

Quadrilateral–Octagon Coordinates for Almost Normal Surfaces

Benjamin A. Burton

CONTENTS

- 1. Introduction
- 2. Normal Surfaces
- 3. Almost Normal Surfaces
- 4. Quadrilateral–octagon Coordinates
- 5. 3-Sphere Recognition
- 6. Enumeration Algorithms
- 7. Measuring Performance
- 8. Joint Coordinates
- Acknowledgments
- References

Normal and almost normal surfaces are essential tools for algorithmic 3-manifold topology, but to use them requires exponentially slow enumeration algorithms in a high-dimensional vector space. The quadrilateral coordinates of Tollefson alleviate this problem considerably for normal surfaces by reducing the dimension of this vector space from $7n$ to $3n$ (where n is the complexity of the underlying triangulation). Here we develop an analogous theory for octagonal almost normal surfaces, using quadrilateral and octagon coordinates to reduce this dimension from $10n$ to $6n$. As an application, we show that quadrilateral–octagon coordinates can be used exclusively in the streamlined 3-sphere recognition algorithm of Jaco, Rubinstein, and Thompson, reducing experimental running times by factors of thousands. We also introduce joint coordinates, a system with only $3n$ dimensions for octagonal almost normal surfaces that has appealing geometric properties.

1. INTRODUCTION

The theory of normal surfaces, introduced in [Kneser 29] and developed in [Haken 61, Haken 62], is central to algorithmic 3-manifold topology. In essence, normal surface theory allows us to search for “interesting” embedded surfaces within a 3-manifold triangulation \mathcal{T} by enumerating the vertices of a polytope in a high-dimensional vector space. Normal surfaces are defined by their intersections with the tetrahedra of \mathcal{T} , which must be collections of disjoint triangles and/or quadrilaterals, collectively referred to as *normal disks*.

In the early 1990s, Rubinstein introduced the concept of an almost normal surface, for use with problems such as 3-sphere recognition and finding Heegaard splittings [Rubinstein 97]. Almost normal surfaces are essentially normal surfaces with a single unusual intersection piece, which may be either an octagon or a tube. Thompson subsequently refined the 3-sphere recognition algorithm to remove any need for tubes [Thompson 94], and since then almost normal surfaces have appeared in algorithms

2000 AMS Subject Classification: Primary 57N10; Secondary 57Q35

Keywords: Normal surfaces, almost normal surfaces, quadrilateral–octagon coordinates, joint coordinates, Q-theory, 3-sphere recognition

such as determining Heegaard genus [Lackenby 08], recognizing small Seifert fibered spaces [Rubinstein 04], and finding bridge surfaces in knot complements [Wilson 08].

In this paper we focus on *octagonal* almost normal surfaces, that is, almost normal surfaces in which the unusual intersection piece is an octagon, not a tube. The reason for this restriction is that octagonal almost normal surfaces are both tractable and useful, and have important applications beyond 3-manifold topology. In detail:

- For practical computation, octagonal almost normal surfaces are significantly easier to deal with than general almost normal surfaces. In particular, the translation between surfaces and high-dimensional vectors becomes much simpler, and the enumeration of these vectors is less fraught with complications.
- As shown by Thompson, octagonal almost normal surfaces are sufficient for running the 3-sphere recognition algorithm [Thompson 94].
- Following on from the previous point, an efficient 3-sphere recognition algorithm is important for computation in *4-manifold* topology. For example, answering even the basic question, “is \mathcal{T} a 4-manifold triangulation?” requires us to run the 3-sphere recognition algorithm over a neighborhood of each vertex of \mathcal{T} . Therefore, improving the efficiency of 3-sphere recognition is an important step toward a general efficient computational framework for working with 4-manifold triangulations.

As suggested above, our focus here is on the *efficiency* of working with almost normal surfaces. The fundamental problem that we face is that the underlying polytope vertex enumeration can grow to become exponentially slow in the number of tetrahedra. This means that in practice, normal surface algorithms cannot be run on large triangulations. Moreover, this exponential growth is not the fault of the algorithms, but an unavoidable feature of the problems that they try to solve. For illustrations of this, see [Burton 10b], which describes cases in which the underlying vertex enumeration problem has exponentially many solutions, or see the proof in [Agol et al. 02] that computing a 3-manifold knot genus (one of the many applications of normal surface theory) is NP-complete.

For almost normal surfaces, our efficiency troubles are even worse than for normal surfaces. This is because the polytope vertex enumeration is exponentially slow not just in the number of tetrahedra n , but also in the dimension of the underlying vector space. For normal

surfaces this dimension is $7n$, whereas for octagonal almost normal surfaces this dimension is $10n$, a significant difference when one is dealing with an exponential algorithm.

In the realm of normal surfaces, much progress has been made in improving the efficiency of enumeration algorithms [Burton 09b, Burton 10a, Tollefson 98]. One key development has been Tollefson’s quadrilateral coordinates [Tollefson 98], in which we work only with quadrilateral normal disks and then reconstruct the triangular disks afterward. This allows us to perform our expensive polytope vertex enumeration in dimension $3n$ instead of $7n$, which yields substantial efficiency improvements.

There are two complications with Tollefson’s approach:

- When reconstructing a normal surface from its quadrilateral disks, we cannot recover any vertex-linking components (these components lie at the frontiers of small regular neighborhoods of vertices of the triangulation). This is typically not a problem, since such components are rarely of interest.
- When we use quadrilateral coordinates for the underlying polytope vertex enumeration, some solutions are “lost.” That is, the resulting set of normal surfaces (called *vertex normal surfaces*) is a strict subset of what we would obtain using the traditional $7n$ -dimensional framework of Haken.

This latter issue can be resolved in two different ways. For some high-level topological algorithms, such as the detection of two-sided incompressible surfaces [Tollefson 98], it has been proven that at least one of the surfaces that we need to find will not be lost. As a more general resolution to this problem, there is a fast quadrilateral-to-standard conversion algorithm through which we can recover all of the lost surfaces [Burton 09b].

The main purpose of this paper is to develop an analogous theory for octagonal almost normal surfaces. Specifically, we show that we can work with only quadrilateral normal disks and octagonal almost normal disks, and then reconstruct the triangular disks afterwards. As a consequence, the dimension for our vertex enumeration drops from $10n$ to $6n$.

We run into the same complications as before—vertex-linking components cannot be recovered, and we may lose some of our original solutions. Here we show that as with quadrilateral coordinates, these are not serious problems. In particular, we show that despite this loss

of information, quadrilateral–octagon coordinates suffice for the 3-sphere recognition algorithm. More generally, we observe that the fast quadrilateral-to-standard conversion algorithm of [Burton 09b] works seamlessly with octagonal almost normal surfaces.

As a practical measure of benefit, we use the software package Regina [Burton 04, Burton 09a] to compare running times for the 3-sphere recognition algorithm with and without quadrilateral–octagon coordinates. Here we see quadrilateral–octagon coordinates improving performance by factors of thousands in several cases. Readers can experiment with quadrilateral–octagon coordinates for themselves by downloading Regina version 4.6 or later.

We finish this paper by introducing *joint coordinates*, in which we exploit natural relationships between quadrilaterals and octagons to reduce our $6n$ dimensions for octagonal almost normal surfaces down to just $3n$ dimensions. Although these coordinates cannot be used with existing enumeration algorithms (due to a loss of convexity in the underlying polytope), they have appealing geometric properties that make them useful for visualization, and which may help develop intuition about the structure of the corresponding solution space.

All of the results in this paper apply only to compact 3-manifold triangulations. In particular, they do not cover the ideal triangulations of [Thurston 78], where the reconstruction of triangular disks can result in pathological (but nevertheless useful) objects such as spun normal surfaces, which contain infinitely many disks [Tillmann 08].

The layout of this paper is as follows. Section 2 begins with an overview of normal surfaces and Tollefson’s quadrilateral coordinates, and Section 3 follows with an overview of almost normal surfaces. In Section 4 we develop the core theory for quadrilateral–octagon coordinates, including necessary and sufficient conditions for a $6n$ -dimensional vector to represent an octagonal almost normal surface.

For the remainder of the paper we focus on applications and extensions of this theory. In Section 5 we describe the streamlined 3-sphere recognition algorithm of Jaco, Rubinstein, and Thompson [Jaco and Rubinstein 03], and show that this algorithm remains correct when we work in quadrilateral–octagon coordinates instead of the original $10n$ -dimensional vector space. Section 6 focuses on the underlying polytope vertex enumeration algorithm, where we observe that state-of-the-art algorithms for enumerating normal surfaces [Burton 09b, Burton 10a] can be used seamlessly with

octagonal almost normal surfaces and quadrilateral–octagon coordinates. In Section 7 we offer experimental measures of running time that show how quadrilateral–octagon coordinates improve the 3-sphere recognition algorithm in practice, and in Section 8 we finish with a discussion of joint coordinates.

2. NORMAL SURFACES

We assume that the reader is already familiar with the theory of normal surfaces (if not, a good overview can be found in [Hass et al. 99]). In this section we outline the relevant aspects of the theory, concentrating on the differences between Haken’s original formulation [Haken 61] and Tollefson’s quadrilateral coordinates [Tollefson 98]. For a more detailed discussion of these two formulations and the relationships between them, the reader is referred to [Burton 09b].

Throughout this paper we assume that we are working with a compact 3-manifold triangulation \mathcal{T} formed from n tetrahedra. By a *compact triangulation*, we mean that every vertex of \mathcal{T} has a small neighborhood whose frontier is a sphere or a disk. This ensures that \mathcal{T} is a triangulation of a compact 3-manifold (possibly with boundary), and rules out the ideal triangulations of [Thurston 78] in which vertices form higher-genus cusps.

To help keep the number of tetrahedra in \mathcal{T} small, we allow different faces of a tetrahedron to be identified (and likewise with edges and vertices). Some authors refer to triangulations with this property as *pseudotriangulations* or *semisimplicial triangulations*. Faces, edges, and vertices of \mathcal{T} that lie entirely within the 3-manifold boundary are called *boundary faces*, *boundary edges*, and *boundary vertices* of \mathcal{T} respectively.

An *embedded normal surface* in \mathcal{T} is a properly embedded surface (possibly disconnected or empty) that intersects each tetrahedron of \mathcal{T} in a collection of disjoint *normal disks*. Each normal disk is either a *triangle* or a *quadrilateral*, with a boundary consisting of three or four




FIGURE 1. Several normal disks within a single tetrahedron.




FIGURE 2. The seven different normal disk types within a tetrahedron.

arcs respectively that cross distinct faces of the tetrahedron. Figure 1 illustrates several disjoint triangles and quadrilaterals within a tetrahedron.

The triangles and quadrilaterals within a tetrahedron can be grouped into seven *normal disk types*, according to which edges of the tetrahedron they intersect. This includes four triangular disk types and three quadrilateral disk types, all of which are illustrated in Figure 2.

Equivalence of normal surfaces is defined by *normal isotopy*, which is an ambient isotopy that preserves each simplex of the triangulation \mathcal{T} . Throughout this paper, any two surfaces that are related by normal isotopy are regarded as the same surface.

Vertex links are normal surfaces that play an important role in the discussion that follows. If V is a vertex of the triangulation \mathcal{T} , then the *vertex link* of V , denoted by $\ell(V)$, is the normal surface at the frontier of a small regular neighborhood of V . This surface is formed entirely from triangular disks (one copy of each triangular disk type surrounding V). Here we follow the nomenclature of [Jaco and Rubinstein 03]; Tollefson refers to vertex links as *trivial surfaces*.

A core strength of normal surface theory is its ability to reduce difficult problems in topology to simpler problems in linear algebra. This is where the formulations of Haken and Tollefson differ, and so we slow down from here onward to give full details. The key difference between the two formulations is that Haken works in a $7n$ -dimensional vector space with coordinates based on triangle and quadrilateral disk types, whereas Tollefson works in a $3n$ -dimensional space based on quadrilateral disk types only.

Definition 2.1. (Vector representations.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let S be an embedded normal surface in \mathcal{T} . For each tetrahedron Δ_i , let $t_{i,1}, t_{i,2}, t_{i,3}$, and $t_{i,4}$ denote the number of triangular disks of S of each type in Δ_i , and let $q_{i,1}, q_{i,2}$, and $q_{i,3}$ denote the number of quadrilateral disks of S of each type in Δ_i .

Then the *standard vector representation* of S , denoted by $\mathbf{v}(S)$, is the $7n$ -dimensional vector

$$\mathbf{v}(S) = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}; t_{2,1}, t_{2,2}, t_{2,3}, t_{2,4}, q_{2,1}, q_{2,2}, q_{2,3}; \dots, q_{n,3}),$$

and the *quadrilateral vector representation* of S , denoted by $\mathbf{q}(S)$, is the $3n$ -dimensional vector

$$\mathbf{q}(S) = (q_{1,1}, q_{1,2}, q_{1,3}; q_{2,1}, q_{2,2}, q_{2,3}; \dots, q_{n,3}).$$

When we are working with $\mathbf{v}(S)$, we say we are working in *standard coordinates* (or *standard normal coordinates* if we wish to distinguish between normal and almost normal surfaces). Likewise, when working with $\mathbf{q}(S)$ we say we are working in *quadrilateral coordinates*. The following uniqueness results are from [Haken 61, Tollefson 98].

Lemma 2.2. *Let \mathcal{T} be a compact 3-manifold triangulation, and let S and S' be embedded normal surfaces in \mathcal{T} .*

- *The standard vector representations $\mathbf{v}(S)$ and $\mathbf{v}(S')$ are equal if and only if the surfaces S and S' are normal isotopic (i.e., they are the “same” normal surface).*
- *The quadrilateral vector representations $\mathbf{q}(S)$ and $\mathbf{q}(S')$ are equal if and only if either (i) S and S' are normal isotopic, or (ii) S and S' can be made normal isotopic by adding or removing vertex-linking components.*

Since we are rarely interested in vertex-linking components, Lemma 2.2 shows that the standard and quadrilateral vector representations each contain everything we might want to know about an embedded normal surface.

Not every integer vector $\mathbf{w} \in \mathbb{R}^{7n}$ or $\mathbf{w} \in \mathbb{R}^{3n}$ is the vector representation of a normal surface. The necessary conditions on \mathbf{w} include a set of *matching equations* as




FIGURE 3. Building the standard matching equations.

well as a set of *quadrilateral constraints*, which we define as follows.

Definition 2.3. (Standard matching equations.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let $\mathbf{w} \in \mathbb{R}^{7n}$ be any $7n$ -dimensional vector whose coordinates we label

$$\mathbf{w} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}; t_{2,1}, t_{2,2}, t_{2,3}, t_{2,4}, q_{2,1}, q_{2,2}, q_{2,3}; \dots, q_{n,3}).$$

For each nonboundary face of \mathcal{T} and each of the three edges surrounding it, we obtain a *standard matching equation* on \mathbf{w} as follows.

Let F be some nonboundary face of \mathcal{T} , and let e be one of the three edges surrounding F . Suppose that Δ_i and Δ_j are the two tetrahedra on either side of F . Then there is precisely one triangular disk type and one quadrilateral disk type in each of Δ_i and Δ_j that meets F in an arc parallel to e , as illustrated in Figure 3. Suppose these disk types correspond to coordinates $t_{i,a}$, $q_{i,b}$, $t_{j,c}$, and $q_{j,d}$ respectively. Then we obtain the matching equation

$$t_{i,a} + q_{i,b} = t_{j,c} + q_{j,d}.$$

Essentially, the standard matching equations ensure that all of the normal disks on either side of a nonboundary face F can be joined together. In Figure 3, the four coordinates are $(t_{i,a}, q_{i,b}, t_{j,c}, q_{j,d}) = (1, 2, 2, 1)$, giving the equation $1+2 = 2+1$, which is indeed satisfied. If \mathcal{T} is a closed triangulation (i.e., it has no boundary), then there are precisely $6n$ standard matching equations for \mathcal{T} (three for each of the $2n$ faces of \mathcal{T}).

Definition 2.4. (Quadrilateral matching equations.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let $\mathbf{w} \in \mathbb{R}^{3n}$ be any $3n$ -dimensional vector, whose coordinates we label

$$\mathbf{w} = (q_{1,1}, q_{1,2}, q_{1,3}; q_{2,1}, q_{2,2}, q_{2,3}; \dots, q_{n,3}).$$

For each nonboundary edge of \mathcal{T} , we obtain a *quadrilateral matching equation* on \mathbf{w} as follows.

Let e be some nonboundary edge of \mathcal{T} , and arbitrarily label the two ends of e *upper* and *lower*. The tetrahedra containing edge e are arranged in a cycle around e , as illustrated in Figure 4. Choose some arbitrary direction around this cycle, and suppose that the tetrahedra that we encounter as we travel in this direction around the cycle are labeled $\Delta_{i_1}, \dots, \Delta_{i_t}$.

For each tetrahedron in this cycle, there are two quadrilateral types meeting edge e : one that rises from the lower end of e to the upper as we travel around the cycle in the chosen direction, and one that falls from the upper end of e to the lower. We call these the *upward quadrilaterals* and *downward quadrilaterals* respectively; these are again illustrated in Figure 4.

Suppose now that the coordinates corresponding to the upward and downward quadrilateral types are $q_{i_1, u_1}, q_{i_2, u_2}, \dots, q_{i_t, u_t}$ and $q_{i_1, d_1}, q_{i_2, d_2}, \dots, q_{i_t, d_t}$ respectively. Then we obtain the matching equation

$$q_{i_1, u_1} + q_{i_2, u_2} + \dots + q_{i_t, u_t} = q_{i_1, d_1} + q_{i_2, d_2} + \dots + q_{i_t, d_t}. \tag{2-1}$$

In other words, the total number of upward quadrilaterals surrounding e equals the total number of downward quadrilaterals surrounding e .

Note that a single tetrahedron might appear multiple times in the cycle around e , in which case a single coordinate $q_{i,j}$ might appear more than once in the equation (2-1). For a closed triangulation \mathcal{T} with v vertices, a quick Euler characteristic calculation shows that we have precisely $n + v$ edges in our triangulation and therefore precisely $n + v$ quadrilateral matching equations.

Definition 2.5. (Quadrilateral constraints.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and consider any vector

$$\mathbf{w} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}; \dots, q_{n,3}) \in \mathbb{R}^{7n}$$

or

$$\mathbf{w} = (q_{1,1}, q_{1,2}, q_{1,3}; \dots, q_{n,3}) \in \mathbb{R}^{3n}.$$

We say that \mathbf{w} satisfies the *quadrilateral constraints* if for every tetrahedron Δ_i , at most one of the quadrilateral coordinates $q_{i,1}$, $q_{i,2}$, and $q_{i,3}$ is nonzero.

We can now describe a full set of necessary and sufficient conditions for a vector $\mathbf{w} \in \mathbb{R}^{7n}$ or $\mathbf{w} \in \mathbb{R}^{3n}$ to




FIGURE 4. Building the quadrilateral matching equations.

be the vector representation of some embedded normal surface. The following result is from [Haken 61, Tollefson 98].

Theorem 2.6. *Let \mathcal{T} be a compact 3-manifold triangulation formed from n tetrahedra. An integer vector ($\mathbf{w} \in \mathbb{R}^{7n}$ or $\mathbf{w} \in \mathbb{R}^{3n}$) is the (standard or quadrilateral) vector representation of an embedded normal surface in \mathcal{T} if and only if:*

- the coordinates of \mathbf{w} are all nonnegative;
- \mathbf{w} satisfies the (standard or quadrilateral) matching equations for \mathcal{T} ;
- \mathbf{w} satisfies the quadrilateral constraints for \mathcal{T} .

Such a vector is referred to as an admissible vector.¹

Essentially, the nonnegativity constraint ensures that the coordinates of \mathbf{w} can be used to count normal disks, the matching equations ensure that these disks can be joined together to form a surface, and the quadrilateral constraint ensures that this surface is embedded (since any two quadrilaterals of different types within the same tetrahedron must intersect).

Many high-level algorithms in 3-manifold topology involve the enumeration of *vertex normal surfaces*, which form a basis from which we can reconstruct all embedded normal surfaces within a triangulation \mathcal{T} . The relevant definitions are as follows.

Definition 2.7. (Projective solution space.) Let \mathcal{T} be a compact 3-manifold triangulation formed from n tetrahedra. The set of all nonnegative vectors in \mathbb{R}^{7n} that satisfy the standard matching equations for \mathcal{T} forms a rational

¹It is sometimes useful to extend the concept of admissibility to rational vectors or even real vectors in \mathbb{R}^{7n} or \mathbb{R}^{3n} , as seen for instance in [Burton 09b]. However, we do not need such extensions in this paper.

polyhedral cone in \mathbb{R}^{7n} . The *standard projective solution space* for \mathcal{T} is the rational polytope formed by intersecting this cone with the hyperplane $\{\mathbf{w} \in \mathbb{R}^{7n} \mid \sum w_i = 1\}$.

The *quadrilateral projective solution space* for \mathcal{T} is defined in a similar fashion by working in \mathbb{R}^{3n} and using the quadrilateral matching equations instead.

Definition 2.8. (Vertex normal surface.) Let \mathcal{T} be a compact 3-manifold triangulation, and let S be an embedded normal surface in \mathcal{T} . If the standard vector representation $\mathbf{v}(S)$ is a positive multiple of some vertex of the standard projective solution space, then we call S a *standard vertex normal surface*. Likewise, if the quadrilateral vector representation $\mathbf{q}(S)$ is a positive multiple of some vertex of the quadrilateral projective solution space, then we call S a *quadrilateral vertex normal surface*.

It should be noted that the definition of a vertex normal surface varies among authors. Definition 2.8 is consistent with [Jaco and Rubinstein 03], as well as earlier work of this author [Burton 09b]. Other authors impose additional conditions, such as Tollefson, who requires S to be connected and two-sided [Tollefson 98], or Jaco and Oertel, who require the elements of $\mathbf{v}(S)$ to have no common factor and who use the alternative name *fundamental edge surface* [Jaco and Oertel 84].

Although vertex normal surfaces can be used as a basis for reconstructing all embedded normal surfaces within a triangulation, this is typically not feasible, since there are infinitely many such surfaces. Instead we frequently find that when searching for an embedded normal surface with some desirable property, we can restrict our attention to *only* vertex normal surfaces. For instance, Jaco and Oertel prove for closed irreducible 3-manifolds that if a two-sided incompressible surface exists, then one can be found as a standard vertex normal surface [Jaco and Oertel 84]. Likewise, Jaco and Tollefson prove that if a 3-manifold contains an essential disk or sphere, then one

can be found as a standard vertex normal surface [Jaco and Tollefson 95].

Using results of this type, a typical high-level algorithm based on normal surface theory includes the following steps:

- (i) Enumerate the (finitely many) vertices of the projective solution space for a given triangulation \mathcal{T} , using techniques from linear programming (see [Burton 10a] for details).
- (ii) Eliminate those vertices that do not satisfy the quadrilateral constraints, and then reconstruct the vertex normal surfaces of \mathcal{T} by taking multiples of those vertices that remain. Although there are infinitely many such multiples, only finitely many will yield *connected* normal surfaces, which is typically what we are searching for.
- (iii) Test each of these vertex normal surfaces for some desirable property (such as incompressibility, or being an essential disk or sphere).

Here we can see the real benefit of working in quadrilateral coordinates—the enumeration of step (i) takes place in a vector space of dimension $7n$ for standard coordinates, but only $3n$ for quadrilateral coordinates. Since both the running time and memory use can become exponential in this dimension [Burton 10a], a reduction from $7n$ to $3n$ can yield dramatic improvements in performance.

However, there is a trade-off for using quadrilateral coordinates. Although every connected quadrilateral vertex normal surface is also a standard vertex normal surface [Burton 09b], the converse is not true in general. Instead, there might be standard vertex normal surfaces (perhaps including the incompressible surfaces, essential disks and spheres, or whatever else we are searching for) that do not show up as quadrilateral vertex normal surfaces. These “lost surfaces” can undermine the correctness of our algorithms, which we maintain in one of two ways:

- We can resolve the problem using theory. This requires us to prove that if the surface we are searching for exists, then it exists not only as a *standard* vertex normal surface but also as a *quadrilateral* vertex normal surface.

Such results can be more difficult to prove in quadrilateral coordinates than in standard coordinates, partly because important functions such as Euler characteristic are no longer linear. Nevertheless, examples can be found—[Tollefson 98] proves such

a result for two-sided incompressible surfaces, and [Jaco et al. 02] refers to similar results for essential disks and spheres.

- We can resolve the problem using algorithms and computation. There is a fast algorithm described in [Burton 09b] that converts a full set of quadrilateral vertex normal surfaces to a full set of standard vertex normal surfaces, thereby recovering those surfaces that were lost. This algorithm is found to have a negligible running time, which means that we are able to work with standard vertex normal surfaces yet still enjoy the significantly greater performance of quadrilateral coordinates.

The main part of this paper is concerned with the development of quadrilateral–octagon coordinates for almost normal surfaces, where we face a similar trade-off. In Section 5 we resolve this problem for the 3-sphere recognition algorithm using the theoretical route, and in Section 6 we show how the more general algorithmic solution can be used.

3. ALMOST NORMAL SURFACES

Almost normal surfaces are an extension of normal surfaces whereby in addition to the usual normal disks, we allow one tetrahedron of the triangulation to contain a single unusual intersection piece. Introduced by Rubinstein for use with the 3-sphere algorithm and related problems [Rubinstein 95, Rubinstein 97], almost normal surfaces also enjoy other applications such as the determination of Heegaard genus [Lackenby 08], the recognition of small Seifert fibered spaces [Rubinstein 04], and finding bridge surfaces in knot complements [Wilson 08].

We begin this section by defining almost normal surfaces, whereupon we restrict our attention to *octagonal* almost normal surfaces. Octagonal almost normal surfaces are significantly easier to deal with, and Thompson has proven that they are sufficient for use with the 3-sphere recognition algorithm [Thompson 94].

In the remainder of this section, we define concepts similar to those seen in Section 2, such as vector representation, matching equations, and vertex almost normal surfaces. These concepts and their corresponding results are well-known extensions to traditional normal surface theory; see [Lackenby 08] or [Rubinstein 97] for a brief sketch. The details, however, are not explicitly laid down in the current literature, and so we present these details here.




FIGURE 5. Examples of exceptional pieces in almost normal surfaces.

Definition 3.1. (Almost normal surface.) Let \mathcal{T} be a compact 3-manifold triangulation, and let Δ be some tetrahedron of \mathcal{T} . A *normal octagon* in Δ is a properly embedded disk in Δ whose boundary consists of eight normal arcs running across the faces of Δ , as illustrated in Figure 5. A *normal tube* in Δ is a properly embedded annulus in Δ consisting of any two disjoint normal disks joined by an unknotted tube, again illustrated in Figure 5.

An *almost normal surface* in \mathcal{T} is a properly embedded surface whose intersection with the tetrahedra of \mathcal{T} consists of (i) zero or more normal disks, plus (ii) in precisely one tetrahedron of \mathcal{T} , either a single normal octagon or a single normal tube (but not both).² This single octagon or tube is referred to as the *exceptional piece* of the almost normal surface.

Although Definition 3.1 requires that almost normal surfaces be properly embedded, for brevity's sake we do not include the word “embedded” in their name. For the remainder of this paper we concern ourselves only with octagonal almost normal surfaces, which are defined as follows.

Definition 3.2. (Octagonal almost normal surface.) An *octagonal almost normal surface* is an almost normal surface whose exceptional piece is a normal octagon (not a tube). For contrast, we will often refer to the almost normal surfaces of Definition 3.1 (where the exceptional piece may be either an octagon or a tube) as *general almost normal surfaces*.

The possible normal octagons within a tetrahedron can be grouped into three *octagon types*, according to how many times they intersect each edge of the tetrahedron. All three octagon types are illustrated in Figure 6.

²Jaco and Rubinstein add the additional constraint that the tube does not join two copies of the same normal surface [Jaco and Rubinstein 03].

As with “embedded,” we will sometimes drop the word “octagonal” from definitions to avoid excessively long names; see for instance the *standard almost normal matching equations* and *vertex almost normal surfaces* (Definitions 3.3 and 3.5), which refer exclusively to octagonal almost normal surfaces.

At this early stage we can already see one reason why octagonal almost normal surfaces are substantially easier to deal with than general almost normal surfaces—while there are only three octagon types within a tetrahedron, there are 25 distinct types of normal tube, giving 28 types of exceptional piece in the general case. Not only is this messier to implement on a computer, but it can lead to significant increases in running time and memory usage. We return to this issue at the end of this section.

Definition 3.3. (Standard vector representation.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let S be an octagonal almost normal surface in \mathcal{T} . For each tetrahedron Δ_i , let $t_{i,1}, t_{i,2}, t_{i,3}$, and $t_{i,4}$ denote the number of triangular disks of each type, let $q_{i,1}, q_{i,2}$, and $q_{i,3}$ denote the number of quadrilateral disks of each type, and let $k_{i,1}, k_{i,2}$, and $k_{i,3}$ denote the number of octagonal disks of each type in Δ_i contained in the surface S .

Then the *standard vector representation* of S , denoted by $\mathbf{v}(S)$, is the $10n$ -dimensional vector

$$\mathbf{v}(S) = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \\ t_{2,1}, t_{2,2}, t_{2,3}, t_{2,4}, q_{2,1}, q_{2,2}, q_{2,3}, k_{2,1}, k_{2,2}, k_{2,3}; \\ \dots, k_{n,3}).$$

Lemma 3.4. *Let \mathcal{T} be a compact 3-manifold triangulation, and let S and S' be octagonal almost normal surfaces in \mathcal{T} . Then the standard vector representations $\mathbf{v}(S)$ and $\mathbf{v}(S')$ are equal if and only if the surfaces S and S' are normal isotopic (i.e., they are the “same” almost normal surface).*

This result is the almost normal counterpart to Lemma 2.2. The proof is the same, and so we do not present the details here. The key observation is that given some number of triangles, quadrilaterals, and/or octagons of various types in a single tetrahedron, if these disks can be packed into the tetrahedron disjointly, then this packing is unique up to normal isotopy.

This brings us to another reason why octagonal almost normal surfaces are simpler to deal with than general almost normal surfaces. In the general case, this packing




FIGURE 6. The three different octagon types within a tetrahedron.

need not be unique. In particular, a tube that joins two normal disks of the same type can be interchanged with some other normal disk of the same type without creating intersections (see Figure 7 for an illustration). Because of this, the extension of Lemma 3.4 to *general* almost normal surfaces fails to hold.

To determine precisely which vectors in \mathbb{R}^{10n} represent octagonal almost normal surfaces, we develop a set of matching equations and quadrilateral–octagon constraints in a similar fashion to Definitions 2.3 and 2.5.

Definition 3.5. (Standard almost normal matching equations.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let $\mathbf{w} \in \mathbb{R}^{10n}$ be any $10n$ -dimensional vector, whose coordinates we label

$$\mathbf{w} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}).$$

For each nonboundary face of \mathcal{T} and each of the three edges surrounding it, we obtain a *standard almost normal matching equation* on \mathbf{w} as follows.

Let F be some nonboundary face of \mathcal{T} , and let e be one of the three edges surrounding F . Suppose that Δ_i and Δ_j are the two tetrahedra on either side of F . Precisely one triangular disk type, one quadrilateral disk type, and two octagonal disk types in each of Δ_i and Δ_j meet F in an arc parallel to e . Suppose these correspond to coordinates $t_{i,a}, q_{i,b}, k_{i,c}$, and $k_{i,d}$ for Δ_i and $t_{j,e}, q_{j,f}, k_{j,g}$,

and $k_{j,h}$ for Δ_j . Then we obtain the matching equation

$$t_{i,a} + q_{i,b} + k_{i,c} + k_{i,d} = t_{j,e} + q_{j,f} + k_{j,g} + k_{j,h}. \quad (3-1)$$

These matching equations are the obvious extension to the original standard matching equations of Definition 2.3—we ensure that all of the disks on one side of a nonboundary face can be joined to all of the disks on the other side. As with normal surfaces, if \mathcal{T} is a closed triangulation, then there are precisely $6n$ standard almost normal matching equations (three for each of the $2n$ faces of \mathcal{T}).

Definition 3.6. (Quadrilateral–octagon constraints.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and consider any vector

$$\mathbf{w} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{10n}.$$

We say that \mathbf{w} satisfies the *quadrilateral–octagon constraints* if:

- (i) For every tetrahedron Δ_i , at most one of the six quadrilateral and octagonal coordinates $q_{i,1}, q_{i,2}, q_{i,3}, k_{i,1}, k_{i,2}$, and $k_{i,3}$ is nonzero;
- (ii) in the entire triangulation \mathcal{T} , at most one of the $3n$ octagonal coordinates $k_{1,1}, \dots, k_{n,3}$ is nonzero.

Like the quadrilateral constraints of Definition 2.5, condition (i) of the quadrilateral–octagon constraints ensures that the disks within a single tetrahedron can be embedded without intersecting. Condition (ii) ensures that we have at most one octagonal *type* within a triangulation. Although this condition is not strong enough to ensure at most one octagonal *disk*, it does have the useful property of invariance under scalar multiplication.

Note that a vector can still satisfy the quadrilateral–octagon constraints even if all its octagonal coordinates




FIGURE 7. Packing a triangle and a tube into a tetrahedron in two distinct ways.

are zero. This is necessary for the vertex enumeration algorithms to function properly; we return to this issue in Section 6.

We can now give a full set of necessary and sufficient conditions for a vector in \mathbb{R}^{10n} to represent an octagonal almost normal surface.

Theorem 3.7. *Let \mathcal{T} be a compact 3-manifold triangulation formed from n tetrahedra. An integer vector $\mathbf{w} \in \mathbb{R}^{10n}$ is the standard vector representation of an octagonal almost normal surface in \mathcal{T} if and only if:*

- *the coordinates of \mathbf{w} are all nonnegative;*
- *\mathbf{w} satisfies the standard almost normal matching equations for \mathcal{T} ;*
- *\mathbf{w} satisfies the quadrilateral–octagon constraints for \mathcal{T} ;*
- *there is precisely one nonzero octagonal coordinate in \mathbf{w} , and this coordinate is set to one.*

Once again, such a vector is called an admissible vector.

Again the proof is essentially the same as for the corresponding theorem in normal surface theory (Theorem 2.6), and so we do not reiterate the details here. The only difference is that we now have a global condition in the quadrilateral–octagon constraints (at most one nonzero octagonal coordinate in the entire triangulation), as well as an extra constraint for admissibility (precisely one nonzero octagonal coordinate with value one). These are to satisfy Definition 3.1, which requires an almost normal surface to have precisely one exceptional piece.

It is occasionally useful to consider surfaces with any number of octagonal disks, though still at most one octagonal disk *type*. In this case the vector representation, matching equations, and quadrilateral–octagon constraints all remain the same; the only change appears in Theorem 3.7, where we remove the final condition (the one that requires a unique nonzero octagonal coordinate with a value of one).

We finish by defining a vertex almost normal surface in a similar fashion to Definition 2.8. We are careful here to specify our coordinate system—in Section 4 we define a similar concept in quadrilateral–octagon coordinates, and (as with normal surfaces) a vertex surface in one coordinate system need not be a vertex surface in another.

Definition 3.8. (Standard vertex almost normal surface.)

Let \mathcal{T} be a compact 3-manifold triangulation formed from

n tetrahedra. The *standard almost normal projective solution space* for \mathcal{T} is the rational polytope formed by (i) taking the polyhedral cone of all nonnegative vectors in \mathbb{R}^{10n} that satisfy the standard almost normal matching equations for \mathcal{T} , and then (ii) intersecting this cone with the hyperplane $\{\mathbf{w} \in \mathbb{R}^{10n} \mid \sum w_i = 1\}$.

Let S be an octagonal almost normal surface in \mathcal{T} . If the standard vector representation $\mathbf{v}(S)$ is a positive multiple of some vertex of the standard almost normal projective solution space, then we call S a *standard vertex almost normal surface*.

As with normal surfaces, we can use the enumeration of vertex almost normal surfaces as a basis for high-level topological algorithms. The streamlined 3-sphere recognition of Jaco, Rubinstein, and Thompson does just this—given a “sufficiently nice” 3-manifold triangulation \mathcal{T} , we (i) enumerate all standard vertex almost normal surfaces within \mathcal{T} , and then (ii) search among these vertex surfaces for an almost normal 2-sphere [Jaco and Rubinstein 03]. We return to this algorithm in detail in Section 5.

This suggests yet another reason to prefer octagonal almost normal surfaces over general almost normal surfaces. Whereas octagonal almost normal surfaces have $10n$ -dimensional vector representations, in the general case we would need $35n$ dimensions (allowing for 25 types of tube in addition to the ten octagons, quadrilaterals, and triangles in each tetrahedron). Since both the running time and memory use for vertex enumeration can grow exponentially in the dimension of the underlying vector space [Burton 10a], increasing this dimension from $10n$ to $35n$ could well have a crippling effect on performance.³

4. QUADRILATERAL–OCTAGON COORDINATES

At this stage we are ready to develop quadrilateral–octagon coordinates, which form the main focus of this paper. Quadrilateral–octagon coordinates act as an almost normal analogy to Tollefson’s quadrilateral coordinates, in that we “forget” all information regarding triangular disks. As with quadrilateral coordinates, we happily find that—except for vertex-linking components—all of the forgotten information can be successfully recovered.

³We can avoid a $35n$ -dimensional vertex enumeration by exploiting the fact that every tube corresponds to a pair of normal disks. However, the enumeration algorithm becomes significantly more complex as a result.

The main results of this section are (i) to show that vectors in quadrilateral–octagon coordinates uniquely identify surfaces up to vertex-linking components (Lemma 4.2), and (ii) to develop a set of necessary and sufficient conditions for a vector in quadrilateral–octagon coordinates to represent an octagonal almost normal surface (Theorem 4.5). Although these mirror Tollefson’s original results in quadrilateral coordinates, the proofs follow a different course—in this sense the author hopes that this paper and [Tollefson 98] make complementary reading.

Definition 4.1. (Quadrilateral–octagon vector representation.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let S be an octagonal almost normal surface in \mathcal{T} . For each tetrahedron Δ_i , let $q_{i,1}$, $q_{i,2}$, and $q_{i,3}$ denote the number of quadrilateral disks of each type, and let $k_{i,1}$, $k_{i,2}$, and $k_{i,3}$ denote the number of octagonal disks of each type in Δ_i contained in the surface S .

Then the *quadrilateral–octagon vector representation* of S , denoted by $\mathbf{k}(S)$, is the $6n$ -dimensional vector

$$\mathbf{k}(S) = (q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; q_{2,1}, q_{2,2}, q_{2,3}, k_{2,1}, k_{2,2}, k_{2,3}; \dots, k_{n,3}).$$

Our first result in quadrilateral–octagon coordinates is a uniqueness lemma, analogous to Lemma 2.2 for normal surfaces and Lemma 3.4 for standard almost normal coordinates.

Lemma 4.2. *Let \mathcal{T} be a compact 3-manifold triangulation, and let S and S' be octagonal almost normal surfaces in \mathcal{T} . Then the quadrilateral–octagon vector representations $\mathbf{k}(S)$ and $\mathbf{k}(S')$ are equal if and only if either (i) the surfaces S and S' are normal isotopic or (ii) S and S' can be made normal isotopic by adding or removing vertex-linking components.*

Proof. The “if” direction is straightforward. If S and S' are normal isotopic or can be made so by adding or removing vertex-linking components, it follows from Lemma 3.4 that their *standard* vector representations $\mathbf{v}(S)$ and $\mathbf{v}(S')$ differ only in their triangular coordinates (since vertex links consist entirely of triangular disks). Therefore the quadrilateral and octagonal coordinates are identical in both $\mathbf{v}(S)$ and $\mathbf{v}(S')$, and we have $\mathbf{k}(S) = \mathbf{k}(S')$.

For the “only if” direction, suppose that $\mathbf{k}(S) = \mathbf{k}(S')$. Let $\mathbf{d} = \mathbf{v}(S) - \mathbf{v}(S')$ in standard almost normal coordi-




FIGURE 8. Adjacent triangles surrounding a common vertex.

nates; it follows then that

$$\mathbf{d} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, 0, 0, 0, 0, 0, 0; t_{2,1}, t_{2,2}, t_{2,3}, t_{2,4}, 0, 0, 0, 0, 0, 0; \dots) \in \mathbb{R}^{10n}$$

for some set of triangular coordinates $\{t_{i,j}\}$. In other words, all of the quadrilateral and octagonal coordinates of \mathbf{d} are zero.

We know from Theorem 2.6 that $\mathbf{v}(S)$ and $\mathbf{v}(S')$ both satisfy the standard almost normal matching equations, and because these equations are linear it follows that \mathbf{d} satisfies them also. However, with the quadrilateral and octagonal coordinates of \mathbf{d} equal to zero, we find that each matching equation (3–1) reduces to the form $t_{i,a} = t_{j,e}$, where $t_{i,a}$ and $t_{j,e}$ represent triangular disk types surrounding a common vertex of the triangulation in adjacent tetrahedra (illustrated in Figure 8).

By following these matching equations around each vertex of the triangulation \mathcal{T} , we find that for each vertex V of \mathcal{T} , the coordinates $\{t_{i,j}\}$ for all triangular disk types surrounding V are equal. That is, $\mathbf{d} = \mathbf{v}(S) - \mathbf{v}(S')$ is a linear combination of standard almost normal vector representations of vertex links. It follows then from Theorem 3.7 that the surfaces S and S' can be made normal isotopic only by adding or removing vertex-linking components.⁴ \square

Following the pattern established in previous sections, we now turn our attention to building a set of necessary and sufficient conditions for a $6n$ -dimensional vector to represent an almost normal surface in quadrilateral–octagon coordinates. These conditions include a set of matching equations modeled on the original quadrilateral matching equations of Tollefson (Definition 4.3) and a recasting of the quadrilateral–octagon constraints

⁴It is important to realize that we *can* in fact add vertex-linking components to an arbitrary surface without causing intersections. This is possible because we can “shrink” a vertex link arbitrarily close to the vertex that it surrounds, allowing us to avoid any other normal or almost normal disks.




FIGURE 9. Building the quadrilateral–octagon matching equations.

in $6n$ dimensions (Definition 4.4). The full set of necessary and sufficient conditions is laid down and proven in Theorem 4.5.

Definition 4.3. (Quadrilateral–octagon matching equations.) Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let $\mathbf{w} \in \mathbb{R}^{6n}$ be any $6n$ -dimensional vector whose coordinates we label

$$\mathbf{w} = (q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}).$$

For each nonboundary edge of \mathcal{T} , we obtain a *quadrilateral–octagon matching equation* on \mathbf{w} as follows.

Let e be some nonboundary edge of \mathcal{T} . As with Tollefson’s original quadrilateral matching equations, we arbitrarily label the two ends of e *upper* and *lower*. The tetrahedra containing edge e are arranged in a cycle around e , as illustrated in the left-hand diagram of Figure 9. Choose some arbitrary direction around this cycle, and suppose that the tetrahedra that we encounter as we travel in this direction around the cycle are labeled $\Delta_{i_1}, \dots, \Delta_{i_t}$.

Consider any tetrahedron Δ_{i_j} in this cycle. Within this tetrahedron, there are two quadrilateral types and two octagon types that meet edge e precisely once. For one quadrilateral and one octagon type, the intersection with e acts as a “hinge” about which two adjacent edges of the disk rise from the lower end of e to the upper end of e as we travel around the cycle in the chosen direction. We call these disk types the *upward quadrilateral* and the

upward octagon in Δ_{i_j} , and we call the remaining two disk types the *downward quadrilateral* and the *downward octagon* in Δ_{i_j} . All four disk types are illustrated in the right-hand portion of Figure 9.

Suppose now that the coordinates corresponding to the upward quadrilateral and octagon types are $q_{i_1, u_1}, q_{i_2, u_2}, \dots, q_{i_t, u_t}$ and $k_{i_1, u'_1}, k_{i_2, u'_2}, \dots, k_{i_t, u'_t}$ respectively, and that the coordinates corresponding to the downward quadrilateral and octagon types are $q_{i_1, d_1}, q_{i_2, d_2}, \dots, q_{i_t, d_t}$ and $k_{i_1, d'_1}, k_{i_2, d'_2}, \dots, k_{i_t, d'_t}$, respectively.⁵ Then we obtain the matching equation

$$\begin{aligned} q_{i_1, u_1} + \dots + q_{i_t, u_t} + k_{i_1, u'_1} + \dots + k_{i_t, u'_t} \\ = q_{i_1, d_1} + \dots + q_{i_t, d_t} + k_{i_1, d'_1} + \dots + k_{i_t, d'_t}. \end{aligned} \quad (4-1)$$

In other words, the total number of upward quadrilaterals and octagons surrounding e equals the total number of downward quadrilaterals and octagons surrounding e .

Note that each tetrahedron surrounding e contains a third quadrilateral type and a third octagon type, neither of which appears in equation (4-1). The third quadrilateral type is missing because it does not intersect with the edge e at all. The third octagon type is missing because although it intersects e twice, these intersections behave in a similar fashion to two triangular disks (one

⁵If we number the quadrilateral and octagon types within each tetrahedron in a natural way, we find that $u'_j = d_j$ and $d'_j = u_j$ for each j . That is, our numbering scheme associates each upward quadrilateral type with a downward octagon type and vice versa. We return to this matter in Section 8.

at each end of e). Details can be found in the proof of Theorem 4.5.

As with Tollefson’s original quadrilateral matching equations, if our triangulation \mathcal{T} is closed and has precisely v vertices, then we obtain a total of $n + v$ quadrilateral–octagon matching equations (one for each of the $n + v$ edges of \mathcal{T}).

Definition 4.4. (Quadrilateral–octagon constraints.)

Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and consider any vector

$$\mathbf{w} = (q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{6n}.$$

We say that \mathbf{w} satisfies the *quadrilateral–octagon constraints* if:

- (i) for every tetrahedron Δ_i , at most one of the six quadrilateral and octagonal coordinates $q_{i,1}, q_{i,2}, q_{i,3}, k_{i,1}, k_{i,2},$ and $k_{i,3}$ is nonzero;
- (ii) in the entire triangulation \mathcal{T} , at most one of the $3n$ octagonal coordinates $k_{1,1}, \dots, k_{n,3}$ is nonzero.

Note that Definition 4.4 is essentially a direct copy of the quadrilateral–octagon constraints for standard almost normal coordinates (Definition 3.6), merely recast in $6n$ dimensions instead of $10n$.

We can now describe the full set of necessary and sufficient conditions for a vector to represent an almost normal surface in quadrilateral–octagon coordinates. The resulting theorem incorporates aspects of both Theorem 2.6 (which uses Tollefson’s original quadrilateral matching equations) and Theorem 3.7 (which introduces the quadrilateral–octagon constraints).

Theorem 4.5. *Let \mathcal{T} be a compact 3-manifold triangulation formed from n tetrahedra. An integer vector $\mathbf{w} \in \mathbb{R}^{6n}$ is the quadrilateral–octagon vector representation of an octagonal almost normal surface in \mathcal{T} if and only if:*

- the coordinates of \mathbf{w} are all nonnegative;
- \mathbf{w} satisfies the quadrilateral–octagon matching equations for \mathcal{T} ;
- \mathbf{w} satisfies the quadrilateral–octagon constraints for \mathcal{T} ;
- there is precisely one nonzero octagonal coordinate in \mathbf{w} , and this coordinate is set to one.




FIGURE 10. Triangles and octagons for the coordinates t_{i_j, a_j} and k_{i_j, b_j} .

Yet again, such a vector is called an *admissible vector*.

Proof. We begin by showing that the four conditions listed in Theorem 4.5 are necessary. Let S be some octagonal almost normal surface in \mathcal{T} . It is clear from Theorem 3.7 that the quadrilateral–octagon vector representation $\mathbf{k}(S)$ is a nonnegative vector that satisfies the quadrilateral–octagon constraints and that there is precisely one nonzero octagonal coordinate in $\mathbf{k}(S)$ whose value is set to one. All that remains then is to show that $\mathbf{k}(S)$ satisfies the quadrilateral–octagon matching equations, which is a simple matter of combining the *standard* almost normal matching equations appropriately. The details are as follows.

Suppose that S has standard vector representation

$$\mathbf{v}(S) = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}).$$

Let e be any nonboundary edge of \mathcal{T} , and arbitrarily label the two ends of e *upper* and *lower*. Following Definition 4.3, let the tetrahedra containing e be labeled $\Delta_{i_1}, \dots, \Delta_{i_t}$ as we cycle in some arbitrary direction around e , let coordinates $q_{i_1, u_1}, q_{i_2, u_2}, \dots, q_{i_t, u_t}$ and $k_{i_1, u'_1}, k_{i_2, u'_2}, \dots, k_{i_t, u'_t}$ correspond to the upward quadrilateral and octagon types, and let coordinates $q_{i_1, d_1}, q_{i_2, d_2}, \dots, q_{i_t, d_t}$ and $k_{i_1, d'_1}, k_{i_2, d'_2}, \dots, k_{i_t, d'_t}$ correspond to the downward quadrilateral and octagon types.

We continue labeling coordinates as follows. Suppose that $t_{i_1, a_1}, t_{i_1, a_2}, \dots, t_{i_1, a_t}$ correspond to the triangular disk types surrounding the upper end of e , as illustrated in the left-hand portion of Figure 10. Furthermore, suppose that $k_{i_1, b_1}, k_{i_2, b_2}, \dots, k_{i_t, b_t}$ correspond to the octagonal disk types in each tetrahedron that are neither upward nor downward octagons, as illustrated in the right-hand portion of Figure 10.

Calling on Theorem 3.7 again, we know that $\mathbf{v}(S)$ satisfies the standard almost normal matching equations (Definition 3.5). Among those matching equations that




FIGURE 11. Writing the coordinates $t_{i,j}$ on the triangulated vertex links.

involve the adjacent pairs of tetrahedra

$$(\Delta_{i_1}, \Delta_{i_2}), (\Delta_{i_2}, \Delta_{i_3}), \dots, (\Delta_{i_t}, \Delta_{i_1}),$$

we find the t equations

$$\begin{aligned} t_{i_1, a_1} + q_{i_1, u_1} + k_{i_1, u'_1} + k_{i_1, b_1} \\ &= t_{i_2, a_2} + q_{i_2, d_2} + k_{i_2, d'_2} + k_{i_2, b_2}, \\ t_{i_2, a_2} + q_{i_2, u_2} + k_{i_2, u'_2} + k_{i_2, b_2} \\ &= t_{i_3, a_3} + q_{i_3, d_3} + k_{i_3, d'_3} + k_{i_3, b_3}, \\ &\dots \\ t_{i_t, a_t} + q_{i_t, u_t} + k_{i_t, u'_t} + k_{i_t, b_t} \\ &= t_{i_1, a_1} + q_{i_1, d_1} + k_{i_1, d'_1} + k_{i_1, b_1}. \end{aligned} \tag{4-2}$$

Summing these together and canceling the common terms $\{t_{i_j, a_j}\}$ and $\{k_{i_j, b_j}\}$, we obtain

$$\begin{aligned} q_{i_1, u_1} + \dots + q_{i_t, u_t} + k_{i_1, u'_1} + \dots + k_{i_t, u'_t} \\ = q_{i_1, d_1} + \dots + q_{i_t, d_t} + k_{i_1, d'_1} + \dots + k_{i_t, d'_t}. \end{aligned}$$

That is, the quadrilateral–octagon vector representation $\mathbf{k}(S)$ satisfies the quadrilateral–octagon matching equations.

We now turn to the more interesting task of proving that our list of conditions is *sufficient* for an integer vector $\mathbf{w} \in \mathbb{R}^{6n}$ to represent an octagonal almost normal surface. Let

$$\mathbf{w} = (q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{6n}$$

be an arbitrary integer vector that satisfies the four conditions listed in the statement of this theorem. Our aim is to extend \mathbf{w} to an integer vector

$$\mathbf{w}' = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{10n}$$

that satisfies the conditions of Theorem 3.7. If we can do this, it will follow from Theorem 3.7 that \mathbf{w}' is the standard almost normal vector representation of some octagonal almost normal surface in \mathcal{T} , whereupon \mathbf{w} must be


the quadrilateral–octagon vector representation of this same surface.

Given our conditions on $\mathbf{w} \in \mathbb{R}^{6n}$, it is clear that any nonnegative extension $\mathbf{w}' \in \mathbb{R}^{10n}$ will satisfy the quadrilateral–octagon constraints, and will have precisely one nonzero octagonal coordinate whose value is set to one. All we must do then is show that we can find a set of nonnegative triangular coordinates $\{t_{i,j}\}$ that satisfy the standard almost normal matching equations of Definition 3.5.

Our broad strategy is to use the vertex links of \mathcal{T} as a “canvas” on which we write the triangular coordinates $t_{i,j}$, and to reformulate the matching equations as local constraints on this canvas. In doing this, we show that the standard almost normal matching equations describe a cochain $\alpha \in C^1(\mathcal{D})$, where \mathcal{D} is the dual polygonal decomposition of the vertex links, and that a solution $\{t_{i,j}\}$ exists if and only if α is a coboundary. Using the quadrilateral–octagon matching equations, we then find that α is a cocycle, whereupon the result follows from the trivial homology of the vertex links. The details are as follows.


Because \mathcal{T} is a compact triangulation, each of its vertex links is a triangulated sphere or disk, as illustrated in the left-hand diagram of Figure 11. Each triangular disk type appears once and only once among the vertex links, and so we can write each integer $t_{i,j}$ on the corresponding vertex link triangle as illustrated in the right-hand diagram of Figure 11. This is the sense in which we use the vertex links as a “canvas.”

We can now reformulate the standard almost normal matching equations as constraints on this canvas. Recall that each standard matching equation involves a face F of \mathcal{T} and arcs parallel to some edge e of this face, as illustrated in the left-hand diagram of Figure 12. We can associate every such equation with a single nonboundary edge g of a triangulated vertex link, where this edge g also appears as an arc of the face F parallel to e , as illustrated in the right-hand diagram of Figure 12. In this




Standard matching equation for F and e Edge g of a vertex link

FIGURE 12. Associating a standard matching equation with an edge of a vertex link.



Equation: $t_{i,a} - t_{j,e} = 3$

FIGURE 13. Representing a standard matching equation by a labeled arrow.



Labels on arrows

Coordinates $t_{i,j}$

FIGURE 14. Solving the standard matching equations.

way, the standard almost normal matching equations and the nonboundary edges of the triangulated vertex links are in one-to-one correspondence.

Now consider some standard matching equation $t_{i,a} + q_{i,b} + k_{i,c} + k_{i,d} = t_{j,e} + q_{j,f} + k_{j,g} + k_{j,h}$ (as seen in Definition 3.5), and let g be the corresponding edge of the triangulated vertex links. The coordinates $t_{i,a}$ and $t_{j,e}$ correspond to the triangles on either side of g , and so we can write this equation in the form

$$t_{i,a} - t_{j,e} = K,$$

where K depends only on the quadrilateral and octagonal coordinates of \mathbf{w} . In other words, K is a fixed quantity (dependent on the chosen edge g) that we can evaluate by looking at our original vector $\mathbf{w} \in \mathbb{R}^{6n}$. We express this equation on our canvas by drawing an arrow from the triangle containing $t_{j,e}$ to the triangle containing $t_{i,a}$,

and by labeling this arrow with the constant K . This procedure is illustrated in Figure 13.

Our situation is now as follows. On our canvas—the triangulated vertex links of \mathcal{T} —we have a labeled arrow crossing each nonboundary edge, and our task is to fill each triangle with an integer such that the difference across each edge matches the label on the corresponding arrow. An example of such a solution for a triangulated disk is illustrated in Figure 14. It is clear at this point that we do not need to worry about our nonnegativity condition, since we can always add a constant to every triangle without changing the differences across the edges.

We can rephrase this using the language of cohomology. Let \mathcal{D} be the dual polygonal decomposition of the set of all vertex links, so that each triangle of a vertex link becomes a vertex of \mathcal{D} and each labeled arrow




FIGURE 15. The triangles surrounding some vertex U of the vertex link $\ell(V)$.

becomes a directed edge of \mathcal{D} . Then together the arrows describe a cochain $\alpha \in C^1(\mathcal{D})$ that maps each dual edge to the corresponding label. A solution $\{t_{i,j}\}$ corresponds to a cochain $\beta \in C^0(\mathcal{D})$ that maps each dual vertex to the integer in the corresponding triangle, and the “difference condition” that such a solution must satisfy is simply $\alpha = \delta\beta$. That is, a solution $\{t_{i,j}\}$ exists if and only if α is a coboundary.

We now turn to the quadrilateral–octagon matching equations, which we assume hold for our original vector $\mathbf{w} \in \mathbb{R}^{6n}$. These equations do not involve the triangular coordinates at all. Instead they tell us about the relations between different quadrilateral and octagonal coordinates of \mathbf{w} , which means they give us information about the labels on our arrows.

Consider some vertex V of the triangulation \mathcal{T} , let U be some nonboundary vertex of the triangulated link $\ell(V)$, and let e be the edge of \mathcal{T} that runs through U and V as illustrated in Figure 15. Let K_1, \dots, K_t be the labels on the arrows surrounding U , as seen in the right-hand diagram of this figure (where we make all arrows point in the same direction around U by reversing arrows and negating labels as necessary). Recall that by construction, each label K_i is a linear combination of two quadrilateral and four octagonal coordinates of \mathbf{w} .

Now consider the quadrilateral–octagon matching equation constructed from edge e . By declaring V to be at the upper end of e , we can invert the procedure used earlier in equation (4-2) to express our matching equation as

$$K_1 + \dots + K_t = 0.$$

In other words, the quadrilateral–octagon matching equations tell us that *around every nonboundary vertex of a triangulated vertex link, the sum of labels on arrows is zero*. We see this for instance in Figure 14: by walking clockwise around each internal vertex and negating labels when arrows point backward, the left internal ver-

tex gives $1 + 0 - (-3) - 4 + 3 + (-1) - 2 = 0$, and the right internal vertex gives $2 + (-1) - 0 - (-2) + (-3) = 0$.

Returning to our cohomology formulation, this simply tells us that $\delta\alpha = 0$, where $\alpha \in C^1(\mathcal{D})$ is the cochain described earlier. That is, α is a cocycle. However, because each vertex link is a sphere or a disk, the cohomology group $H^1(\mathcal{D})$ is trivial. Therefore α is also a coboundary, as required. \square

The final step of this proof shows why we must exclude the ideal triangulations of [Thurston 78] from our consideration. In an ideal triangulation, vertices form higher-genus cusps, whereupon the vertex links become higher-genus surfaces with nontrivial homology. Therefore, although the quadrilateral–octagon matching equations still show that $\alpha \in C^1(\mathcal{D})$ is a cocycle in the proof above, we can no longer conclude from this that α is a coboundary and that the solution $\{t_{i,j}\}$ exists.

To finish this section, we define a vertex surface in our new coordinate system using the same pattern that we have employed several times already.

Definition 4.6. (Quadrilateral–octagon vertex almost normal surface.) Let \mathcal{T} be a compact 3-manifold triangulation. The quadrilateral–octagon projective solution space for \mathcal{T} is the rational polytope formed by (i) taking the polyhedral cone of all nonnegative vectors in \mathbb{R}^{6n} that satisfy the quadrilateral–octagon matching equations for \mathcal{T} , and then (ii) intersecting this cone with the hyperplane $\{\mathbf{w} \in \mathbb{R}^{6n} \mid \sum w_i = 1\}$.

Let S be an octagonal almost normal surface in \mathcal{T} . If the quadrilateral–octagon vector representation $\mathbf{k}(S)$ is a positive multiple of some vertex of the quadrilateral–octagon projective solution space, then we call S a quadrilateral–octagon vertex almost normal surface.

It should be noted that while it can be shown that a connected quadrilateral–octagon vertex almost normal surface is also a standard vertex almost normal surface,⁶ the converse is not necessarily true. We address this problem for the 3-sphere recognition algorithm in the following section by proving that the surface we seek does indeed appear as a vertex surface in quadrilateral–octagon coordinates. More generally, we describe in Section 6 how the conversion algorithm of [Burton 09b] can reconstruct the set of all standard vertex almost normal surfaces, given the set of all quadrilateral–octagon vertex almost normal surfaces as input.

⁶The proof is identical to the corresponding result for normal surfaces; see [Burton 09b] for details.

5. 3-SPHERE RECOGNITION

The algorithm to recognize the 3-sphere has seen significant evolution since it was first introduced by Rubinstein in 1992. Rubinstein’s original algorithm involved finding a maximal disjoint collection of embedded normal 2-spheres within a triangulation \mathcal{T} , slicing \mathcal{T} open along these 2-spheres, and then searching for almost normal 2-spheres within the complementary regions [Rubinstein 97]. Thompson gave an alternative proof of this algorithm using Gabai’s concept of *thin position*, and also showed that it was necessary to consider only octagonal almost normal surfaces [Thompson 94].

The algorithm at this stage remained extremely slow⁷ and fiendishly difficult to implement. The main problems were (i) the need to locate and deal with many normal and almost normal surfaces simultaneously and (ii) the need to locate almost normal surfaces in complementary regions of \mathcal{T} containing not only tetrahedra but also sliced and truncated *pieces* of tetrahedra. Fortunately, this algorithm was simplified enormously by Jaco and Rubinstein using the concept of *0-efficient triangulations* [Jaco and Rubinstein 03], to the point where a computer implementation became practicable. The first real implementation of 3-sphere recognition was in the software package Regina in 2004 [Burton 04], over a decade after the algorithm was first introduced.

We begin this section with a brief discussion of the theory behind the final algorithm of [Jaco and Rubinstein 03], followed by the algorithm itself (Algorithm 5.4). A key step of this algorithm (and indeed its bottleneck) is an enumeration of standard vertex almost normal surfaces. The main result of this section is Theorem 5.6, in which we show that we can restrict our attention to *quadrilateral–octagon* vertex normal surfaces instead.

As noted in the introduction, the enumeration of normal and almost normal surfaces can grow to become exponentially slow in the dimension of the underlying vector space [Burton 10a]. Using Theorem 5.6 we are able to reduce this dimension from $10n$ to $6n$, which in theory should cut down the running time substantially. In Section 7 we test this experimentally, and indeed we find that the speed of 3-sphere recognition improves by orders of magnitude for the cases that we examine.

We turn our attention now to the most recent form of the 3-sphere recognition algorithm, as given in [Jaco and Rubinstein 03]. The advantages of this algorithm over

its predecessors are due to the use of 0-efficient triangulations, which are defined as follows.

Definition 5.1. (0-Efficiency.) Let \mathcal{T} be a closed compact 3-manifold triangulation. We say that \mathcal{T} is *0-efficient* if the only embedded normal 2-spheres in \mathcal{T} are vertex links.

It turns out that 0-efficient triangulations are relatively common, in that they exist for all closed orientable irreducible 3-manifolds except for $\mathbb{R}P^3$ [Jaco and Rubinstein 03, Theorem 5.5]. Moreover, Jaco and Rubinstein provide a procedure for explicitly constructing a 0-efficient triangulation of such a manifold. More generally, Jaco and Rubinstein prove the following:

Theorem 5.2. *Let \mathcal{T} be a closed compact 3-manifold triangulation representing some (unknown) orientable 3-manifold M . Then there is a procedure to express M as a connected sum $M = M_1 \# \cdots \# M_t$, where each M_i is either given by a 0-efficient triangulation \mathcal{T}_i or is one of the special spaces $S^2 \times S^1$, $\mathbb{R}P^3$, and the lens space $L(3, 1)$.*

The details of this procedure can be found in [Jaco and Rubinstein 03, Theorems 5.9 and 5.10] and surrounding comments. The key idea is to repeatedly locate embedded normal 2-spheres and crush them until no such 2-spheres can be found. Note that we might still be unable to identify the constituent manifolds $\{M_i\}$, but with the 0-efficient triangulations $\{\mathcal{T}_i\}$ we may be better placed to learn more about them. We do not expand further on this decomposition procedure of Jaco and Rubinstein. Although it plays a key role in the 3-sphere recognition algorithm, our focus for this paper is on a different part of the algorithm.

The core result behind Jaco and Rubinstein’s version of the 3-sphere recognition algorithm is the following theorem, which builds on earlier work of Rubinstein and Thompson [Rubinstein 97, Thompson 94] by exploiting properties of 0-efficiency. The various components of this theorem can be found in [Jaco and Rubinstein 03, Proposition 5.12] and surrounding comments.

Theorem 5.3. *Let \mathcal{T} be a closed compact 3-manifold triangulation that is orientable and 0-efficient. Then the following statements are equivalent:*

- \mathcal{T} is a triangulation of the 3-sphere;
- \mathcal{T} has more than one vertex, or \mathcal{T} contains an octagonal almost normal 2-sphere;

⁷In theory, of course, since at that stage a computer implementation did not exist.

- \mathcal{T} has more than one vertex, or \mathcal{T} contains an octagonal almost normal 2-sphere that is a standard vertex almost normal surface.

Based on this result, the full 3-sphere recognition algorithm of Jaco and Rubinstein runs as follows.

Algorithm 5.4. (3-Sphere recognition.) Let \mathcal{T} be a closed compact 3-manifold triangulation, and let M be the 3-manifold that \mathcal{T} represents. The following algorithm decides whether M is the 3-sphere S^3 :

1. Test whether M is orientable and has trivial first homology. If not, then terminate with the result $M \neq S^3$.
2. Using the procedure of Theorem 5.2, express the underlying 3-manifold M as a connected sum decomposition $M_1 \# M_2 \# \cdots \# M_t$, where each M_i is given by a 0-efficient triangulation \mathcal{T}_i . If this list is empty (i.e., $t = 0$), then terminate with the result $M = S^3$.
3. Of the 0-efficient triangulations $\mathcal{T}_1, \dots, \mathcal{T}_t$, ignore those with more than one vertex. For each one-vertex triangulation \mathcal{T}_i :
 - (i) Enumerate the standard vertex almost normal surfaces of \mathcal{T}_i .
 - (ii) Search through the resulting list of surfaces for an almost normal 2-sphere. If one cannot be found, then terminate with the result $M \neq S^3$.
4. If we have not yet terminated, then every 0-efficient triangulation \mathcal{T}_i has either more than one vertex or an almost normal 2-sphere. In this case we conclude that $M = S^3$.

Remark 5.5. There are some points worth noting about this algorithm:

- In step 2, we do not account for the special spaces $S^2 \times S^1$, $L(3, 1)$, and $\mathbb{R}P^3$ that can arise in the decomposition procedure of Theorem 5.2. This is because the homology test in step 1 prevents any of these special spaces from appearing.
- The enumeration of surfaces in step 3 involves a modified double description method, which is described fully in [Burton 10a]. We return to the enumeration algorithm in Section 6, where we discuss it from the perspective of quadrilateral–octagon coordinates.

We come now to the main result of this section, which is a quadrilateral–octagon analogue for the earlier Theorem 5.3. What we essentially show is that for the enumeration of vertex almost normal surfaces in step 3 of the algorithm above, we can work in quadrilateral–octagon coordinates instead of standard coordinates (in other words, $6n$ dimensions instead of $10n$). This is important from a practical perspective, since experience indicates that this enumeration step is typically the bottleneck for the entire 3-sphere recognition algorithm.⁸

Theorem 5.6. *Let \mathcal{T} be a closed compact 3-manifold triangulation that is orientable and 0-efficient. Then the following statements are equivalent:*

- \mathcal{T} is a triangulation of the 3-sphere;
- \mathcal{T} has more than one vertex, or \mathcal{T} contains an octagonal almost normal 2-sphere that is a quadrilateral–octagon vertex almost normal surface.

Proof. We assume that \mathcal{T} is a one-vertex triangulation, since otherwise the result follows immediately from Theorem 5.3. Given this, it is clear from Theorem 5.3 that \mathcal{T} triangulates the 3-sphere if and only if \mathcal{T} contains an octagonal almost normal 2-sphere. All we need to show is that if \mathcal{T} contains an octagonal almost normal 2-sphere, then it contains one as a quadrilateral–octagon vertex almost normal surface.

Our proof is based on an idea of Casson, used also by Jaco and Rubinstein, to prove the corresponding claim in standard coordinates. We work within a face of the projective solution space and show that the maximum of $\chi'(\mathbf{u}) - O(\mathbf{u})$ occurs at a vertex, where $\chi'(\cdot)$ represents the Euler characteristic and $O(\cdot)$ is the sum of octagonal coordinates. One complication that we face in quadrilateral–octagon coordinates is that in contrast to the situation in standard coordinates, the Euler characteristic is not a linear functional. Nevertheless, we are able to work around this difficulty by falling back on convexity instead. The details are as follows.

Suppose that \mathcal{T} contains some octagonal almost normal 2-sphere S . Let $\mathcal{P}_{QO} \subseteq \mathbb{R}^{6n}$ denote the quadrilateral–octagon projective solution space (Definition 4.6), and let F be the minimal-dimensional face of \mathcal{P}_{QO} containing the vector representation $\mathbf{k}(S)$. This face F is the face in which we plan to work.

⁸If the manifold M is a connected sum of several high-complexity homology 3-spheres, then the decomposition procedure of Jaco and Rubinstein becomes a greater problem for performance. However, it is reasonable to suggest that such cases are rare in “ordinary” applications.

We begin by showing that every point $\mathbf{u} \in F$ satisfies the quadrilateral–octagon constraints. In contrast, suppose that some $\mathbf{u} \in F$ does not satisfy these constraints. Then for some coordinate position $i \in \{1, \dots, 6n\}$, we must have $u_i > 0$ where $\mathbf{k}(S)_i = 0$. Let H be the hyperplane $H = \{\mathbf{w} \in \mathbb{R}^{6n} \mid w_i = 0\}$; it is clear that H is a supporting hyperplane for \mathcal{P}_{QO} , and so $H \cap F$ is a subspace of F containing $\mathbf{k}(S)$ but not \mathbf{u} , contradicting the minimality of F .

In order to define the Euler characteristic function $\chi' : F \rightarrow \mathbb{R}$, we must understand the relationship between standard and quadrilateral–octagon vector representations. With this in mind, we define the *projection map* $\pi : \mathbb{R}^{10n} \rightarrow \mathbb{R}^{6n}$ and the *extension map* $\varepsilon : F \rightarrow \mathbb{R}^{10n}$ as follows:⁹

- (i) For a vector $\mathbf{v} \in \mathbb{R}^{10n}$, the projection $\pi(\mathbf{v})$ is the vector \mathbf{v} with all triangular coordinates removed. That is, if

$$\mathbf{v} = (t_{1,1}, t_{1,2}, t_{1,3}, t_{1,4}, q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{10n},$$

then

$$\pi(\mathbf{v}) = (q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{6n}.$$

- (ii) For a vector $\mathbf{u} \in F \subset \mathbb{R}^{6n}$, the extension $\varepsilon(\mathbf{u})$ is defined as follows. Because $F \subseteq \mathcal{P}_{QO}$, we know that \mathbf{u} satisfies the quadrilateral–octagon matching equations. By the same argument used in the proof of Theorem 4.5, we can therefore solve the *standard* almost normal matching equations to obtain values for the missing triangular coordinates, giving us an extension $\mathbf{x} \in \mathbb{R}^{10n}$ that satisfies the standard almost normal matching equations and for which $\pi(\mathbf{x}) = \mathbf{u}$.

By the same argument used in the proof of Lemma 4.2, this extension is unique up to multiples of vertex links. We therefore define $\varepsilon(\mathbf{u})$ to be the “minimal” extension, in the sense that we subtract the largest possible multiple of each vertex link without allowing any coordinates to become negative. In other words, every coordinate of $\varepsilon(\mathbf{u})$ is nonnegative, and for every vertex link $\ell(V)$, the coordinate for some triangular disk type in $\ell(V)$ is zero.

⁹These maps are the almost normal analogues of *quadrilateral projection* and *canonical extension*, which are defined in [Burton 09b] for the context of embedded normal surfaces.

It is important to note that, based on the way in which we solve the standard almost normal matching equations, if \mathbf{u} is an integer vector then $\varepsilon(\mathbf{u})$ is an integer vector also.

It is clear that $\pi : \mathbb{R}^{10n} \rightarrow \mathbb{R}^{6n}$ is a linear map. For $\varepsilon : F \rightarrow \mathbb{R}^{10n}$ the situation is a little more complex. By the linearity of the matching equations, it is clear that

$$\varepsilon(\lambda\mathbf{u}) = \lambda\varepsilon(\mathbf{u}) \quad (5-1)$$

for any $\lambda \geq 0$. On the other hand, for arbitrary $\mathbf{u}, \mathbf{w} \in F$, we know only that $\varepsilon(\mathbf{u} + \mathbf{w})$ and $\varepsilon(\mathbf{u}) + \varepsilon(\mathbf{w})$ are related by adding or subtracting multiples of vertex links. Since both $\varepsilon(\mathbf{u})$ and $\varepsilon(\mathbf{w})$ are nonnegative vectors, $\varepsilon(\mathbf{u} + \mathbf{w})$ can only *subtract* vertex links from their sum, yielding the nonlinear relation

$$\varepsilon(\mathbf{u} + \mathbf{w}) = \varepsilon(\mathbf{u}) + \varepsilon(\mathbf{w}) - \sum \lambda_i \mathbf{v}(\ell(V_i)), \quad (5-2)$$

where each $\ell(V_i)$ is a vertex-linking surface and each λ_i is nonnegative.

We can now define our Euler characteristic function as follows. It is well known that the Euler characteristic is a linear functional in *standard* coordinates: for an almost normal surface S the Euler characteristic $\chi(S)$ is a linear function of the coordinates $\{t_{i,j}\}$, $\{q_{i,j}\}$, and $\{k_{i,j}\}$,¹⁰ and we simply extend this to a linear functional $\chi : \mathbb{R}^{10n} \rightarrow \mathbb{R}$. On our face $F \subseteq \mathcal{P}_{QO}$ we then define the Euler characteristic function $\chi' : F \rightarrow \mathbb{R}$ by

$$\chi'(\mathbf{u}) = \chi(\varepsilon(\mathbf{u})).$$

Although χ' is not linear on F , we can observe that each vertex link $\ell(V_i)$ is a 2-sphere, and so $\chi(\ell(V_i)) > 0$. Therefore, equations (5-1) and (5-2) give

$$\begin{aligned} \chi'(\lambda\mathbf{u}) &= \lambda\chi'(\mathbf{u}) \quad \text{for all } \mathbf{u} \in F \text{ and } \lambda \geq 0; \\ \chi'(\mathbf{u} + \mathbf{w}) &\leq \chi'(\mathbf{u}) + \chi'(\mathbf{w}) \quad \text{for all } \mathbf{u}, \mathbf{w} \in F. \end{aligned} \quad (5-3)$$

That is, χ' is a *convex* function on F .

We are now able to exploit an analogue of the functional that Casson uses in standard coordinates. Define the function $g : F \rightarrow \mathbb{R}$ by $g(\mathbf{u}) = \chi'(\mathbf{u}) - O(\mathbf{u})$, where $O(\mathbf{u})$ is the sum of all octagonal coordinates in \mathbf{u} . Since χ' is convex and O is clearly linear, it follows that g is convex also. Therefore the maximum of g is achieved at a vertex of the face F . Let this vertex be $\mathbf{m} \in F$.

¹⁰The number of faces in S is simply $\sum t_{i,j} + \sum q_{i,j} + \sum k_{i,j}$. The number of vertices in S is $\sum w(e_i)$, where $w(e_i)$ is the number of times S intersects the edge e_i of \mathcal{T} , and where $w(e_i)$ can be written as a linear function of the disks in some arbitrary tetrahedron containing e_i . Edges of S are dealt with in a similar way.

Our original almost normal 2-sphere S has $g(\mathbf{k}(S)) = 1$, since S has Euler characteristic two, precisely one octagonal disk, and no vertex-linking components. Given that $\mathbf{k}(S) \in F$, it follows that $g(\mathbf{m}) > 0$ also. Using the fact that \mathcal{P}_{QO} is a rational polytope, we can define $\mathbf{m}' \in \mathbb{Z}^{6n}$ to be the smallest positive multiple of \mathbf{m} with all integer coordinates.

Given that $F \subseteq \mathcal{P}_{QO}$ and that every vector in F satisfies the quadrilateral–octagon constraints, it follows that the extension $\varepsilon(\mathbf{m}')$ satisfies all the conditions of admissibility in \mathbb{R}^{10n} except perhaps the requirement that the unique octagonal coordinate be set to one; instead we might have multiple octagonal disks, or we might have none at all. We can therefore reconstruct an embedded surface S' with standard vector representation $\mathbf{v}(S') = \varepsilon(\mathbf{m}')$, where S' is one of the following:

- an octagonal almost normal surface;
- a surface like an octagonal almost normal surface but with more than one octagonal disk;
- an embedded normal surface with no octagonal disks at all.

We can show that the surface S' is connected as follows. Suppose that S' consists of distinct components S'_1, \dots, S'_t , where $t > 1$. Then in quadrilateral–octagon coordinates we have $\mathbf{m}' = \mathbf{k}(S') = \sum \mathbf{k}(S'_i)$, and since \mathbf{m}' is the smallest integer multiple of a vertex of \mathcal{P}_{QO} , it follows that all but one of the integer vectors $\mathbf{k}(S'_1), \dots, \mathbf{k}(S'_t)$ must be zero. Therefore all but one of the components S'_i are vertex links, which is impossible because the standard vector representation $\mathbf{v}(S')$ was constructed using the extension map ε .

From equation (5–3) we have $\chi'(\mathbf{m}') - O(\mathbf{m}') > 0$, and because S' is connected it follows that $2 \geq \chi(S') > O(\mathbf{m}') \geq 0$. We must therefore be in one of the following situations:

- (i) $\chi(S') = 2$ and $O(\mathbf{m}') = 0$.

In this case S' is an embedded normal 2-sphere. Since our triangulation \mathcal{T} is 0-efficient, it follows that S' is a vertex link and therefore $\mathbf{k}(S') = \mathbf{0}$, contradicting the fact that $\mathbf{k}(S')$ is a positive multiple of the vertex $\mathbf{m} \in \mathcal{P}_{QO}$.

- (ii) $\chi(S') = 1$ and $O(\mathbf{m}') = 0$.

In this case S' is an embedded normal projective plane. Since \mathcal{T} is orientable, S' must be a one-sided surface that doubles to an embedded normal sphere, giving the same contradiction as above.

- (iii) $\chi(S') = 2$ and $O(\mathbf{m}') = 1$.

In this case S' has precisely one octagonal disk, and is therefore an octagonal almost normal 2-sphere.

The only case that does not yield a contradiction is (iii). Since $\mathbf{k}(S')$ is a positive multiple of the vertex $\mathbf{m} \in \mathcal{P}_{QO}$, it follows that S' is the quadrilateral–octagon vertex almost normal 2-sphere that we seek. □

6. ENUMERATION ALGORITHMS

In this section we examine the practical issue of enumerating vertex almost normal surfaces. We do not go into the full details of the enumeration algorithms, since they are intricate enough to form the subjects of papers themselves [Burton 09b, Burton 10a]. However, we do explain in broad terms why the algorithms used for enumerating *normal* surfaces can also be used to enumerate *almost normal* surfaces in both standard and quadrilateral–octagon coordinates, with no unexpected changes.

The layout of this section is as follows. We begin in Section 6.1 with the direct enumeration algorithm, which is based on a filtered double description method. In Section 6.2 we discuss the conversion algorithm from quadrilateral–octagon to standard coordinates, which allows us to enumerate vertex surfaces in standard coordinates substantially faster than through a direct enumeration. We conclude in Section 6.3 with some further notes on the implementation and use of these algorithms.

The following are the key observations that we make for quadrilateral–octagon coordinates:

- (i) Enumerating vertex surfaces in quadrilateral–octagon coordinates is a simple matter of applying the direct enumeration algorithm of [Burton 10a] “out of the box,” though we cannot enforce the “one and only one octagon” constraint until the algorithm has finished.
- (ii) Likewise, we can use the conversion algorithm of [Burton 09b] out of the box to convert the vertices of the quadrilateral–octagon projective solution space into the vertices of the standard projective solution space, though again we must be careful with our use of the “one and only one octagon” constraint.
- (iii) As a consequence of (ii), we can use quadrilateral–octagon coordinates to substantially improve the speed of high-level topological algorithms, even *without* specific results such as Theorem 5.6 that allow us to focus only on quadrilateral–octagon coordinates.

6.1 Direct Enumeration

At its core, the enumeration of vertex *normal* surfaces uses a combination of the double description method of [Motzkin et al. 53] and the filtering method of Letscher. The details can be found in [Burton 10a], but essentially the algorithm runs as follows.

Suppose we are working in the vector space \mathbb{R}^d with g matching equations (so for a closed one-vertex triangulation we have $d = 7n$ and $g = 6n$ in standard coordinates, or $d = 3n$ and $g = n+1$ in quadrilateral coordinates). We inductively create a series of polytopes $P_0, \dots, P_g \subseteq \mathbb{R}^d$ described by their vertex sets V_0, \dots, V_g according to the following procedure:

- The polytope P_0 is the intersection of the nonnegative orthant in \mathbb{R}^d with the projective hyperplane $\{\mathbf{x} \in \mathbb{R}^d \mid \sum x_i = 1\}$, and the corresponding vertex set V_0 consists of all unit vectors in \mathbb{R}^d .
- The polytope P_i is created by intersecting P_{i-1} with a hyperplane corresponding to the i th matching equation. The vertex set V_i consists of vertices $\mathbf{v} \in V_{i-1}$ that lie inside this hyperplane, as well as combinations of pairs of vertices $\mathbf{u}, \mathbf{v} \in V_{i-1}$ that lie on opposite sides of this hyperplane.

The final polytope P_g is the projective solution space, and by rescaling the vertex set V_g into integer coordinates we can reconstruct the corresponding vertex normal surfaces.

Although this procedure accounts for nonnegativity and the matching equations, we have not made use of the quadrilateral constraints. This is where the filtering method of Letscher comes into play. The key idea is to enforce the quadrilateral constraints at every stage of the double description method. Specifically, we strip all vertices from each set V_i that do not satisfy the quadrilateral constraints. Although this means that each set V_i does not give a complete representation of the polytope P_i , by filtering out “bad” vertices at every stage of the algorithm we can tame the exponential explosion in the size of the vertex sets V_i , improving the performance of the algorithm in practice by a substantial amount.

It is useful to understand *why* this enumeration algorithm works, so that we can see whether it can also be used with almost normal surfaces. In essence, the key reasons are as follows:

- The double description method of Motzkin et al. works because the projective solution space is a convex polytope, defined as the intersection of the

nonnegative orthant with the projective hyperplane $\sum x_i = 1$ and an additional hyperplane for each matching equation.

- The filtering method of Letscher works because the quadrilateral constraints satisfy the following key properties:

Property A: The quadrilateral constraints are satisfied on a union of faces of the nonnegative orthant, and therefore on a union of faces of the projective solution space.

Property B: Let \mathbf{u} and \mathbf{v} be nonnegative vectors in \mathbb{R}^d . If either \mathbf{u} or \mathbf{v} does not satisfy the quadrilateral constraints, then the combination $\alpha\mathbf{u} + \beta\mathbf{v}$ can never satisfy the quadrilateral constraints for any $\alpha, \beta > 0$.

Note that property B is an immediate consequence of property A, and that property A holds because each constraint is of the form “at most one of the coordinates $\{x_i \mid i \in C\}$ may be nonzero,” where $C \subseteq \{1, \dots, d\}$ is some set of coordinate positions.

We now turn our attention to the enumeration of vertex *almost normal* surfaces, in both standard almost normal coordinates and quadrilateral–octagon coordinates.

Once again, the projective solution space is the intersection of the nonnegative orthant with the projective hyperplane $\sum x_i = 1$ and an additional hyperplane for each matching equation. As a result, the double description method of Motzkin et al. works seamlessly with almost normal surfaces.

Like the original quadrilateral constraints, the quadrilateral–octagonal constraints for almost normal surfaces are each of the form “at most one of the coordinates $\{x_i \mid i \in C\}$ may be nonzero,” where $C \subseteq \{1, \dots, d\}$ is some set of coordinate positions. As a result, both of the above properties A and B hold, and we can seamlessly use the filtering method of Letscher to enforce the quadrilateral–octagon constraints at each stage of the double description method.

However, Theorems 3.7 and 4.5 show that octagonal almost normal surfaces come with an additional constraint:

Constraint (\star): For \mathbf{v} to be the vector representation of an octagonal almost normal surface, there must be some nonzero octagonal coordinate in \mathbf{v} , and this coordinate must be set to one.

It is clear that we cannot enforce (\star) on the projective solution space, since there the coordinates of each vector are rationals (not integers) that *sum* to one. From the viewpoint of the projective solution space, this constraint is not so much a property of a vector \mathbf{v} , but rather a property of the smallest multiple of \mathbf{v} with integer coordinates. It follows that the final constraint (\star) cannot be inserted verbatim into the filtering process.

We might instead consider enforcing a weaker version of (\star) , whereby every vector $\mathbf{v} \in V_i$ must have some nonzero octagonal coordinate (therefore eliminating vectors that yield no octagons at all). However, this variant is also unsuitable for filtering, since it satisfies neither of the properties A and B. In essence, the reason we must keep track of *normal* surfaces (with no octagons) is so that we can combine them with old almost normal surfaces to create new almost normal surfaces.

The conclusion then is that we must forget the final condition (\star) while the algorithm is running, and enforce it only once we have our final set of vertices V_g . Note that this is not a severe penalty—the quadrilateral–octagon constraints already ensure that we have at most one octagon *type* in each vector, and so our only inefficiency is that we must carry around vectors that yield too many octagons of a single type, or that yield no octagons at all.

As a final note, the paper [Burton 10a] offers a number of additional optimizations to the core filtered double description method. As with the core algorithm, these optimizations can also be used seamlessly with octagonal almost normal surfaces, as long as we remember to delay the constraint (\star) until after the algorithm has finished.

6.2 The Conversion Algorithm

The paper [Burton 09b] describes a conversion algorithm from quadrilateral to standard coordinates for normal surfaces. The purpose of this algorithm is not just to convert *vectors* between coordinate systems (which is fairly straightforward), but to convert entire *solution sets*. That is, the algorithm begins with the set of all vertices of the *quadrilateral* projective solution space that satisfy the quadrilateral constraints, and converts this to the (typically much larger) set of all vertices of the *standard* projective solution space that satisfy the quadrilateral constraints. We are therefore able to recover the standard vertex normal surfaces that are “lost” in quadrilateral coordinates.

As a result, this algorithm allows us to enumerate all standard vertex normal surfaces using the following two-step procedure:

1. Use direct enumeration (as described in Section 6.1) to enumerate all vertices of the *quadrilateral* projective solution space that satisfy the quadrilateral constraints.
2. Use the conversion algorithm (as described below) to recover all vertices of the *standard* projective solution space that satisfy the quadrilateral constraints, and thereby the set of all standard vertex normal surfaces.

Experimentation shows the conversion algorithm to have negligible running time, and as a result this two-step procedure is found to be orders of magnitude faster than a direct enumeration in standard coordinates [Burton 09b].

The overall outcome is that we can harness the *speed* of quadrilateral coordinates without the need to prove additional *theorems* in quadrilateral coordinates (such as we do here for quadrilateral–octagon coordinates in Theorem 5.6).

Broadly speaking, the conversion algorithm operates as follows. Suppose the triangulation \mathcal{T} is formed from n tetrahedra, and contains the m vertices V_1, \dots, V_m . We inductively construct lists of vectors $L_0, \dots, L_m \subset \mathbb{R}^{7n}$ according to the following procedure:

- The list L_0 contains the input for the algorithm, which consists of all vertices of the quadrilateral projective solution space that satisfy the quadrilateral constraints. Each vector is extended from \mathbb{R}^{3n} to \mathbb{R}^{7n} by solving the standard matching equations.
- Each subsequent list L_i generates all nonnegative vectors in \mathbb{R}^{7n} that satisfy the quadrilateral constraints and that can be formed by (i) combining vectors from the previous list L_{i-1} and then (ii) adding or subtracting a multiple of the vertex-linking vector $\mathbf{v}(\ell(V_i))$. This list L_i is constructed from L_{i-1} using an algorithm similar to the filtered double description method of Section 6.1, though there are additional complications.

The final list L_m becomes the set of all vertices of the standard projective solution space that satisfy the quadrilateral constraints.

The key reason why the conversion algorithm works (in addition to those reasons discussed earlier in Section 6.1) is the following relationship between standard and quadrilateral coordinates:

Property C: The projection from standard to quadrilateral coordinates (whereby we simply remove the triangular coordinates $\{t_{i,j}\}$) is a linear map from the standard projective solution space to the quadrilateral projective solution space. Moreover, the kernel of this map is generated by the standard vector representations of the vertex links.

We can now see why the conversion algorithm works seamlessly for almost normal surfaces. If we replace standard and quadrilateral *normal* coordinates with standard and quadrilateral–octagon *almost normal* coordinates, the critical property C still holds. We can thereby follow through the algorithm and its proof as presented in [Burton 09b], and we find that the algorithm works as expected.

Specifically, what this algorithm achieves for almost normal surfaces is to begin with the set of all vertices of the *quadrilateral–octagon* projective solution space that satisfy the quadrilateral–octagon constraints, and to convert this to the (again typically much larger) set of all vertices of the *standard almost normal* projective solution space that satisfy the quadrilateral–octagon constraints.

As with direct enumeration, there is a catch involving the constraint (\star) , which we recall insists that each vector contain a nonzero octagonal coordinate whose value is set to one. For the same reasons as discussed in Section 6.1, we cannot enforce the constraint (\star) at each stage of the conversion algorithm. More importantly, *we cannot enforce (\star) on the set of input vectors*—the input must be the set of *all* vertices of the quadrilateral–octagon solution space that satisfy the quadrilateral–octagon constraints, whether these vertices yield many octagonal disks or whether they yield none. Once again, we must delay the enforcement of (\star) until the entire algorithm has finished running and we are ready to present our final results.

As a final note, we observe that the conversion algorithm allows us to enumerate all standard vertex almost normal surfaces using the following two-step procedure:

1. Use direct enumeration to enumerate all vertices of the *quadrilateral–octagon* projective solution space that satisfy the quadrilateral–octagon constraints, taking care not to enforce the extra constraint (\star) .

2. Use the conversion algorithm to recover all vertices of the *standard almost normal* projective solution space that satisfy the quadrilateral–octagon constraints, and thereby the set of all standard vertex almost normal surfaces.

As is the case with normal surfaces, experimentation shows that this two-step procedure runs orders of magnitude faster than a direct enumeration in standard coordinates.

6.3 Further Notes

We finish with some additional notes on the implementation and use of the enumeration and conversion algorithms.

Our first observation is the following. Although we work in $10n$ and $6n$ dimensions for standard almost normal and quadrilateral–octagon coordinates respectively, these large dimensions seem wasteful. The quadrilateral–octagon constraints guarantee at most one nonzero octagonal coordinate for each vector, so a different possibility might be to “select” a desired octagonal disk type and then work in $7n + 1$ or $3n + 1$ dimensions instead.

Casson has suggested such a technique [Jaco et al. 02], where we iterate through all $3n$ possible octagonal disk types, and for each such disk type we augment a traditional coordinate system for *normal* surfaces with a single coordinate for this octagon. As a result we obtain $3n$ distinct projective solution spaces, each with the significantly smaller dimension $7n + 1$ or $3n + 1$.

Although this reduction in dimensions is appealing, in practice both procedures essentially perform the same computations—by working in a full set of standard almost normal or quadrilateral–octagon coordinates, we are simply performing the $3n$ smaller enumerations of Casson “simultaneously.” This is because the quadrilateral–octagon constraints enforce at most one nonzero octagonal coordinate, and so the set of vertices at each stage of the enumeration algorithm is essentially the union of all $3n$ vertex sets in Casson’s scheme, with no additional “junk” vertices that must later be thrown away.

More importantly, however, any enumeration of vertex *almost normal* surfaces includes an implicit enumeration of vertex *normal* surfaces, since the quadrilateral–octagon constraints allow surfaces with no octagons at all. To this end, a single “simultaneous” enumeration in $10n$ or $6n$ dimensions should be more efficient, for if we run $3n$ independent enumerations in $3n$ different projective solution spaces, then we effectively perform this

implicit (and potentially slow [Burton 09b]) normal surface enumeration $3n$ distinct times.

Our second observation involves the constraint (\star) from Section 6.1—recall that this is the final condition of Theorems 3.7 and 4.5, where we insist that there be some nonzero octagonal coordinate, and that this coordinate be set to one. We have already observed that (\star) cannot be enforced during either the enumeration or conversion algorithm, and that we must instead apply it as a filter after the algorithms have finished.

It is worth noting that there are situations in which we do not want to enforce (\star) at all, even after the algorithms have run. We have already seen one example in Section 6.2, where the conversion algorithm requires that we do not enforce (\star) on the vertices in quadrilateral–octagon coordinates. Another example arises in applications in which we use the *vertex* almost normal surfaces as a basis to generate *all* almost normal surfaces (possibly with some limitations such as genus to keep the list finite).¹¹

In this case we cannot enforce (\star) either, since it is possible to obtain new admissible vectors through combinations of old vectors that break (\star) . For instance, we could combine an almost normal surface with a plain normal surface (having no octagonal disks) to obtain a new almost normal surface, or we could combine a surface with two octagons with a plain normal surface to obtain the double of a new almost normal surface, whereupon we simply divide by two.

7. MEASURING PERFORMANCE

In this section we measure the practical benefits of using quadrilateral–octagon coordinates. We do this by experimentally comparing running times for the 3-sphere recognition algorithm, using different coordinate systems for the critical step in which we enumerate vertex almost normal surfaces.

For our experiments we use the 15 smallest-volume homology 3-spheres from the closed hyperbolic census of [Hodgson and Weeks 94]. The reason for choosing homology 3-spheres is that we want to focus on almost normal surface enumeration—manifolds with nontrivial homology are eliminated in the first step of the 3-sphere recognition algorithm, and experience suggests that most *real* 3-spheres simplify to trivially small pieces during the

¹¹Such applications do appear in the literature; see [Lackenby 08] and [Rubinstein 04] for examples.

Hyperbolic Volume	Dehn Filling	Tetrahedra
1.39850888	m004(1, 2)	10
1.91221025	m011(2, 3)	11
2.22671790	m015(−3, 2)	11
2.25976713	m038(1, 2)	11
2.51622138	m081(3, 2)	12
2.62940540	m032(5, 2)	12
2.71245881	m120(−3, 2)	12
2.86563023	m137(−5, 1)	13
2.98683705	m137(5, 1)	13
3.08052001	m154(−2, 3)	12
3.08386105	m137(−6, 1)	14
3.16236729	m137(6, 1)	14
3.40043687	m222(−3, 2)	13
3.44586464	m199(−5, 1)	14
3.54091542	m260(−3, 2)	13

TABLE 1. The 15 homology 3-spheres used for experimentation.

decomposition procedure in the second step of the algorithm.¹²

We use 0-efficient triangulations of these homology 3-spheres, with sizes ranging from 10 to 14 tetrahedra. Table 1 lists the volume of each manifold, the size of each triangulation, and the Dehn filling given by Hodgson and Weeks to reconstruct each manifold. Each Dehn filling is applied to a cusped manifold from the hyperbolic census of [Hildebrand and Weeks 89].

For each of our 15 triangulations, we compare the running times for the following two procedures:

- 3-sphere recognition as given in Algorithm 5.4, using standard almost normal coordinates for the vertex enumeration in step 3 of the algorithm.
- The same algorithm, but using quadrilateral–octagon coordinates for the vertex enumeration in step 3, as authorized by Theorem 5.6.

All experiments were performed on a single 2.3 GHz AMD Opteron processor using the software package Regina [Burton 04, Burton 09a].

The running times are plotted in Figure 16 using log scales, and the results are extremely pleasing. Even in the worst case, quadrilateral–octagon coordinates still improve the running time by a factor of 30. At the other extreme, for several triangulations we find that quadrilateral–octagon coordinates improve the running time by factors of thousands, with an increase of just under 5000 times the speed for the best example.

¹²It is, however, possible to construct arbitrarily large 0-efficient triangulations of the 3-sphere [Jaco and Rubinstein 03].

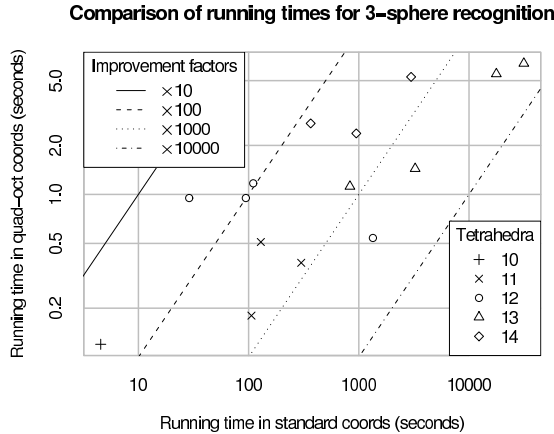


FIGURE 16. Performance comparisons for 3-sphere recognition.

8. JOINT COORDINATES

We finish this paper with an exploratory discussion of *joint coordinates* for octagonal almost normal surfaces. While quadrilateral–octagon coordinates reduce the dimension of the underlying vector space from $10n$ to $6n$, joint coordinates reduce this even further from $6n$ to $3n$. The key idea is to use negative coordinates for octagons and positive coordinates for quadrilaterals, noting from the quadrilateral–octagon constraints that the two cannot occur together within the same tetrahedron.

Joint coordinates have a number of appealing properties. Not only is their dimension small, but they carry the same information as quadrilateral–octagon coordinates (in contrast to the step from standard to quadrilateral–octagon coordinates, where we lose information about vertex links). Moreover, joint coordinates adhere to almost the same constraints in \mathbb{R}^{3n} as Tollefson’s quadrilateral coordinates for *normal* surfaces. In particular, they satisfy the original quadrilateral matching equations and quadrilateral constraints from Section 2.

There is a cost, however, which is the loss of convexity. For joint coordinates, we must allow one coordinate (but no more) to become negative. As a result, we no longer work in the nonnegative orthant of \mathbb{R}^{3n} , but rather the nonnegative orthant *and* the $3n$ “almost nonnegative” orthants that border it. This has severe consequences for the enumeration algorithms described in Section 6, which rely on convexity as a core requirement.

Nevertheless, it is pleasing to be able to express octagonal almost normal surfaces using essentially the same coordinate system as normal surfaces, and to do so in a way that portrays them as a natural extension of Tollefson’s original framework (where our extension involves simply stepping “just outside” the nonnegative orthant).

The layout of this section is as follows. We begin by describing the way in which we number quadrilateral and octagon types within each tetrahedron, which must be done carefully for joint coordinates to work. Following this, we define joint coordinates and develop the corresponding uniqueness and admissibility results. We then present an example using a one-tetrahedron triangulation, in which we show graphically how the vector representations of normal and almost normal surfaces appear in the corresponding solution space in \mathbb{R}^3 . To finish, we discuss how the loss of convexity affects both the projective solution space and the enumeration algorithms.

Definition 8.1. (Quadrilateral and octagon numbering.)

Let Δ be any tetrahedron in some compact 3-manifold triangulation. Within Δ , we number the quadrilateral and octagon types 1, 2, and 3 so that for each $i \in \{1, 2, 3\}$, the two edges of Δ that quadrilaterals of type i *never* meet are the same two edges of Δ that octagons of type i meet *twice*. This correspondence between quadrilaterals and octagons is illustrated in Figure 17.

This numbering scheme is very natural, in that the correspondence between quadrilaterals and octagons reflects the natural symmetries of these disks within a tetrahedron. One can also think of an octagon of type i as being obtained from a quadrilateral of type i by pulling each edge of the quadrilateral around and over the nearby vertex of the tetrahedron.

With our numbering scheme in place, we are now ready to define joint coordinates. Since we plan to present joint coordinates as a natural extension of Tollefson’s quadrilateral coordinates, we define them for *both* normal and almost normal surfaces as follows.

Definition 8.2. (Joint vector representation.)

Let \mathcal{T} be a compact 3-manifold triangulation formed from the n tetrahedra $\Delta_1, \dots, \Delta_n$, and let S be an embedded normal or octagonal almost normal surface in \mathcal{T} . For each tetrahedron Δ_i , let $q_{i,1}, q_{i,2}$, and $q_{i,3}$ denote the number of quadrilateral disks of each type, and let $k_{i,1}, k_{i,2}$, and $k_{i,3}$ denote the number of octagonal disks of each type in Δ_i contained in the surface S .

Then the *joint vector representation* of S , denoted by $\mathbf{j}(S)$, is the $3n$ -dimensional vector

$$\mathbf{j}(S) = (q_{1,1} - k_{1,1}, q_{1,2} - k_{1,2}, q_{1,3} - k_{1,3}; \\ q_{2,1} - k_{2,1}, q_{2,2} - k_{2,2}, q_{2,3} - k_{2,3}; \\ \dots, q_{n,3} - k_{n,3}).$$




FIGURE 17. The correspondence between quadrilaterals and octagons.

Our first task is to show that joint coordinates in $3n$ dimensions retain all of the information carried by quadrilateral–octagon coordinates in $6n$ dimensions.

Lemma 8.3. *Let \mathcal{T} be a compact 3-manifold triangulation, and let S and S' be embedded normal or octagonal almost normal surfaces in \mathcal{T} . Then the joint vector representations $\mathbf{j}(S)$ and $\mathbf{j}(S')$ are equal if and only if the quadrilateral–octagon vector representations $\mathbf{k}(S)$ and $\mathbf{k}(S')$ are equal.*

Here we define the quadrilateral–octagon vector representation for an embedded normal surface in the obvious way, by setting all octagonal coordinates to zero.

Proof. It is clear that if $\mathbf{k}(S) = \mathbf{k}(S')$, then $\mathbf{j}(S) = \mathbf{j}(S')$. Suppose conversely that $\mathbf{j}(S) = \mathbf{j}(S')$, and consider the (i, t) th coordinate $j_{i,t} = q_{i,t} - k_{i,t}$.

For both S and S' we know that $q_{i,t} \geq 0$ and $k_{i,t} \geq 0$. Moreover, since S and S' satisfy the quadrilateral–octagon constraints, we know that they can each have $q_{i,t} > 0$ or $k_{i,t} > 0$ but not both. It follows that for both S and S' we have one of the following situations:

- $j_{i,t} = 0$, in which case $q_{i,t} = k_{i,t} = 0$;
- $j_{i,t} = K > 0$, in which case $q_{i,t} = K$ and $k_{i,t} = 0$;
- $j_{i,t} = -K < 0$, in which case $q_{i,t} = 0$ and $k_{i,t} = K$.

That is, we can reconstruct the individual constituents $q_{i,t}$ and $k_{i,t}$ from the joint coordinate $j_{i,t}$, whereupon we obtain $\mathbf{k}(S) = \mathbf{k}(S')$. \square

As an immediate consequence of Lemmas 8.3 and 4.2, we obtain the following uniqueness result for joint vector representations:

Corollary 8.4. *Let \mathcal{T} be a compact 3-manifold triangulation, and let S and S' be embedded normal or octagonal almost normal surfaces in \mathcal{T} . Then the joint vector representations $\mathbf{j}(S)$ and $\mathbf{j}(S')$ are equal if and only if either (i) the surfaces S and S' are normal isotopic, or (ii) S and S' can be made normal isotopic by adding or removing vertex-linking components.*

We proceed now to give a complete classification of joint vector representations of embedded normal and octagonal almost normal surfaces. As indicated earlier, one of the appealing features of joint coordinates is that this classification corresponds precisely to Tollefson’s theorem for embedded *normal* surfaces (Theorem 2.6), except for the fact that we must allow one coordinate to become negative.

Theorem 8.5. *Let \mathcal{T} be a compact 3-manifold triangulation formed from n tetrahedra. An integer vector $\mathbf{w} \in \mathbb{R}^{3n}$ is the joint vector representation of an embedded normal or octagonal almost normal surface in \mathcal{T} if and only if:*

- at most one coordinate of \mathbf{w} is negative;
- \mathbf{w} satisfies the quadrilateral matching equations for \mathcal{T} (Definition 2.4);

- \mathbf{w} satisfies the quadrilateral constraints for \mathcal{T} (Definition 2.5);
- if there is a negative coordinate in \mathbf{w} , then this coordinate is set to -1 .

Moreover, such a vector represents an embedded normal surface in \mathcal{T} if and only if all of its coordinates are nonnegative.

It is worth pointing out that we interpret the quadrilateral matching equations and the quadrilateral constraints literally for any $3n$ -dimensional vector. We do not try to “reconstruct” quadrilateral coordinates from \mathbf{w} , but instead we read Definitions 2.4 and 2.5 precisely as given. In particular, the vector

$$\mathbf{w} = (j_{1,1}, j_{1,2}, j_{1,3}; \dots, j_{n,3}) \in \mathbb{R}^{3n}$$

is deemed to satisfy the quadrilateral constraints if at most one of $j_{i,1}$, $j_{i,2}$, and $j_{i,3}$ is nonzero for any given i . Likewise, \mathbf{w} satisfies the quadrilateral matching equations if for each nonboundary edge e of \mathcal{T} we have

$$j_{i_1, u_1} + j_{i_2, u_2} + \dots + j_{i_t, u_t} = j_{i_1, d_1} + j_{i_2, d_2} + \dots + j_{i_t, d_t},$$

where each u_k is the number of an upward quadrilateral type meeting e in the i_k th tetrahedron of \mathcal{T} , and each d_k is the number of a downward quadrilateral type meeting e in the i_k th tetrahedron of \mathcal{T} .

Proof of Theorem 8.5. Normal surfaces (as opposed to almost normal surfaces) are easily dealt with. Suppose that S is some embedded normal surface in \mathcal{T} . Then we have $\mathbf{j}(S) = \mathbf{q}(S)$, and it is clear from Theorem 2.6 that $\mathbf{j}(S)$ satisfies the four conditions given in this theorem, and that every coordinate of $\mathbf{j}(S)$ is nonnegative. Conversely, suppose that some integer vector $\mathbf{w} \in \mathbb{R}^{3n}$ satisfies these four conditions, and that all of its coordinates are nonnegative. Then \mathbf{w} satisfies the conditions of Theorem 2.6, whereupon it follows that \mathbf{w} is the quadrilateral vector representation—and therefore also the joint vector representation—of some embedded normal surface in \mathcal{T} .

We turn our attention now to the more interesting case of octagonal almost normal surfaces. The key observation is the following. Consider the quadrilateral–octagon matching equation derived from some nonboundary edge e of the triangulation, as described in Definition 4.3, and let Δ be some tetrahedron containing e . If we use the numbering scheme of Definition 8.1, then the i th quadrilateral type in Δ is an *upward* quadrilateral if and only

if the i th octagon type in Δ is a *downward* octagon, and vice versa. This is easily verified by examining Figure 9.

Using this observation, we can reduce each quadrilateral–octagon matching equation (4–1) to

$$\begin{aligned} & (q_{i_1, u_1} - k_{i_1, u_1}) + \dots + (q_{i_t, u_t} - k_{i_t, u_t}) \\ & = (q_{i_1, d_1} - k_{i_1, d_1}) + \dots + (q_{i_t, d_t} - k_{i_t, d_t}), \end{aligned} \quad (8-1)$$

where the coordinates $q_{i_1, u_1}, q_{i_2, u_2}, \dots, q_{i_t, u_t}$ and $q_{i_1, d_1}, q_{i_2, d_2}, \dots, q_{i_t, d_t}$ correspond to the upward and downward quadrilaterals respectively about the edge e . Translated into joint coordinates, this reduces further to

$$j_{i_1, u_1} + \dots + j_{i_t, u_t} = j_{i_1, d_1} + \dots + j_{i_t, d_t}, \quad (8-2)$$

which is identical to the corresponding quadrilateral matching equation in \mathbb{R}^{3n} .

We can now finish the proof of Theorem 8.5. Suppose that S is some octagonal almost normal surface in \mathcal{T} . Then the following observations follow immediately from Theorem 4.5:

- Precisely one octagonal coordinate in $\mathbf{k}(S)$ is nonzero, and the corresponding quadrilateral coordinate in $\mathbf{k}(S)$ must be zero as a result. Therefore precisely one coordinate of $\mathbf{j}(S)$ is negative.
- The quadrilateral–octagon vector representation $\mathbf{k}(S)$ satisfies each quadrilateral–octagon matching equation as described by (8–1). Therefore the joint vector representation $\mathbf{j}(S)$ satisfies each *quadrilateral* matching equation, as described by (8–2).
- For each tetrahedron of \mathcal{T} , at most one of the six corresponding quadrilateral and octagonal coordinates in $\mathbf{k}(S)$ is nonzero, and so at most one of the three corresponding joint coordinates in $\mathbf{j}(S)$ is nonzero. Therefore $\mathbf{j}(S)$ satisfies the quadrilateral constraints.
- The unique nonzero octagonal coordinate in $\mathbf{k}(S)$ has value $+1$, and so the unique negative coordinate in $\mathbf{j}(S)$ has value -1 .

Therefore the joint vector representation $\mathbf{j}(S)$ satisfies all four conditions listed in the statement of this theorem.

Conversely, suppose that some integer vector

$$\mathbf{w} = (j_{1,1}, j_{1,2}, j_{1,3}; \dots, j_{n,3}) \in \mathbb{R}^{3n}$$

satisfies all four conditions listed in this theorem statement and that one of its coordinates is negative (recall that the nonnegative case was dealt with at the beginning of this proof). We define the $6n$ -dimensional vector

$$\mathbf{w}' = (q_{1,1}, q_{1,2}, q_{1,3}, k_{1,1}, k_{1,2}, k_{1,3}; \dots, k_{n,3}) \in \mathbb{R}^{6n}$$




FIGURE 18. The one-tetrahedron triangulation \mathcal{S} .

by setting

$$q_{i,t} = \begin{cases} j_{i,t} & \text{if } j_{i,t} \geq 0; \\ 0 & \text{if } j_{i,t} < 0, \end{cases}$$

and

$$k_{i,t} = \begin{cases} 0 & \text{if } j_{i,t} \geq 0; \\ -j_{i,t} & \text{if } j_{i,t} < 0. \end{cases}$$

Using the four conditions of this theorem statement and following the previous argument in reverse, it is simple to show that \mathbf{w}' satisfies the conditions of Theorem 4.5. It follows then that \mathbf{w}' is the quadrilateral–octagon vector representation of some octagonal almost normal surface in \mathcal{T} , and so \mathbf{w} is the joint vector representation of this same surface. \square

Because joint coordinates are $3n$ -dimensional, we are able to visualize them explicitly in \mathbb{R}^3 for a one-tetrahedron triangulation. We do just this in the following example to illustrate the various conditions of Theorem 8.5.

Example 8.6. Let \mathcal{S} be the following compact 3-manifold triangulation (which we will shortly prove represents the 3-sphere). We begin with the single tetrahedron $ABCD$ and identify faces $ABC \leftrightarrow BCD$ (with a twist) and $ABD \leftrightarrow ACD$ (folded directly over the common edge AD), as illustrated in Figure 18. The resulting triangulation has one tetrahedron, one vertex (since $A, B, C,$ and D are all identified), and two edges (where $AB, BC, CD, BD,$ and AC are all identified, and AD is left in a class of its own).

Let Δ represent the sole tetrahedron of \mathcal{S} , and number the three quadrilateral types in Δ so that types 1, 2, and 3 separate the edge pairs $(AB, CD), (AC, BD),$ and (AD, BC) respectively. We find then that both joint matching equations for \mathcal{S} (one for each edge) reduce to the form $j_{1,1} = j_{1,2}$.

We plot the resulting solution space in \mathbb{R}^3 in Figure 19. In the top-left diagram, we shade the region in which

$\mathbf{w} = (j_{1,1}, j_{1,2}, j_{1,3}) \in \mathbb{R}^3$ has at most one negative coordinate (for clarity, we restrict our shading to the interior of a sphere around the origin). In the top-right diagram, we shade the intersection of this region with the hyperplane $j_{1,1} = j_{1,2}$, which gives us the closed half-plane $H = \{(x, x, z) \mid x \geq 0\}$. If we wish to enforce the quadrilateral constraints, then we must restrict our attention to the three coordinate axes (where at most one coordinate is nonzero); the final intersection of H with these three axes is plotted in the bottom-left diagram. The resulting solution space is simply the entire $j_{1,3}$ -axis, taken in both directions.

From Theorem 8.5 it follows that if we ignore vertex-linking components, then the normal and octagonal normal surfaces in \mathcal{S} correspond precisely to the integer points

$$\{(0, 0, k) \mid k \in \mathbb{Z}, k \geq -1\}.$$

With some further investigation we can classify these surfaces as follows:

- $(0, 0, 0)$ represents the empty surface;
- $(0, 0, k)$ for $k \geq 1$ represents k copies of the embedded normal torus surrounding the edge AD , which is referred to in [Jaco and Rubinstein 03] as a *thin edge link*;
- $(0, 0, -1)$ is an octagonal almost normal 2-sphere.

These surfaces are individually plotted in the bottom-right diagram of Figure 19.

To finish, we note that (i) there are no normal 2-spheres (aside from the vertex link that we have ignored), and so \mathcal{S} is a 0-efficient triangulation, and that (ii) \mathcal{S} contains an octagonal almost normal 2-sphere. Using Theorem 5.3 and noting that \mathcal{S} is orientable, it follows that \mathcal{S} is in fact a triangulation of the 3-sphere.

At this point, it is natural to ask what becomes of the projective solution space in joint coordinates. Recall that in other coordinate systems, the nonnegative orthant and the matching equations intersect to give a convex polyhedral cone, and that the projective solution space is a cross-section of this cone, taken by intersecting the cone with the hyperplane $\sum x_i = 1$.

The difficulty we face with joint coordinates is that we no longer have a convex polyhedral cone to work with. Instead we begin with the union of $3n + 1$ orthants in \mathbb{R}^{3n} (where at most one coordinate is nonnegative), which is not even a convex set. Upon intersecting this with the joint matching equations, we obtain a set \mathcal{P} with the

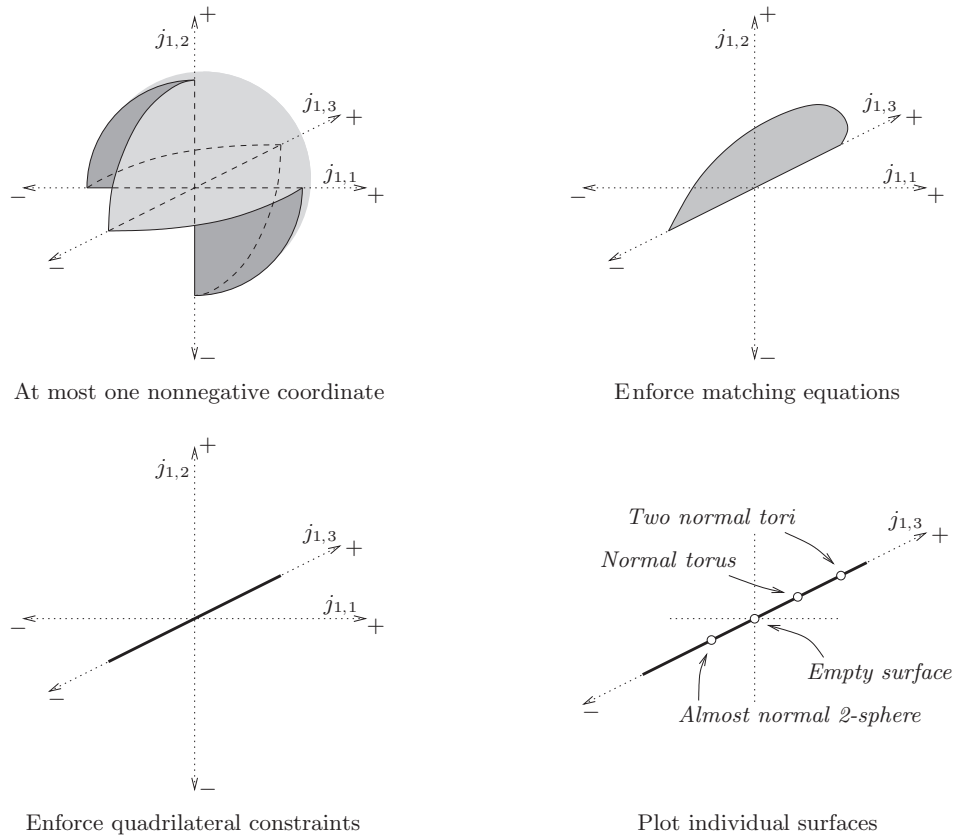


FIGURE 19. Building the solution space for the triangulation \mathcal{S} .

following properties. \mathcal{P} is a cone in the sense that $\mathbf{x} \in \mathcal{P}$ implies that $\lambda \mathbf{x} \in \mathcal{P}$ for any $\lambda \geq 0$, but like the union of orthants before it, \mathcal{P} might not be convex (although in Example 8.6 it happens to be). More importantly, \mathcal{P} can contain diametrically opposite points (such as $(0, 0, \pm 1)$ in our example), and so in general we cannot form a cross-section by slicing through \mathcal{P} with a hyperplane.

We could perhaps take a cross-section using the unit sphere, but this would lift us out of the world of polytopes, making it difficult to design algorithms. Perhaps the simplest solution is to take a cross-section using the “polyhedral unit sphere” $\sum |x_i| = 1$, as illustrated in Figure 20. Continuing with Example 8.6, the left-hand diagram of Figure 20 shows the intersection of our four original orthants with the “sphere” $|j_{1,1}| + |j_{1,2}| + |j_{1,3}| = 1$, and the right-hand diagram shows the subsequent intersection with the matching equation $j_{1,1} = j_{1,2}$.

Although we are now able to define the projective solution space in a sensible way, we do not obtain a single convex polytope as a result. Instead we obtain $3n + 1$ distinct convex polytopes—one for each of the original orthants—joined together along their boundaries. To

enumerate the vertices of this structure would therefore require $3n + 1$ distinct passes through the vertex enumeration algorithm,¹³ bringing us back to the scheme of Casson that we discussed in Section 6.3. It is worth noting again that the polytope of this structure that sits within the nonnegative orthant is precisely Tollefson’s quadrilateral projective solution space for *normal* surfaces.

It follows then that joint coordinates do not appear practical for use in enumeration algorithms. Nevertheless, they have appealing geometric properties that may render them useful for other purposes:

- They live in a remarkably small number of dimensions.
- They express the space of admissible vectors for octagonal almost normal surfaces as a natural geometric extension of Tollefson’s space for normal surfaces, obtained simply by expanding our scope from the

¹³In fact, only $3n$ passes are required. We can ignore the nonnegative orthant, since it can be shown that every vertex of the nonnegative orthant is also a vertex of one of the $3n$ adjacent “almost nonnegative” orthants.

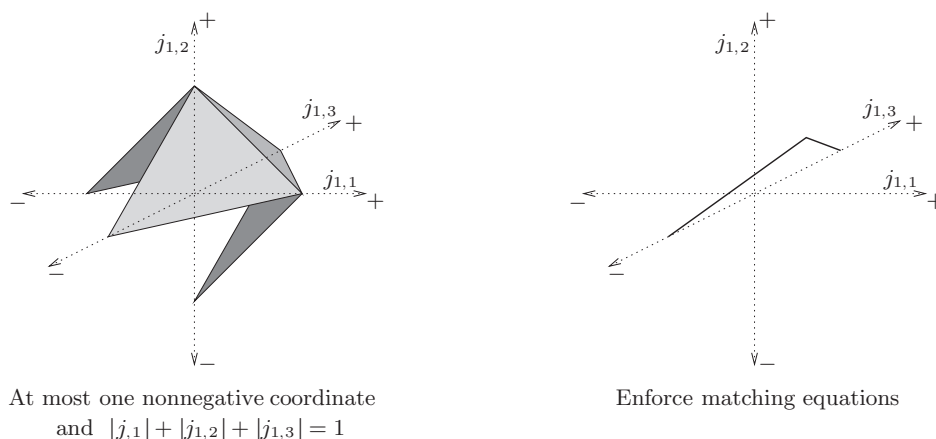


FIGURE 20. Constructing the “projective solution space” in joint coordinates.

nonnegative orthant to include the neighboring “almost nonnegative” orthants.

As an immediate application, these properties make joint coordinates a useful tool for visualizing the almost normal solution space. More generally, they could perhaps open the way for new *theoretical* insights into the structure of the solution space.

To illustrate the latter point, we can draw analogies with Casson’s edge weight coordinates for normal surfaces, which are developed and exploited in [Burton 03]. Like joint coordinates, edge weight coordinates use very few dimensions and are geometrically appealing, but a loss of convexity makes them impractical for use in algorithms. Nevertheless, their tight geometric structure has led to new theoretical and combinatorial insights, and we hope that joint coordinates can offer the same.

ACKNOWLEDGMENTS

The author is grateful to the Victorian Partnership for Advanced Computing for the use of their excellent computing resources, to the University of Melbourne for their continued support for the software package Regina, and to the anonymous referees for their thoughtful suggestions. The author is supported by the Australian Research Council under the Discovery Projects funding scheme (project DP1094516).

REFERENCES

- [Agol et al. 02] Ian Agol, Joel Hass, and William Thurston. “3-Manifold Knot Genus is NP-Complete.” In *STOC ’02: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, pp. 761–766. New York: ACM Press, 2002.
- [Burton 03] Benjamin A. Burton. “Minimal Triangulations and Normal Surfaces.” PhD thesis, University of Melbourne, 2003. Available online (<http://regina.sourceforge.net/>).
- [Burton 04] Benjamin A. Burton. “Introducing Regina, the 3-Manifold Topology Software.” *Experiment. Math.* 13:3 (2004), 267–272.
- [Burton 09a] Benjamin A. Burton. “Regina: Normal Surface and 3-Manifold Topology Software.” Available online (<http://regina.sourceforge.net/>), 1999–2009.
- [Burton 09b] Benjamin A. Burton. “Converting between Quadrilateral and Standard Solution Sets in Normal Surface Theory.” *Algebr. Geom. Topol.* 9:4 (2009), 2121–2174.
- [Burton 10a] Benjamin A. Burton. “Optimizing the Double Description Method for Normal Surface Enumeration.” *Math. Comp.* 79 (2010), 453–484.
- [Burton 10b] Benjamin A. Burton. “The Complexity of the Normal Surface Solution Space.” In *SCG ’10: Proceedings of the Twenty-Sixth Annual Symposium on Computational Geometry*, pp. 201–209. New York: ACM, 2010.
- [Haken 61] Wolfgang Haken. “Theorie der Normalflächen.” *Acta Math.* 105 (1961), 245–375.
- [Haken 62] Wolfgang Haken. “Über das Homöomorphieproblem der 3-Mannigfaltigkeiten I.” *Math. Z.* 80 (1962), 89–120.
- [Hass et al. 99] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. “The Computational Complexity of Knot and Link Problems.” *J. Assoc. Comput. Mach.* 46:2 (1999), 185–211.
- [Hildebrand and Weeks 89] Martin V. Hildebrand and Jeffrey R. Weeks. “A Computer Generated Census of Cusped Hyperbolic 3-Manifolds.” In *Computers and Mathematics (Cambridge, MA, 1989)*, pp. 53–59. New York: Springer, 1989.

- [Hodgson and Weeks 94] Craig D. Hodgson and Jeffrey R. Weeks. “Symmetries, Isometries and Length Spectra of Closed Hyperbolic Three-Manifolds.” *Experiment. Math.* 3:4 (1994), 261–274.
- [Jaco and Oertel 84] William Jaco and Ulrich Oertel. “An Algorithm to Decide If a 3-Manifold Is a Haken Manifold.” *Topology* 23:2 (1984), 195–209.
- [Jaco and Rubinstein 03] William Jaco and J. Hyam Rubinstein. “0-Efficient Triangulations of 3-Manifolds.” *J. Differential Geom.* 65:1 (2003), 61–168.
- [Jaco and Tollefson 95] William Jaco and Jeffrey L. Tollefson. “Algorithms for the Complete Decomposition of a Closed 3-Manifold.” *Illinois J. Math.* 39:3 (1995), 358–406.
- [Jaco et al. 02] William Jaco, David Letscher, and J. Hyam Rubinstein. “Algorithms for Essential Surfaces in 3-Manifolds.” In *Topology and Geometry: Commemorating SISTAG*, Contemporary Mathematics 314, pp. 107–124. Providence: Amer. Math. Soc., 2002.
- [Kneser 29] Hellmuth Kneser. “Geschlossene Flächen in dreidimensionalen Mannigfaltigkeiten.” *Jahresbericht der Deutschen Math. Verein.* 38 (1929), 248–260.
- [Lackenby 08] Marc Lackenby. “An Algorithm to Determine the Heegaard Genus of Simple 3-Manifolds with Nonempty Boundary.” *Algebr. Geom. Topol.* 8:2 (2008), 911–934.
- [Motzkin et al. 53] T. S. Motzkin, H. Raiffa, G. L. Thompson, and R. M. Thrall. “The Double Description Method.” In *Contributions to the Theory of Games, Vol. II*, edited by H. W. Kuhn and A. W. Tucker, Annals of Mathematics Studies 28, pp. 51–73. Princeton: Princeton University Press, 1953.
- [Rubinstein 95] J. Hyam Rubinstein. “An Algorithm to Recognize the 3-Sphere.” In *Proceedings of the International Congress of Mathematicians (Zürich, 1994)*, Vol. 1, pp. 601–611. Basel: Birkhäuser, 1995.
- [Rubinstein 97] J. Hyam Rubinstein. “Polyhedral Minimal Surfaces, Heegaard Splittings and Decision Problems for 3-Dimensional Manifolds.” In *Geometric Topology (Athens, GA, 1993)*, AMS/IP Stud. Adv. Math., Vol. 2, pp. 1–20. Providence: Amer. Math. Soc., 1997.
- [Rubinstein 04] J. Hyam Rubinstein. “An Algorithm to Recognise Small Seifert Fiber Spaces.” *Turkish J. Math.* 28:1 (2004), 75–87.
- [Thompson 94] Abigail Thompson. “Thin Position and the Recognition Problem for S^3 .” *Math. Res. Lett.* 1:5 (1994), 613–630.
- [Thurston 78] William P. Thurston. “The Geometry and Topology of 3-Manifolds.” Lecture notes, Princeton University, 1978.
- [Tillmann 08] Stephan Tillmann. “Normal Surfaces in Topologically Finite 3-Manifolds.” *Enseign. Math. (2)* 54 (2008), 329–380.
- [Tollefson 98] Jeffrey L. Tollefson. “Normal Surface Q -Theory.” *Pacific J. Math.* 183:2 (1998), 359–374.
- [Wilson 08] Robin Todd Wilson. “Meridional Almost Normal Surfaces in Knot Complements.” *Algebr. Geom. Topol.* 8:3 (2008), 1717–1740.

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

Received May 6, 2009; accepted July 10, 2009.