applying Lemma 3.7 again we find that $\mathbf{v}$ is an admissible vertex of $\mathscr{S}$ not in F but compatible with every vertex of F. $\square$

We digress briefly to make a simple observation based on Lemma 3.10. Recall the *vertex links* from Section 2.1, which correspond to normal surfaces that surround the vertices of the triangulation $\mathcal{T}$ and consist entirely of triangular discs.

Corollary 3.11. *In the standard solution cone $\mathscr{S}^\vee$, every maximal admissible face contains every vertex link from the underlying triangulation.*

Proof. Vertex links represent surfaces with only triangular discs, and so the corresponding vectors in $\mathbb{R}^{7n}$ do not contain any non-zero quadrilateral coordinates at all. Therefore every vertex link is admissible and compatible with *every* point $\mathbf{x} \in \mathscr{S}^\vee$, and so by Lemma 3.10 every vertex link must belong to every maximal admissible face of $\mathscr{S}^\vee$. $\square$

It should be noted that Corollary 3.11 extends to the standard projective solution space $\mathscr{S}$ if we replace each vertex link $\mathbf{v}$ with the scaled multiple $\mathbf{v}/\sum v_i$. However, it does not extend to the quadrilateral projective solution space $\mathscr{Q}$, since in quadrilateral coordinates every vertex link projects to the zero vector.

For our final result of this section, we extend the categorisation of Lemma 3.10 to apply to arbitrary sets of points within the solution spaces.

Corollary 3.12. *Let $X \subseteq \mathscr{S}$ be any set of points. Then the following conditions are equivalent:*

(i) X is a maximal admissible face of $\mathscr{S}$;

(ii) X is a maximal set of admissible and pairwise compatible points in $\mathscr{S}$;

(iii) X is the convex hull of a maximal set of admissible and pairwise compatible vertices of $\mathscr{S}$.

In conditions (ii) and (iii), "maximal" is used in the context of set inclusion. For instance, in (ii) it means that there is no larger set $X' \supset X$ of admissible and pairwise compatible points in $\mathscr{S}$.

These equivalences remain true if we replace $\mathscr{S}$ with $\mathscr{Q}$. In the solution cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$, conditions (i) and (ii) are equivalent but again we cannot use (iii).

Proof. Steps (i) $\Rightarrow$ (ii) and (i) $\Rightarrow$ (iii) follow immediately from Lemma 3.10. To prove the remaining steps (ii) $\Rightarrow$ (i) and (iii) $\Rightarrow$ (i) we work in $\mathscr{S}$ as always, since the arguments are identical for $\mathscr{Q}$, and also $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ where applicable.

For (ii) $\Rightarrow$ (i), let X be some maximal set of admissible and pairwise compatible points in $\mathscr{S}$. By Lemma 3.9 there is some maximal admissible face $F \supseteq X$, and if $F \neq X$ then Lemma 3.7 contradicts the maximality of our original set X.

For (iii) $\Rightarrow$ (i), let $X = \mathrm{conv}(V)$ where V is a maximal set of admissible and pairwise compatible vertices of $\mathscr{S}$. Again Lemma 3.9 gives some maximal admissible face $F \supseteq V$. Because F is the convex hull of its vertices, if $F \neq X$ then F must have some additional vertex $\mathbf{v} \notin V$. By Lemma 3.7 it follows that $\mathbf{v}$ is admissible and compatible with every vertex in V, contradicting the maximality of V. $\square$

4 Bounds for general polytopes

Our ultimate aim is to place bounds on the complexity of the *admissible* face semilattice for the projective solution space. To do this, we must first understand the complexity of the *full* face lattice for an arbitrary polytope.

We begin this section by examining the behaviour of McMullen's upper bound as we change the number of facets k (Lemma 4.1) and the dimension d (Lemma 4.2). We follow with an asymptotic summation result that will prove useful in later sections (Corollary 4.4).

Notation. *For any integers $2 \le d < k$, let $M_{d,k}$ denote McMullen's upper bound as expressed in Theorem 2.3:*

$$M_{d,k} = \binom{k - \lfloor \frac{d+1}{2} \rfloor}{k-d} + \binom{k - \lfloor \frac{d+2}{2} \rfloor}{k-d}.$$

A simple rearrangement gives the equivalent expression:

$$M_{d,k} = \begin{cases} \binom{k-\frac{d}{2}}{\frac{d}{2}} + \binom{k-\frac{d}{2}-1}{\frac{d}{2}-1} & \text{if } d \text{ is even;} \\ 2\binom{k-\frac{d+1}{2}}{\frac{d+1}{2}-1} & \text{if } d \text{ is odd.} \end{cases} \tag{4.1}$$

Our first simple result describes the behaviour of $M_{d,k}$ as we vary the number of facets.

Lemma 4.1. *For any integers $2 \le d < k < k'$, we have $M_{d,k} < M_{d,k'}$. That is, increasing the number of facets of a polytope will always increase McMullen's upper bound.*

Proof. This follows immediately from equation (4.1), using the relations $\binom{m}{i} < \binom{m+1}{i}$ for $1 \le i \le m$ and $\binom{m}{0} = \binom{m+1}{0}$ for $0 \le m$. □

Varying the dimension is a little more complicated. McMullen's bound is not a monotonic function of d, and in general there can be many local maxima and minima as d ranges from 2 to $k-1$; Figure 6 illustrates this for $k = 100$ facets. However, $M_{d,k}$ is well-behaved for $d \le k/2$, as shown by the following result.

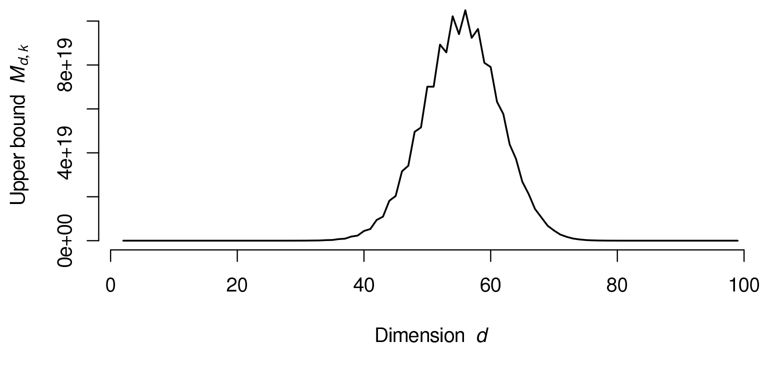

Figure 6: McMullen's upper bound $M_{d,k}$ for $k = 100$ facets

Lemma 4.2. *For any integers d, k with $2 \le d \le k/2$, we have $M_{d,k} \le M_{d+1,k}$. That is, increasing the dimension of a polytope will not decrease McMullen's upper bound, as long as there are sufficiently many facets.*

Proof. We begin by noting that $2 \le d \le k/2$ implies $d+1 < k$, so both $M_{d,k}$ and $M_{d+1,k}$ are defined. Our proof relies on a straightforward expansion of the binomial coefficients in equation (4.1). As with equation (4.1), we treat even and odd d separately.

If d is even, let $d = 2s$. Then $M_{d,k} \le M_{d+1,k}$ expands to $\binom{k-s}{s} + \binom{k-s-1}{s-1} \le 2\binom{k-s-1}{s}$, or

$$\frac{(k-s)!}{s!(k-2s)!} + \frac{(k-s-1)!}{(s-1)!(k-2s)!} \le \frac{2(k-s-1)!}{s!(k-2s-1)!}.$$

Cancelling common factors reduces this to $(k-s)+s \le 2(k-2s)$; that is, $4s \le k$, which is immediate from our initial condition $d \le k/2$.

If d is odd, let $d = 2s-1$. Now $M_{d,k} \le M_{d+1,k}$ expands to $2\binom{k-s}{s-1} \le \binom{k-s}{s} + \binom{k-s-1}{s-1}$, or

$$\frac{2(k-s)!}{(s-1)!(k-2s+1)!} \le \frac{(k-s)!}{s!(k-2s)!} + \frac{(k-s-1)!}{(s-1)!(k-2s)!}.$$

This simplifies to $2(k-s)s \le (k-s)(k-2s+1) + s(k-2s+1)$, which in turn can be rearranged to $k^2 - k \le 2(k-s)^2$.

The odd case therefore gives $M_{d,k} \le M_{d+1,k}$ if and only if $k^2 - k \le 2(k-s)^2$, and again we prove this latter inequality from our initial conditions. Using $2 \le d \le k/2$ we obtain $s \le (k+2)/4$, and so $k - s \ge (3k-2)/4 > 0$. From this we obtain

$$2(k-s)^2 \ge 2\left(\frac{3k-2}{4}\right)^2 = k^2 - k + \frac{1}{8}(k-2)^2 \ge k^2 - k,$$

and the result $M_{d,k} \le M_{d+1,k}$ is established. □

We finish this section by studying sums of the form $\sum_d \alpha^d M_{d,k}$; these sums reappear in sections 5 and 6 of this paper. Our focus is on the asymptotic growth of these sums as a function of k. We approach this by first examining the binomial coefficients $\binom{m-i}{i}$, and then returning to the sums $\sum_d \alpha^d M_{d,k}$ in Corollary 4.4.

Lemma 4.3. *For any integer $m \ge 0$ and any real $\alpha > 0$, define*

$$S_\alpha(m) = \sum_{i=0}^{\lfloor m/2 \rfloor} \alpha^i \binom{m-i}{i}.$$

Then S_α satisfies the recurrence relation $S_\alpha(m) = S_\alpha(m-1) + \alpha S_\alpha(m-2)$ for all $m \ge 2$, and the asymptotic growth rate of S_α relative to m is

$$S_\alpha(m) \in \Theta\left(\left[\frac{1+\sqrt{1+4\alpha}}{2}\right]^m\right).$$

Proof. First we note that $S_\alpha(m)$ can be written as a sum over all $i \in \mathbb{Z}$, since $\binom{m-i}{i} = 0$ whenever $i < 0$ or $i > \lfloor m/2 \rfloor$. Using the identity $\binom{m-i}{i} = \binom{m-i-1}{i} + \binom{m-i-1}{i-1}$, we have

$$\begin{aligned} S_\alpha(m) = \sum_{i\in\mathbb{Z}} \alpha^i \binom{m-i}{i} &= \sum_{i\in\mathbb{Z}} \alpha^i \binom{m-i-1}{i} + \sum_{i\in\mathbb{Z}} \alpha^i \binom{m-i-1}{i-1} \\ &= \sum_{i\in\mathbb{Z}} \alpha^i \binom{(m-1)-i}{i} + \alpha \sum_{i\in\mathbb{Z}} \alpha^{i-1} \binom{(m-2)-(i-1)}{i-1} \\ &= S_\alpha(m-1) + \alpha S_\alpha(m-2), \end{aligned}$$

thereby establishing our recurrence relation.

The characteristic equation for this recurrence is $x^2 - x - \alpha = 0$, with roots $r_1 = \frac{1-\sqrt{1+4\alpha}}{2}$ and $r_2 = \frac{1+\sqrt{1+4\alpha}}{2}$; it is clear that $r_1 < 0 < r_2$ and $0 < |r_1| < |r_2|$. It follows that

$S_\alpha(m) = c_1 r_1^m + c_2 r_2^m$ for some non-zero coefficients c_1, c_2 depending only on α, and that the growth rate of $S_\alpha(m)$ relative to m is therefore

$$S_\alpha(m) \in \Theta(r_2^m) = \Theta\left(\left[\frac{1+\sqrt{1+4\alpha}}{2}\right]^m\right).$$

□

Corollary 4.4. *For any real α in the range $0 < \alpha \leq 1$, consider the sum $\sum_{d=2}^{k-1} \alpha^d M_{d,k}$ as a function of the integer $k > 2$. This sum has an asymptotic growth rate of*

$$\sum_{d=2}^{k-1} \alpha^d M_{d,k} \in \Theta\left(\left[\frac{1+\sqrt{1+4\alpha^2}}{2}\right]^k\right).$$

Proof. Using equation (4.1) and setting $d = 2i$ or $d = 2i - 1$ for even or odd d respectively, we obtain the following identity:

$$\begin{aligned}
\sum_{d=2}^{k-1} \alpha^d M_{d,k} &= \sum_{\substack{2 \leq d < k \\ d \text{ even}}} \alpha^d \left[\binom{k-\frac{d}{2}}{\frac{d}{2}} + \binom{k-\frac{d}{2}-1}{\frac{d}{2}-1}\right] + 2\sum_{\substack{3 \leq d < k \\ d \text{ odd}}} \alpha^d \binom{k-\frac{d+1}{2}}{\frac{d+1}{2}-1} \\
&= \sum_{i=1}^{\lfloor (k-1)/2 \rfloor} \alpha^{2i}\binom{k-i}{i} + \sum_{i=1}^{\lfloor (k-1)/2 \rfloor} \alpha^{2i}\binom{k-i-1}{i-1} + 2\sum_{i=2}^{\lfloor k/2 \rfloor} \alpha^{2i-1}\binom{k-i}{i-1} \\
&= \sum_{i \in \mathbb{Z}} (\alpha^2)^i \binom{k-i}{i} - 1 - \{\alpha^k \text{ if } k \text{ is even}\} \\
&\quad + \alpha^2 \sum_{i \in \mathbb{Z}} (\alpha^2)^{i-1} \binom{(k-2)-(i-1)}{i-1} - \{\alpha^k \text{ if } k \text{ is even}\} \\
&\quad + 2\alpha \sum_{i \in \mathbb{Z}} (\alpha^2)^{i-1} \binom{(k-1)-(i-1)}{i-1} - 2\alpha - \{2\alpha^k \text{ if } k \text{ is odd}\} \\
&= S_{\alpha^2}(k) + \alpha^2 S_{\alpha^2}(k-2) + 2\alpha S_{\alpha^2}(k-1) - 2\alpha^k - 2\alpha - 1,
\end{aligned}$$

where $S_{\alpha^2}(\cdot)$ is the function defined earlier in Lemma 4.3 (though note that the subscript is now squared). Because each $S_{\alpha^2}(k)$ is non-negative and $|\alpha| \leq 1$, it follows immediately from Lemma 4.3 that

$$\sum_{d=2}^{k-1} \alpha^d M_{d,k} \in \Theta\left(\left[\frac{1+\sqrt{1+4\alpha^2}}{2}\right]^k\right).$$

□

5 The quadrilateral solution set

In this section we combine the structural results of Section 3 with the asymptotic bounds of Section 4 to yield our first main result: a new bound on the size of the quadrilateral solution set.

Recall that the *quadrilateral solution set* is the set of all admissible vertices of the quadrilateral projective solution space $\mathscr{Q}$. Little is currently known about the size of this set; the only theoretical bound to date is 4^n, as outlined in Section 2.1.

In this paper we employ more sophisticated techniques to bring this bound down to approximately $O(3.303^n)$. Our broad strategy is as follows. We first bound the number of maximal admissible faces of each dimension; in particular, we show that there are at most 3^{n-1-d} maximal admissible faces of each dimension $d \leq n-1$, and no maximal admissible faces of any dimension $d \geq n$. We then convert these results into a bound on the number of admissible vertices using McMullen's theorem and the asymptotic results of Section 4.

Throughout this section we denote the coordinates of a vector $\mathbf{x} \in \mathbb{R}^{3n}$ by

$$\mathbf{x} = (x_{1,1}, x_{1,2}, x_{1,3},\ x_{2,1}, x_{2,2}, x_{2,3},\ \ldots,\ x_{n,1}, x_{n,2}, x_{n,3}),$$

where $x_{i,j}$ is the coordinate representing the jth quadrilateral type within the ith tetrahedron. We also make repeated use of the tetrahedral solutions $\tau^{(1)}, \ldots, \tau^{(n)} \in \mathscr{Q}^\vee$; recall from Section 2.1 that the kth tetrahedral solution $\tau^{(k)}$ has $\tau^{(k)}_{k,1} = \tau^{(k)}_{k,2} = \tau^{(k)}_{k,3} = 1$ and all $(3n-3)$ remaining coordinates set to zero.

Lemma 5.1. *Every admissible face of the quadrilateral projective solution space has dimension $\leq n-1$.*

Proof. Let F be some d-dimensional admissible face of the quadrilateral projective solution space $\mathscr{Q}$, and let $F^\vee$ be the corresponding $(d+1)$-dimensional admissible face of the quadrilateral solution cone $\mathscr{Q}^\vee$. Every pair of points in $F^\vee$ must be compatible (Lemma 3.7), and so for each $i = 1, \ldots, n$ at least two of the three coordinates $x_{i,1}, x_{i,2}, x_{i,3}$ must be simultaneously zero for *all* points $\mathbf{x} \in F^\vee$.

It follows that the entire face $F^\vee$ lies within some n-dimensional subspace $S \subseteq \mathbb{R}^{3n}$ defined by setting $2n$ coordinates equal to zero. We therefore have $\dim F^\vee \leq \dim S$; that is, $d+1 \leq n$, or $d \leq n-1$. ☐

Lemma 5.2. *For each $d \in \{0, \ldots, n-1\}$, the number of maximal admissible faces of dimension d in the quadrilateral projective solution space is at most 3^{n-1-d}.*

Proof. Let $F_1, \ldots, F_k$ be distinct maximal admissible d-faces within the quadrilateral projective solution space $\mathscr{Q}$, where $k > 3^{n-1-d}$. For convenience we work in the quadrilateral solution cone $\mathscr{Q}^\vee$ instead, using the corresponding maximal admissible faces $F_1^\vee, \ldots, F_k^\vee$ each of dimension $d+1$.

Our strategy is to construct a decreasing sequence of linear subspaces $\mathbb{R}^{3n} \supset S_0 \supset S_1 \supset \ldots \supset S_n$ with the following properties:

(i) Each subspace S_i contains all of the tetrahedral solutions $\tau^{(i+1)}, \ldots, \tau^{(n)}$.

(ii) For each subspace S_i, there is some integer $t_i \geq 0$ for which S_i has dimension $\leq 2n - i - t_i$, and for which S_i contains strictly more than $3^{n-1-d-t_i}$ of the maximal admissible faces $F_1^\vee, \ldots, F_k^\vee$.

(iii) For each subspace S_i and each integer $j = 1, \ldots, i$, the subspace S_i is contained in at least two of the three hyperplanes $x_{j,1} = 0$, $x_{j,2} = 0$ and $x_{j,3} = 0$. In other words, for each of the first i tetrahedra, at least two of the three corresponding quadrilateral coordinates are simultaneously zero for all points in S_i.

We construct this sequence inductively as follows:

- We set the initial subspace S_0 to be the solution space to the quadrilateral matching equations. Property (i) holds because $\tau^{(1)}, \ldots, \tau^{(n)} \in \mathscr{Q}^\vee \subseteq S_0$. Property (ii) holds with $t_0 = 0$, since we have $\dim S_0 = 2n$ from Theorem 2.1, and since all $k > 3^{n-1-d}$ of our maximal admissible faces are contained within $\mathscr{Q}^\vee \subseteq S_0$. Property (iii) is vacuously satisfied for $i = 0$.

- For each $i > 0$, we construct S_i from S_{i-1} as follows. Let $X = \{F_j^\vee \mid F_j^\vee \subseteq S_{i-1}\}$; that is, the set of all maximal admissible faces from our original collection that are contained within the previous subspace S_{i-1}. Because each $F_j^\vee$ is an admissible face, we know from Lemma 3.7 that each $F_j^\vee$ lies in at least two (and possibly all three) of the hyperplanes $x_{i,1} = 0$, $x_{i,2} = 0$ and $x_{i,3} = 0$ (though *which* of these hyperplanes $F_j^\vee$ belongs to will typically depend on j). Consider the following two cases:

  (a) Suppose that all $F_j^\vee \in X$ are *simultaneously* contained in at least two of the three hyperplanes $x_{i,1} = 0$, $x_{i,2} = 0$ and $x_{i,3} = 0$; that is, this choice does not depend on j. Without loss of generality, let these two hyperplanes be $x_{i,2} = 0$ and $x_{i,3} = 0$.

  In this case we let S_i be the intersection of the subspace S_{i-1} with the hyperplanes $x_{i,2} = 0$ and $x_{i,3} = 0$. Note that every face $F_j^\vee \in X$ belongs to the subspace S_i as a result.

  Property (i) holds for S_i because each of the tetrahedral solutions $\tau^{(i+1)}, \ldots, \tau^{(n)}$ belongs to S_{i-1} as well as all three hyperplanes $x_{i,1} = 0$, $x_{i,2} = 0$ and $x_{i,3} = 0$. Property (iii) for S_i follows immediately from our construction.

  Property (ii) for S_i is established as follows. Let $t_i = t_{i-1}$. We note that S_i is a *strict* subspace of S_{i-1}, because the tetrahedral solution $\tau^{(i)}$ lies in S_{i-1} (from property (i) for S_{i-1}) but not S_i (because $\tau_{i,2}^{(i)}, \tau_{i,3}^{(i)} \neq 0$). It follows that $\dim S_i \leq \dim S_{i-1} - 1 \leq 2n - (i-1) - t_{i-1} - 1 = 2n - i - t_i$. Furthermore, our construction ensures that every face $F_j^\vee \in X$ lies within S_i, and using property (ii) for S_{i-1} there are strictly more than $3^{n-1-d-t_{i-1}} = 3^{n-1-d-t_i}$ such faces.

  (b) Otherwise, all $F_j^\vee \in X$ are not simultaneously contained in at least two of the three hyperplanes $x_{i,1} = 0$, $x_{i,2} = 0$ and $x_{i,3} = 0$. Consider the three sets

  $$\begin{aligned} X_1 &= \{F_j^\vee \in X \mid F_j^\vee \text{ lies in both hyperplanes } x_{i,2} = 0,\ x_{i,3} = 0\}; \\ X_2 &= \{F_j^\vee \in X \mid F_j^\vee \text{ lies in both hyperplanes } x_{i,3} = 0,\ x_{i,1} = 0\}; \\ X_3 &= \{F_j^\vee \in X \mid F_j^\vee \text{ lies in both hyperplanes } x_{i,1} = 0,\ x_{i,2} = 0\}. \end{aligned}$$

  We know from our earlier comments that $X = X_1 \cup X_2 \cup X_3$. Without loss of generality suppose that X_1 is the largest of these three sets; in particular, $|X_1| \geq |X|/3$.

  For this case we define S_i to be the intersection of the subspace S_{i-1} and the two hyperplanes $x_{i,2} = 0$ and $x_{i,3} = 0$. Note that the faces $F_j^\vee$ that lie within S_i are precisely those in the set X_1.

  Once again properties (i) and (iii) for S_i are simple consequences of our construction. To establish property (ii) for S_i, we let $t_i = t_{i-1} + 1$. The number of faces $F_j^\vee$ in S_i is $|X_1| \geq |X|/3 > 3^{n-1-d-t_{i-1}}/3 = 3^{n-1-d-t_i}$ as required. Bounding the dimension of S_i requires a little more work.

  We know that there is some face $F_a^\vee \in X$ that is not in the set X_1 (otherwise we would have fallen back to case (a)). However, this face $F_a^\vee$ must belong to one of X_1, X_2 or X_3; without loss of generality suppose that $F_a^\vee \in X_2$. Let S_i' be the intersection of the subspace S_{i-1} with the hyperplane $x_{i,3} = 0$. Because $\tau^{(i)} \in S_{i-1}$ but $\tau_{i,3}^{(i)} \neq 0$ it follows that S_i' is a strict subspace of S_{i-1}, and we have $\dim S_i' \leq \dim S_{i-1} - 1$.

  Now we find that S_i is the intersection of S_i' with the hyperplane $x_{i,2} = 0$. The face $F_a^\vee$ lies within the hyperplane $x_{i,3} = 0$ and therefore lies in S_i'; however, because $F_a^\vee \notin X_1$ it cannot also lie in the hyperplane $x_{i,2} = 0$, which means that $F_a^\vee$ does not lie in S_i. Therefore S_i is a strict subspace of S_i', and we have $\dim S_i \leq \dim S_i' - 1 \leq \dim S_{i-1} - 2$, giving a final dimension $\dim S_i \leq 2n - (i-1) - t_{i-1} - 2 = 2n - i - t_i$.

This establishes properties (i)–(iii) for our sequence of linear subspaces $\mathbb{R}^{3n} \supset S_0 \supset S_1 \supset \ldots \supset S_n$. We finish our proof by considering the final subspace S_n.

From property (ii) we know that S_n contains at least one of the maximal admissible faces $F_1^\vee, \ldots, F_k^\vee$, and so $\dim S_n \geq d+1$. The dimension constraint of property (ii) then gives $t_n \leq n - 1 - d$, whereupon we find that S_n contains strictly more than $3^{n-1-d-t_n} \geq 1$ of the maximal admissible faces $F_1^\vee, \ldots, F_k^\vee$. That is, S_n must contain *at least two* of these faces. Let these faces be $F_a^\vee$ and $F_b^\vee$.

By property (iii) we know that all points in S_n are pairwise compatible, and so every point in $F_a^\vee$ must be compatible with every point in $F_b^\vee$. However, from Corollary 3.12 we know that $F_a^\vee$ and $F_b^\vee$ are each maximal sets of admissible and pairwise compatible points in $\mathscr{Q}^\vee$, giving $F_a^\vee = F_b^\vee$ and a contradiction. □

This bound of $\leq 3^{n-1-d}$ maximal admissible faces of dimension d appears to be tight for large dimensions d (in particular, for $d \geq \frac{n}{2} - 1$ as we discuss in Section 7). Nevertheless, even for large dimensions this not the entire story. We might be able to achieve equality for some large dimensions d, but we cannot achieve equality for *all* large dimensions simultaneously, as indicated by the following result.

Lemma 5.3. *If the quadrilateral projective solution space has a maximal admissible face of dimension $n-1$, then this is the* only *maximal admissible face (of any dimension).*

Proof. Suppose that we have two distinct maximal admissible faces $F, G \subseteq \mathscr{Q}$ where $\dim F = n-1$. Once again we work in the quadrilateral solution cone $\mathscr{Q}^\vee$, using the corresponding maximal admissible faces $F^\vee, G^\vee$ with $\dim F^\vee = n$.

For each $i = 1, \ldots, n$, Lemma 3.7 shows that face $F^\vee$ must lie within at least two of the three hyperplanes $x_{i,1} = 0$, $x_{i,2} = 0$ and $x_{i,3} = 0$. Likewise, $G^\vee$ must lie within at least two of these hyperplanes, and so both $F^\vee$ and $G^\vee$ must *simultaneously* lie in at least one of the hyperplanes $x_{i,1} = 0$, $x_{i,2} = 0$ or $x_{i,3} = 0$. Without loss of generality let this common hyperplane be $x_{i,1} = 0$.

Let S be the solution space to the quadrilateral matching equations in $\mathbb{R}^{3n}$; by Theorem 2.1 we have $\dim S = 2n$. Let S' be the subspace of S formed by intersecting S with each of the hyperplanes $x_{i,1} = 0$ for $i = 1, \ldots, n$.

Each of the tetrahedral solutions $\tau^{(i)}$ belongs to S but not S'. It is clear that the tetrahedral solutions are linearly independent (their non-zero coordinates appear in distinct positions), and so $\dim S' \leq \dim S - n = n$. Faces $F^\vee$ and $G^\vee$ still lie within S' however, and because $\dim F^\vee = n$ it follows that $\dim S' = n$ and that S' is the affine hull of $F^\vee$.

We now see that the face $G^\vee$ lies within the affine hull of the face $F^\vee$; it follows that $G^\vee$ must be a subface of $F^\vee$, contradicting the maximality of $G^\vee$. □

Lemmata 5.1 and 5.2 together bound the number of maximal admissible faces of every dimension in $\mathscr{Q}$. We can now use these results to prove our main theorem, which is a new bound on the size of the quadrilateral solution set (that is, the number of admissible vertices of $\mathscr{Q}$).

Theorem 5.4. *The size of the quadrilateral solution set is asymptotically bounded above by*

$$O\left(\left[\frac{3+\sqrt{13}}{2}\right]^n\right) \quad \simeq \quad O(3.303^n).$$

Proof. Let κ denote the number of admissible vertices of the quadrilateral projective solution space $\mathscr{Q}$. Our strategy is to bound κ by working through the maximal admissible faces of each dimension. To avoid small-case irregularities, we assume that $n \geq 3$.

More specifically, each admissible vertex must belong to some maximal admissible face of dimension ≥ 0. We can therefore bound κ by (i) computing McMullen's bound for the number of vertices of each maximal admissible face, and then (ii) summing these bounds over all maximal admissible faces of all dimensions. We might count some vertices multiple times in this sum, but each vertex will be counted at least once.

We piece this sum together one dimension at a time, using Lemma 5.2 to bound the number of maximal admissible d-faces for each d.

- There are $\leq 3^{n-1}$ maximal admissible 0-faces, adding 3^{n-1} admissible vertices to our sum.
- There are $\leq 3^{n-2}$ maximal admissible 1-faces, adding $2 \cdot 3^{n-2}$ admissible vertices to our sum (since each 1-face is an edge, and has precisely two vertices).
- For each d in the range $2 \leq d \leq n-1$, there are $\leq 3^{n-1-d}$ maximal admissible d-faces. Each of these d-faces has at most n facets (Corollary 3.8) and therefore at most $M_{d,n}$ vertices (Theorem 2.3 and Lemma 4.1). This adds $\leq 3^{n-1-d} \cdot M_{d,n}$ admissible vertices to our sum.

By Lemma 5.1 there are no admissible d-faces for any dimension $d \geq n$, and so our final bound on κ becomes

$$\begin{aligned}
\kappa \quad &\leq \quad 3^{n-1} + 2 \cdot 3^{n-2} + \sum_{d=2}^{n-1} 3^{n-1-d} \cdot M_{d,n} \\
&= \quad 3^{n-1} + 2 \cdot 3^{n-2} + 3^{n-1} \sum_{d=2}^{n-1} (1/3)^d \cdot M_{d,n} \\
&\in \quad O\left(3^n + 3^n \cdot \left[\frac{1+\sqrt{1+4/9}}{2}\right]^n\right),
\end{aligned}$$

using the asymptotic bound from Corollary 4.4. The second term in this final expression dominates the first, and we have

$$\kappa \in O\left(\left[3 \cdot \frac{1+\sqrt{13/9}}{2}\right]^n\right) = O\left(\left[\frac{3+\sqrt{13}}{2}\right]^n\right) \simeq O(3.303^n).$$

□

6 The standard solution set

Having established new bounds for the quadrilateral projective solution space $\mathscr{Q} \subseteq \mathbb{R}^{3n}$, we can now transport this information to the standard projective solution space $\mathscr{S} \subseteq \mathbb{R}^{7n}$.

As noted in the introduction, the first upper bound on the number of admissible vertices of $\mathscr{S}$ was 128^n, proven by Hass et al. [14]. The best bound known to date is approximately $O(29.03^n)$, proven by the author [4]. The argument by Hass et al. relies on the fact that each vertex can be described as an intersection of facets of $\mathscr{S}$, and with $\leq 7n$ facets there can be at most $2^{7n} = 128^n$ such intersections. The bound of $O(29.03^n)$ was obtained by deriving a simple asymptotic extension to McMullen's upper bound theorem.

In this paper we tighten the best upper bound in standard coordinates to approximately $O(14.556^n)$ admissible vertices. Our strategy is to draw on our earlier results in quadrilateral coordinates. We begin by describing a bijection between maximal admissible faces of $\mathscr{Q}$ and $\mathscr{S}$, and then once again we aggregate over faces of varying dimensions.

As a further application of these techniques, we examine the special but important case of a one-vertex triangulation. The author sketches a proof in [6] that for a one-vertex triangulation the solution space $\mathscr{S}$ has approximately $O(15^n/\sqrt{n})$ vertices. Our final result of this paper is to tighten this bound to approximately $O(4.852^n)$.

Lemma 6.1. *Let v be the number of vertices in the underlying triangulation $\mathcal{T}$. Then there is a bijection between the maximal admissible faces of $\mathscr{Q}$ and the maximal admissible faces of $\mathscr{S}$ that maps i-faces of $\mathscr{Q}$ to $(i+v)$-faces of $\mathscr{S}$ for every i.*

Proof. For convenience we work in the solution cones $\mathscr{S}^\vee$ and $\mathscr{Q}^\vee$ instead of the projective solution spaces $\mathscr{S}$ and $\mathscr{Q}$; Lemma 3.6 shows this formulation to be equivalent. We establish our bijection in the direction from $\mathscr{S}^\vee$ to $\mathscr{Q}^\vee$ using the (linear) quadrilateral projection map $\pi\colon \mathbb{R}^{7n} \to \mathbb{R}^{3n}$. Recall from Section 2.1 that π is an onto map that preserves admissibility and inadmissibility, as well as compatibility and incompatibility.

We can apply the map π to sets of points (and in particular, faces of $\mathscr{S}^\vee$). Let $\pi(X)$ denote the image $\{\pi(\mathbf{x}) \mid \mathbf{x} \in X\}$ for any set $X \subseteq \mathscr{S}^\vee$. Although π might not map faces to faces in general, we claim that it does map *maximal admissible faces* of $\mathscr{S}^\vee$ to maximal admissible faces of $\mathscr{Q}^\vee$. Moreover, we claim that π is in fact the bijection that we seek. We prove these claims in stages.

- *π maps maximal admissible faces of $\mathscr{S}^\vee$ to maximal admissible faces of $\mathscr{Q}^\vee$.*

  Let F be a maximal admissible face of $\mathscr{S}^\vee$. Because π preserves admissibility and compatibility, all points in $\pi(F)$ are admissible and pairwise compatible. It follows from Lemma 3.9 that there is some maximal admissible face G of $\mathscr{Q}^\vee$ for which $\pi(F) \subseteq G$.

  If $\pi(F)$ is not itself a maximal admissible face then we can find some admissible point $\mathbf{g} \in G\backslash\pi(F)$. We know that $\mathbf{g}$ is compatible with every point in $\pi(F)$ (Lemma 3.7), and because π preserves inadmissibility and incompatibility it follows that every point in the preimage $\pi^{-1}(\mathbf{g}) \subseteq \mathscr{S}^\vee\backslash F$ is admissible and compatible with every point in F. This contradicts the assumption that F is a maximal admissible face of $\mathscr{S}^\vee$ (Corollary 3.12), and it follows that $\pi(F)$ must indeed be a maximal admissible face of $\mathscr{Q}^\vee$.

- *As a map between maximal admissible faces, π is one-to-one. That is, for every two distinct maximal admissible faces $F, G \subseteq \mathscr{S}^\vee$, we have $\pi(F) \neq \pi(G)$.*

  Let F and G be distinct maximal admissible faces of $\mathscr{S}^\vee$. By Corollary 3.12 there exist admissible and *incompatible* points $\mathbf{f} \in F$ and $\mathbf{g} \in G$. Because π preserves incompatibility it follows that $\pi(\mathbf{f})$ and $\pi(\mathbf{g})$ are incompatible points in $\mathscr{Q}^\vee$. That is, we have two incompatible points $\pi(\mathbf{f})$ and $\pi(\mathbf{g})$ in the maximal admissible faces $\pi(F)$ and $\pi(G)$ respectively, and from Corollary 3.12 again it follows that $\pi(F) \neq \pi(G)$.

- *As a map between maximal admissible faces, π is onto. That is, for every maximal admissible face $G \subseteq \mathscr{Q}^\vee$, there is a maximal admissible face $F \subseteq \mathscr{S}^\vee$ for which $\pi(F) = G$.*

  Let G be any maximal admissible face of $\mathscr{Q}^\vee$, and consider the preimage $\pi^{-1}(G)$. Because π preserves inadmissibility and incompatibility, $\pi^{-1}(G)$ must be a collection of admissible and pairwise compatible points in $\mathscr{S}^\vee$. By Lemma 3.9 there is some maximal admissible face $F \subseteq \mathscr{S}^\vee$ for which $F \supseteq \pi^{-1}(G)$. This gives us $\pi(F) \supseteq G$, and because both $\pi(F)$ and G are maximal admissible faces of $\mathscr{Q}^\vee$ it follows that $\pi(F) = G$.

This shows that π yields a bijection between the maximal admissible faces of $\mathscr{S}^\vee$ and the maximal admissible faces of $\mathscr{Q}^\vee$. All that remains now is to establish how π affects the dimensions of these faces.

Let F be some maximal admissible face in $\mathscr{S}^\vee$. We know from Section 2.1 that the kernel of the linear map π is generated by the v linearly independent vertex links (where v is the number of vertices in the underlying triangulation). Moreover, Corollary 3.11 shows that all v vertex links belong to the maximal admissible face F. Therefore we must have $\dim F = \dim \pi(F) + v$. $\square$

It should be noted that $\mathscr{Q}$ may contain no admissible points at all; in this case $\mathscr{Q}$ has a single maximal admissible face of dimension -1 (the empty face). In standard coordinates, $\mathscr{S}$ will always have admissible points (in particular, we always have the vertex links).

Now that we are equipped with this bijection, we aim to bound the dimensions of the maximal admissible faces of $\mathscr{S}$. To do this, we must place a bound on the number of vertices v of the underlying triangulation.

Lemma 6.2. *Any closed and connected 3-manifold triangulation with $n > 2$ tetrahedra can have at most $n + 1$ vertices.*

Proof. Let $\mathcal{T}$ be such a triangulation, and let G denote the *face pairing graph* of $\mathcal{T}$. This is the connected 4-valent multigraph whose vertices represent tetrahedra of $\mathcal{T}$ and whose edges represent identifications between tetrahedron faces (in particular, loops and multiple edges are allowed). See [2] for further discussion and explicit examples of face pairing graphs.[9]

Let S be a spanning tree within G, and let $\mathcal{T}_S$ denote the "partial triangulation" constructed from the same n tetrahedra by making *only* the face identifications described by the edges of S. This means that $\mathcal{T}_S$ is a connected simplicial complex formed from n tetrahedra by identifying precisely $n - 1$ pairs of faces. Moreover, the original triangulation $\mathcal{T}$ can be obtained from $\mathcal{T}_S$ by identifying the remaining $n + 1$ pairs of faces that correspond to the edges of $G\backslash S$. Figure 7 illustrates a face pairing graph G with a spanning tree S, and shows how the partial triangulation $\mathcal{T}_S$ might appear.

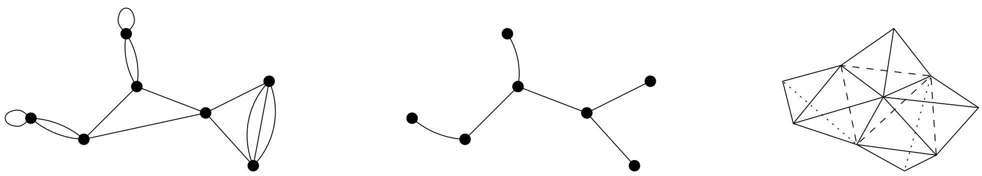

Face pairing graph G Spanning tree S Partial triangulation $\mathcal{T}_S$

Figure 7: The partial triangulation $\mathcal{T}_S$ corresponding to a spanning tree in G

Let v and v_S denote the number of vertices in $\mathcal{T}$ and $\mathcal{T}_S$ respectively. It is clear that $v \leq v_S$, since we obtain $\mathcal{T}$ from $\mathcal{T}_S$ by making additional face identifications (which may identify vertices of $\mathcal{T}_S$ together to reduce the total vertex count) but never adding new tetrahedra (and therefore never increasing the total vertex count).

It is straightforward to count the number of vertices in $\mathcal{T}_S$. Because S is a spanning tree, we construct $\mathcal{T}_S$ as follows:

- Begin with some initial tetrahedron Δ_1, which gives us four initial vertices for $\mathcal{T}_S$.

[9] G can also be thought of as the *dual 1-skeleton* of $\mathcal{T}$, with a *dual vertex* at the centre of every tetrahedron of $\mathcal{T}$ and a *dual edge* running through every face of $\mathcal{T}$.

- Follow by joining some new tetrahedron Δ_2 to Δ_1 along a single face. This introduces precisely one additional vertex to $\mathcal{T}_S$, since the other three vertices of Δ_2 (those on the joining face) become identified with the original vertices from Δ_1.
- Next we join some new tetrahedron Δ_3 to *either* Δ_2 or Δ_1 along a single face. Again this introduces precisely one new vertex to $\mathcal{T}_S$ (the vertex of Δ_3 not on the joining face).
- We continue this procedure, joining the remaining tetrahedra $\Delta_4, \ldots, \Delta_n$ into our structure along a single face each, creating one new vertex for $\mathcal{T}_S$ every time.

It follows that the number of vertices in $\mathcal{T}_S$ is precisely $v_S = n + 3$, and we obtain $v \leq n + 3$ as a result.

We can reduce our bound from $n + 3$ to $n + 1$ by studying the *leaves* of the tree S; that is, vertices of the tree with only one incident edge. Each leaf corresponds to a tetrahedron of $\mathcal{T}_S$ with only one face joined to the remainder of the structure. Moreover, the vertex opposite this face is not (yet) identified with any other vertices of any tetrahedron at all; we call this the *isolated vertex* of the leaf. This situation is illustrated in Figure 8.

Spanning tree S Partial triangulation $\mathcal{T}_S$

Figure 8: A tetrahedron of $\mathcal{T}_S$ corresponding to a leaf in the spanning tree S

Every tree of size $n > 2$ has at least two leaves; let ℓ be one such leaf, and let Δ_ℓ be the corresponding tetrahedron in $\mathcal{T}_S$. Consider the three faces of Δ_ℓ that surround the isolated vertex of ℓ. At least one of these faces must be joined to face of a *different* tetrahedron in the final triangulation $\mathcal{T}$; as a consequence, the isolated vertex of ℓ will be identified with some other tetrahedron vertex and we will have $v \leq v_S - 1 = n + 2$ vertices in total.

We can repeat this argument upon a second leaf ℓ' to lower our bound once more, establishing the final result $v \leq v_S - 2 = n + 1$. The only way this argument can fail is if both “new” vertex identifications are the same; that is, from our first leaf we find that the isolated vertex of ℓ is identified with the isolated vertex of ℓ', and then from our second leaf we find that the isolated vertex of ℓ' is identified with the isolated vertex of ℓ.

We are only forced into this redundancy if *every* additional edge from ℓ in the complementary graph $S \backslash G$ runs to ℓ' or is a loop back to ℓ; likewise, every additional edge from ℓ' in $S \backslash G$ must run to ℓ or be a loop back to ℓ'. In other words, we must have one of the two scenarios depicted in Figure 9.

Even still, we can avoid this redundancy if the tree S has three or more leaves (we simply replace ℓ' with a different selection). In fact, given that we can choose *any* spanning tree S, we are only forced into this redundancy if *every* spanning tree within G has precisely two leaves and gives one of the scenarios of Figure 9. The only such connected 4-valent multigraph G on $n > 2$ vertices is the graph depicted in Figure 10; that is, a single n-cycle with a loop at every vertex.

For such a face pairing graph we can lower our bound from $n+3$ to $n+1$ as follows. Let i be a non-leaf vertex of the tree S. The full graph G has a loop at vertex i, which means that two distinct vertices of the corresponding tetrahedron in $\mathcal{T}_S$ will be identified in the final triangulation $\mathcal{T}$. This identification does not involve the isolated vertices of the leaves,

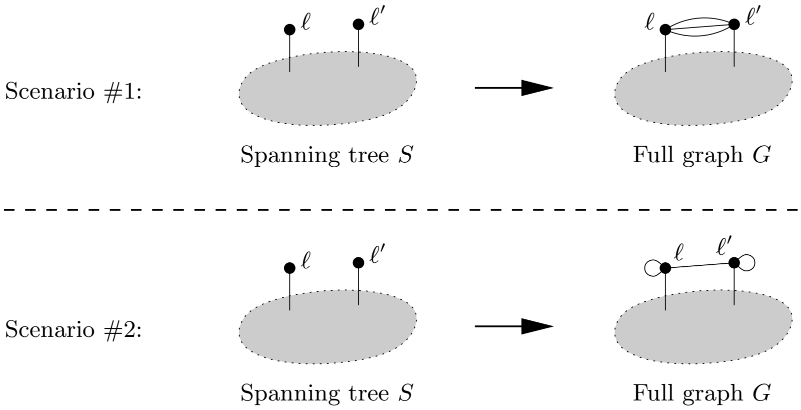

Figure 9: The two “redundant” scenarios in our analysis of leaves

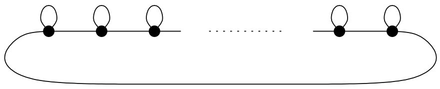

Figure 10: The only face pairing graph that forces redundancy in our leaf analysis

and so we can now return to our earlier argument on a single leaf to find a second (and different) identification between distinct vertices of $\mathcal{T}_S$, showing that $v \leq v_S - 2 = n+1$. $\square$

It can in fact be shown that this bound of $v \leq n+1$ is tight; the proof involves a general construction for arbitrary n, and we omit the details here. For $n = 2$ there is a closed 3-manifold triangulation with $n + 2 = 4$ vertices (this is the triangulation of the 3-sphere obtained by identifying the boundaries of two tetrahedra using the identity mapping).

We proceed now to the main result of this section, which is a new bound on the asymptotic growth rate of the size of the standard solution set (that is, the number of vertices of the standard projective solution space $\mathscr{S}$).

Theorem 6.3. *The size of the standard solution set is asymptotically bounded above by*

$$O\left(\left[9 \cdot \left(\frac{1+\sqrt{13/9}}{2}\right)^5\right]^n\right) \quad \simeq \quad O(14.556^n).$$

Proof. Let σ denote the number of admissible vertices of the standard projective solution space $\mathscr{S}$. Following the analogous result in quadrilateral coordinates (Theorem 5.4), our strategy is to bound σ by working through the maximal admissible faces of each dimension. As usual, we let v denote the number of vertices in the underlying triangulation $\mathcal{T}$.

Once again we assume that $n \geq 3$ to avoid small-case anomalies. Furthermore, we assume that the quadrilateral projective solution space $\mathscr{Q}$ has at least one admissible vertex (otherwise it is simple to show that there are precisely $v \leq n+1$ admissible vertices in $\mathscr{S}$, corresponding to the v vertex links in $\mathcal{T}$).

Let F be any maximal admissible face of $\mathscr{S}$. From Corollary 3.8 we know that F has at most $5n$ facets. Furthermore, Lemma 5.1 and Lemma 6.1 together show that F has dimension $d+v$ for some d in the range $0 \leq d \leq n-1$. Our immediate aim is to bound the number of vertices of F. There are two cases to consider:

- If $d > 0$ or $v > 1$ then the dimension of F is ≥ 2, and we can combine McMullen's theorem with Lemma 4.1 to show that F has at most $M_{d+v,5n}$ vertices. Using Lemma 6.2

we then have $d+v \le d+n+1 \le 2n < 5n/2$, whereupon Lemma 4.2 gives us $M_{d+v,5n} \le M_{d+n+1,5n}$. It follows that F has at most $M_{d+n+1,5n}$ vertices.

- If $d=0$ and $v=1$ then F is a 1-face (an edge) with precisely 2 vertices. It is simple to show that $2 \le M_{n+1,5n} = M_{d+n+1,5n}$, so again F has at most $M_{d+n+1,5n}$ vertices.

Once more we observe that each admissible vertex of $\mathscr{S}$ is a vertex of some maximal admissible face, and so we can bound σ by summing this bound of $M_{d+n+1,5n}$ over all maximal admissible faces. Lemma 5.2 and Lemma 6.1 together show that $\mathscr{S}$ has at most 3^{n-1-d} maximal admissible faces of dimension $d+v$ for each d, and so we have

$$\sigma \le \sum_{d=0}^{n-1} 3^{n-1-d} \cdot M_{d+n+1,5n} = \sum_{e=n+1}^{2n} 3^{2n-e} \cdot M_{e,5n}. \tag{6.1}$$

We can loosen this bound by extending the summation index e to the full range $2 \le e < 5n$, yielding

$$\sigma \le \sum_{e=2}^{5n-1} 3^{2n-e} \cdot M_{e,5n} = 9^n \sum_{e=2}^{5n-1} (1/3)^e \cdot M_{e,5n},$$

whereupon Corollary 4.4 gives us an asymptotic growth rate of

$$\sigma \in O\left(9^n \cdot \left[\frac{1+\sqrt{1+4/9}}{2}\right]^{5n}\right) = O\left(\left[9 \cdot \left(\frac{1+\sqrt{13/9}}{2}\right)^5\right]^n\right) \simeq O(14.556^n).$$

□

We finish this section by applying our techniques to the important case of a one-vertex triangulation. In this case we are able to strip an extra 3^n from our bound, yielding the following asymptotic result.

Theorem 6.4. *If we restrict our attention to triangulations with precisely one vertex, then the size of the standard solution set is asymptotically bounded above by*

$$O\left(\left[3 \cdot \left(\frac{1+\sqrt{13/9}}{2}\right)^5\right]^n\right) \quad \simeq \quad O(4.852^n).$$

Proof. The argument is almost identical to the proof of Theorem 6.3, and we do not repeat the details here. The main difference arises in the derivation of equation (6.1):

- For the case $d > 0$, we replace the bound $v \le n+1$ with the more precise $v = 1$, allowing us to replace the term $M_{d+n+1,5n}$ with the tighter $M_{d+1,5n}$.
- For the case $d = 0$, we cannot use McMullen's bound at all since we are looking at maximal admissible faces of dimension $d+v=1$. Instead we note that every 1-face is an edge with precisely two vertices, and we replace $M_{d+n+1,5n}$ with the constant 2.

Separating out the cases $d > 0$ and $d = 0$, equation (6.1) then becomes

$$\sigma \le 2 \cdot 3^{n-1} + \sum_{d=1}^{n-1} 3^{n-1-d} \cdot M_{d+1,5n} = \frac{2}{3} \cdot 3^n + \sum_{e=2}^{n} 3^{n-e} \cdot M_{e,5n}.$$

Again we extend the summation index e to the full range $2 \le e < 5n$, giving

$$\sigma \le \frac{2}{3} \cdot 3^n + \sum_{e=2}^{5n-1} 3^{n-e} \cdot M_{e,5n} = \frac{2}{3} \cdot 3^n + 3^n \sum_{e=2}^{5n-1} (1/3)^e \cdot M_{e,5n},$$

whereupon Corollary 4.4 shows the asymptotic growth rate to be

$$\sigma \in O\left(3^n + 3^n \cdot \left[\frac{1+\sqrt{1+4/9}}{2}\right]^{5n}\right) = O\left(\left[3 \cdot \left(\frac{1+\sqrt{13/9}}{2}\right)^5\right]^n\right) \simeq O(4.852^n).$$

□

7 Discussion

The complexity bounds of Sections 5 and 6 are significant improvements upon the prior state of the art. The reason for this success is because we have been able to integrate admissibility (in particular, the quadrilateral constraints) with the high-powered machinery of polytope theory (in particular, McMullen's upper bound theorem). Previous results have either used polytope theory on only a superficial level [14], or else drawn on deeper polytope theory but without any use of admissibility at all [4, 6].

The difficulty in integrating admissibility with polytope theory arises because the quadrilateral constraints are non-linear, and the admissible region of each projective solution space is far from being a convex polytope. In this paper we circumvent these difficulties by working with maximal admissible faces. However, this leads to certain inefficiencies, as we discuss further below.

It is known that any complexity bound on the size of the standard and quadrilateral solution sets must be exponential, even if we restrict our attention to one-vertex triangulations [4, 5]. However, the new bounds in this paper still leave significant room to move. In standard coordinates the worst known cases grow with complexity $O(17^{n/4}) \simeq O(2.03^n)$ in comparison to our $O(14.556^n)$; see [4] for details.[10] In quadrilateral coordinates, comprehensive experimental evidence from [5] suggests that the worst cases grow with complexity well below $O(\phi^n) \simeq O(1.618^n)$, in contrast to our current bound of $O(3.303^n)$.

This gap between theory and practice suggests that further research into theoretical bounds could be fruitful. The methods of this paper suggest several potential avenues for improvement:

- Because the proofs of Theorems 5.4 and 6.3 iterate through each maximal admissible face, it is likely that we count each admissible vertex many times over. Finding a mechanism to avoid this multiple-counting could help tighten our bounds further.

- The key to all of the new bounds in this paper is Lemma 5.2, where we show that $\mathscr{Q}$ has at most 3^{n-1-d} maximal admissible faces of each dimension $d \leq n-1$. This bound has been empirically tested against all ~ 150 million closed 3-manifold triangulations of size $n \leq 9$ (the same census used in [4]), with intriguing results.

  The outcomes of this testing are summarised in Table 1. For high dimensions $d \geq \frac{n}{2}-1$, the bound of $\leq 3^{n-1-d}$ maximal admissible faces appears to be tight (these numbers appear in bold in the table). For low dimensions $d < \frac{n}{2} - 1$ the number of maximal admissible faces drops away significantly, right down to what appears to be $O(n)$ maximal admissible faces of dimension 0.

  As an exploratory exercise, for each $n \leq 9$ we can work through the original proof of Theorem 5.4 but replace each bound of 3^{n-1-d} maximal admissible d-faces with the corresponding figure from Table 1. The resulting bounds on the number of admissible vertices of $\mathscr{Q}$ are shown in Table 2, and their growth rate settles down to roughly

[10] These cases are constructed and analysed for all $n > 5$, and experimental evidence supports the conjecture that these are the worst cases possible.

Number of tetrahedra (n)	Most maximal admissible faces of dimension ...									Number of triangulations
	0	1	2	3	4	5	6	7	8	
1	**1**									4
2	**3**	**1**								17
3	4	**3**	**1**							81
4	5	**9**	**3**	**1**						577
5	6	15	**9**	**3**	**1**					5 184
6	4	22	**27**	**9**	**3**	**1**				57 753
7	8	31	50	**27**	**9**	**3**	**1**			722 765
8	6	40	78	**81**	**27**	**9**	**3**	**1**		9 787 509
9	4	48	118	180	**81**	**27**	**9**	**3**	**1**	139 103 032

Table 1: The largest number of maximal admissible faces of various dimensions

Number of tetrahedra (n)	1	2	3	4	5	6	7	8	9
Max. number of admissible vertices	1	5	13	39	104	315	859	2458	7018

Table 2: Empirical complexity bounds based on the results of Table 1

$O(2.86^n)$, well below our current bound of $O(3.303^n)$. This suggests that, if we can tighten Lemma 5.2 for low dimensions, we can significantly improve our bounds again.

- Finally, even for high-dimensional faces where Lemma 5.2 does appear to be tight, we know from Lemma 5.3 that equality cannot hold for all high dimensions *simultaneously.* Empirical testing again suggests that Lemma 5.3 is merely one example of a larger set of constraints, and exploring these constraints may yield more useful information about the structure and number of maximal admissible faces.

For a final observation, we return to the worst known cases in standard coordinates. These are pathological triangulations of the 3-sphere for arbitrary $n > 5$, each with $O(17^{n/4})$ admissible vertices in $\mathscr{S}$, and there is strong empirical evidence [4] to suggest that this family of triangulations yields the largest number of vertices for all n.

What is interesting about these cases is each triangulation has *only one* maximal admissible face. In quadrilateral coordinates this maximal face is just an $(n-1)$-simplex, and the quadrilateral projective solution space $\mathscr{Q}$ has only n admissible vertices in total. In other words, for these cases the pathological complexity only appears in the extension to standard coordinates. These observations suggest that a better understanding of the relationships between the face lattices in $\mathscr{S}$ and $\mathscr{Q}$ could be an important step in achieving stronger bounds on the complexities of these polytopes.

Acknowledgements

The author is grateful to the Australian Research Council for their support under the Discovery Projects funding scheme (project DP1094516).

References

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

[2] Benjamin A. Burton, *Face pairing graphs and 3-manifold enumeration*, J. Knot Theory Ramifications **13** (2004), no. 8, 1057–1101.

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

[4] ———, *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.

[5] ———, *Extreme cases in normal surface enumeration*, In preparation, 2010.

[6] ———, *Optimizing the double description method for normal surface enumeration*, Math. Comp. **79** (2010), no. 269, 453–484.

[7] ———, *Quadrilateral-octagon coordinates for almost normal surfaces*, Experiment. Math. **19** (2010), no. 3, 285–315.

[8] Benjamin A. Burton, J. Hyam Rubinstein, and Stephan Tillmann, *The Weber-Seifert dodecahedral space is non-Haken*, To appear in Trans. Amer. Math. Soc., `arXiv:0909.4625`, September 2009.

[9] Patrick J. Callahan, Martin V. Hildebrand, and Jeffrey R. Weeks, *A census of cusped hyperbolic 3-manifolds*, Math. Comp. **68** (1999), no. 225, 321–332.

[10] M. E. Dyer, *The complexity of vertex enumeration methods*, Math. Oper. Res. **8** (1983), no. 3, 381–402.

[11] Branko Grünbaum, *Convex polytopes*, 2nd ed., Graduate Texts in Mathematics, no. 221, Springer, New York, 2003.

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

[13] ———, *Über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I*, Math. Z. **80** (1962), 89–120.

[14] 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.

[15] John Hempel, *3-manifolds*, Annals of Mathematics Studies, no. 86, Princeton University Press, Princeton, NJ, 1976.

[16] William Jaco and Ulrich Oertel, *An algorithm to decide if a* 3*-manifold is a Haken manifold*, Topology **23** (1984), no. 2, 195–209.

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

[18] 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.

[19] Ensil Kang and J. Hyam Rubinstein, *Ideal triangulations of 3-manifolds I; Spun normal surface theory*, Proceedings of the Casson Fest, Geom. Topol. Monogr., vol. 7, Geom. Topol. Publ., Coventry, 2004, pp. 235–265.

[20] Leonid Khachiyan, Endre Boros, Konrad Borys, Khaled Elbassioni, and Vladimir Gurvich, *Generating all vertices of a polyhedron is hard*, Discrete Comput. Geom. **39** (2008), no. 1-3, 174–190.

[21] Hellmuth Kneser, *Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten*, Jahresbericht der Deut. Math. Verein. **38** (1929), 248–260.

[22] Bruno Martelli and Carlo Petronio, *A new decomposition theorem for 3-manifolds*, Illinois J. Math. **46** (2002), 755–780.

[23] Sergei V. Matveev, *Complexity theory of three-dimensional manifolds*, Acta Appl. Math. **19** (1990), no. 2, 101–130.

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

[25] J. Hyam Rubinstein, *An algorithm to recognize the* 3*-sphere*, Proceedings of the International Congress of Mathematicians (Zürich, 1994), vol. 1, Birkhäuser, 1995, pp. 601–611.

[26] Richard P. Stanley, *Enumerative combinatorics, Vol. 1*, Cambridge Studies in Advanced Mathematics, no. 49, Cambridge University Press, Cambridge, 1997.

[27] Abigail Thompson, *Thin position and the recognition problem for S^3*, Math. Res. Lett. **1** (1994), no. 5, 613–630.

[28] Stephan Tillmann, *Normal surfaces in topologically finite 3-manifolds*, Enseign. Math. (2) **54** (2008), 329–380.

[29] Jeffrey L. Tollefson, *Normal surface Q-theory*, Pacific J. Math. **183** (1998), no. 2, 359–374.

[30] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, no. 152, Springer-Verlag, New York, 1995.

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