Computing Immersed Normal Surfaces in the Figure-Eight Knot Complement

Richard Rannard

CONTENTS

The theory of (embedded) normal surfaces is a powerful technique in 3-manifold topology. There has been much recent interest in extending the theory to immersed surfaces, in particular to attack the word problem for 3-manifolds. Progress in this area has been hindered by the lack of nontrivial examples. This paper and the related work [Matsumoto and Rannard 1997] cover a particular example in depth, using methods which may be generalized. We give detailed information on the existence of immersed surfaces in the figure-eight knot complement using nontrivial computational techniques. After an introduction to the theory, introducing some new concepts, we discuss some strategies for enumerating surfaces of low genus, which have been implemented in software written by the author. The results are tabulated and an unusual example discussed.

1. INTRODUCTION

In this paper we apply computational methods to determine data on the existence of immersed normal surfaces in the figure-eight knot complement. The difficulties of these computations are such that only two examples have previously appeared in the literature; see [Aitchison et al. 1997]. By investigating a single example in detail, we aim to furnish examples (and counter-examples) to aid in the development of the theory. The computational methods used are also of interest; we show how the problem of enumerating surfaces can be reduced to enumerating the elements of a permutation group which have certain properties. This problem may be transformed into the question of searching a certain graph, for which various existing strategies are compared. All the algorithms mentioned have been implemented in software by the author.

The data obtained in this paper are used in [Matsumoto and Rannard 1997]; those applications motivate some of the methods used here.

This paper is structured as follows. Section 2 gives a short introduction to the "classical" theory

Keywords: 3-manifolds, immersed normal surfaces, normal surface theory, algorithms

1991 Mathematics Subject Classification: 57N35, 57M50

1058-6458/1999 $0.50 per page

of (embedded) normal surfaces. Section 3 describes an extension of normal surface theory to cover immersed surfaces. Section 4 describes an extension of normal describes the application of this theory to the figure-eight knot complement. Section 5 contains a description of the algorithms used in the enumeration of surfaces, and their implementation in software. Section 6 describes and interprets the results obtained by these methods. The concluding section takes up one particularly interesting immersed normal surface.

2. EMBEDDED NORMAL SURFACE THEORY

The theory of embedded normal surfaces was first developed by Haken [1961] and used by him to prove several important results in the theory of 3-manifolds; see also [Jaco and Oertel 1984]. We use a different but equivalent approach for which good references are [Hemion 1992; Tollefson 1998; Jaco and Rubinstein 1989]. None of the material in this section is original.

The basis of normal surface theory is the following sequence of ideas. Let M be a compact 3-manifold with or without boundary, and suppose Δ is a fixed triangulation of M. We allow Δ to be a proper triangulation or a pseudotriangulation, and the tetrahedra in Δ may or may not be ideal. Given a surface F embedded in M, try to isotope F so that it meets each 3-simplex of the triangulation Δ in a set of standard pieces, called *normal disks.* Such a surface is called a *normal surface*; many "interesting" surfaces can be made normal. A normal surface generates a sequence of numbers which are the number of times each possible normal disk occurs. We think of this sequence as a point in a vector space over $\mathbb{R}$. The set of points which may come from surfaces lie in a cone in the vector space called the *solution space.*

We denote a 3-manifold M with triangulation Δ by (M, Δ). This paper deals with one such M and Δ, the ideal triangulation of the figure-eight knot complement given in [Thurston 1997].

Let τ be a tetrahedron forming part of Δ.

Definition 2.1. A *normal arc* in τ is an embedded arc on some face of $\partial\tau$ which has endpoints on different edges of the face. A *normal disk* in τ is a disk properly embedded in τ which intersects each face of τ in at most one normal arc. Thus a normal disk is either

1. a triangle cutting off a vertex, or
2. a quadrilateral separating a pair of vertices,

whose edges have endpoints on different edges of the tetrahedron.

Two normal disks are of the same *type* if one can be isotoped into another while preserving these conditions and without an edge of a disk passing over a vertex of a tetrahedron. Thus there are exactly seven types of normal disks for each tetrahedron in Δ, as shown in Figure 1. Normal disks may have only three or four sides; those with three sides are called *T-disks*, while those with four sides are called *Q-disks.* This terminology is due to Tollefson [1998].

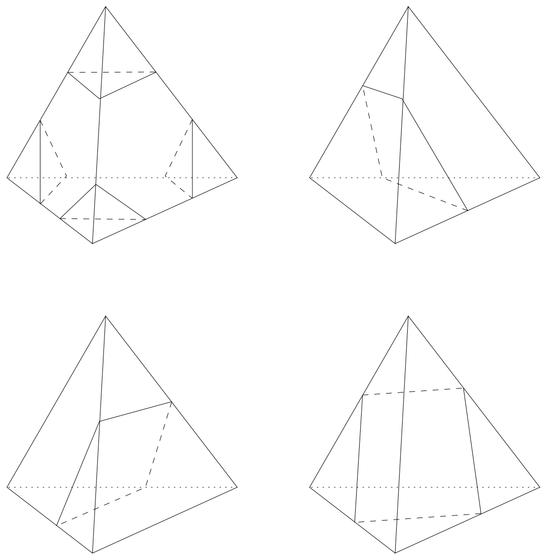

FIGURE 1. The seven combinatorial types of normal disks in a tetrahedron.

Definition 2.2. A surface F properly embedded in M is a *normal surface* if F meets each tetrahedron of Δ in a collection of normal disks.

Whether or not a surface is normal thus depends on the triangulation used.

A surface may be perturbed slightly by an isotopy and still intersect the triangulation in the same set of normal disks. A *normal isotopy* is an ambient isotopy of F which doesn't change the type of any of the normal disks of F. We shall not distinguish between normal surfaces which are equivalent under normal isotopy.

We shall consider only closed normal surfaces in this paper, corresponding to surfaces in S^3 which do not meet the figure-eight knot.

If there are t tetrahedra in a triangulation Δ, there will be $7t$ different normal disks in Δ. Let $W(M_8, \Delta)$ be a copy of $\mathbb{R}^{14}$, and associate with each normal disk type a different coordinate of $W(M_8, \Delta)$ (in Section 4 we will give a numbering of the disks in our example). We interpret the normal surface having w_i normal disks of type i as corresponding to the point $(w_1, \ldots, w_{14})$ in $W(M_8, \Delta)$, where $w_i \geq 0$, for $i = 1, \ldots, 14$.

We may determine the set of possible vectors corresponding to surfaces in the following way. Choose a face of Δ common to two tetrahedra τ_1 and τ_2. For any normal surface F in M, each normal disk in τ_1 which meets the face is joined to a normal disk in τ_2 along an edge (see Figure 2), and these two normal disks meet the face in the same normal arc type. Hence each normal arc type gives a linear equation of the form $w_i + w_j = w_k + w_l$, where the terms are the numbers of normal disks of fixed type. The set of all such equations from all normal arc types of Δ is called the set of *normal matching equations*.

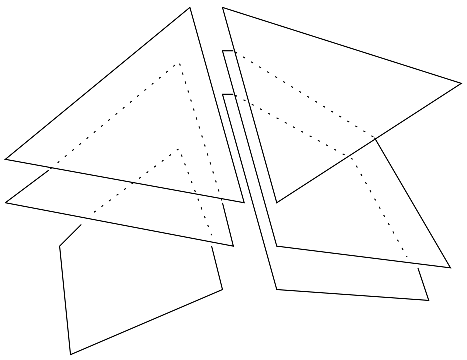

FIGURE 2. A set of normal disks satisfying the normal matching equation for one normal arc type. The disks join along arcs of that normal arc type.

Since the normal matching equations are linear, the set of points which satisfy them is a linear subspace of $W(M_8, \Delta)$. We are not interested in points which have negative or noninteger coordinates, since these cannot correspond to surfaces. The *normal solution space* is the set of vectors in this subspace which have all coordinates nonnegative. A *solution vector* is a vector in the solution space with integer coordinates. Each normal surface corresponds to a single solution vector.

We can think of the matching equations as being a set of linear Diophantine equations with solutions the solution vectors. Then the theory of Diophantine equations implies that there is a finite set of solution vectors which generate all solution vectors by addition; see [Hemion 1992]. We call the solution vectors in the generating set *fundamental solution vectors*; every solution vector can be written as a sum of nonnegative multiples of the fundamental solution vectors.

We may ask whether all solution vectors correspond to normal surfaces. This is not the case if any tetrahedron of Δ contains two Q-disks of different types, since any two such disks must intersect (Figure 3).

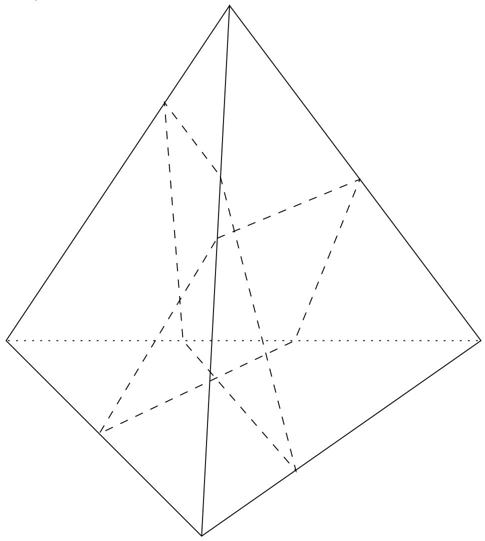

FIGURE 3. Two Q-disks of different types in the same tetrahedron must intersect.

Any solution vector which doesn't have this property corresponds to unique embedded normal surface; see [Hemion 1992].

This criterion does not apply in the case of immersed normal surfaces.

The *projective solution space* is the intersection of the normal solution space with the hyperplane of codimension 1 defined by the equation

$$x_1 + \cdots + x_{14} = 1.$$

It is a convex polytope of dimension one less than the solution space. The projective solution space carries a surprisingly large amount of information about the topology of the manifold; see [Jaco and Oertel 1984; Jaco and Tollefson 1995].

3. IMMERSED NORMAL SURFACE THEORY

In this section we study an extension of the above theory to immersed surfaces. We begin with a definition of an immersed normal surface.

Definition 3.1. A closed surface F properly immersed in (M, Δ) is an *immersed normal surface* if it meets the 2-skeleton of Δ transversally and meets each tetrahedron of Δ in a set of normal disks.

Clearly we can associate to each immersed normal surface a vector as for embedded normal surfaces. Then we can utilize all the apparatus from classical normal surface theory, in particular the normal matching equations and solution space.

We relax the last condition in Section 2, that we cannot have two Q-disks of different type in the same tetrahedron, since this is clearly no longer necessary for immersed surfaces. We now consider every vector in the normal solution space as possibly giving an immersed normal surface. To derive a surface from a vector, we take the normal disks corresponding to the vector and try to join them together along their edges.

3A. The Gluing Group

The different ways of joining normal disks are in 1-1 correspondence with a certain permutation group.

Fix one particular position for the normal disks. A normal disk can only be joined to another normal disk where they meet a face of the triangulation in a common normal arc. The possible ways of gluing disks along a given normal arc type are in 1-1 correspondence with a permutation group. If there are n normal arcs of a given type, then there are n normal disks in one tetrahedron which have an edge of that normal arc type, glued to n normal disks in the adjacent tetrahedron. The possible ways of joining them are parametrized by the symmetric group on n letters.

For each normal arc type, the disks can be glued in any manner, since we are ignoring for the moment the question of whether the resulting 2-complex has singularities. Hence, the set of all possible ways of sticking the disks together corresponds to a direct product of permutation groups, one for each normal arc type.

Definition 3.2. The *gluing group* G_H for a solution vector H is the direct product of the symmetric groups corresponding to each normal arc type in the triangulation, together with a numbering of the normal disks. A *gluing* is an element of the gluing group (i.e., a permutation), corresponding to a specific joining of the normal disks along their edges.

A gluing is in a sense the fundamental object in immersed normal surface theory. We shall call a gluing and its associated 2-complex by the same symbol, and rely on context to distinguish them.

Note that the gluing group is defined entirely by a normal vector, since it depends only on the number of each type of normal arc, which in turn depend only on the numbers of each type of normal disk. A presentation for it can easily be written down.

A useful way to visualize immersed normal surfaces in this context is to imagine isotoping all the normal disks so that their vertices lie on the midpoints of the edges of the tetrahedra (with respect to some well-behaved continuous polyhedral metric, such as used in [Jaco and Rubinstein 1989]). The surface is then not in general position, in fact normal disks of the same type will be superimposed. All the normal disks which have a common arc type will meet in a single arc of that type. The idea is to hold the double arcs more firmly by pushing them into the 2-skeleton of the manifold; otherwise normal isotopies of the surface can push a double arc from one tetrahedron into another, making two normally isotopic surfaces appear different. We call this positioning of the normal disks the *standard position* of the gluing.

There is a "canonical" element in every gluing group, as we now show. We can modify any gluing g by removing all intersections of the normal disks which take place along normal arcs (thinking of g being in standard position). The result is a gluing where the only normal disks which cross are Q-disks of different types. This gluing is clearly unique, and from now on we shall number the normal disks so that this gluing corresponds to the identity element of the gluing group.

We shall see below that normal surfaces and gluings are not equivalent concepts. In general there may be gluings which do not correspond to nonsingular immersions of surfaces, and multiple gluings may represent the same surface.

3B. Regular Gluings

We are really only interested in gluings which are topologically 2-manifolds, that is images of nonsingular immersions of surfaces.

Definition 3.3. A *regular vector* is a solution vector which corresponds to a set of normal disks which may be joined to give a nonsingular immersed surface. Such a way of joining the normal disks is called a *regular gluing*.

A central problem in immersed normal surface theory is to find conditions on a solution vector which imply it must be regular. This is a very difficult problem; for partial results see [Letscher 1997; Matsumoto and Rannard 1997; Rannard 1997]. Note that several nonisotopic regular gluings may exist for a single regular vector.

For brevity we use the term "regular surface" to mean "regular immersed normal surface".

The possible singularities of a gluing are of a particular type. Let p be a point where the gluing meets the 1-skeleton of the triangulation, and consider how the normal disks in the gluing with p as a vertex meet. If the gluing represents a nonsingular immersion of a surface, then these normal disks will form a disk around p. However, in forming the gluing by joining normal disks together, there is no guarantee that this will happen. It may be that the normal disks will "go around p several times", or more formally, that the link of p will have winding number about the edge greater than 1. Let the winding number be n; we call this a *branch point* of *order* n. Figure 4 shows a branch point of order 2. If $n = 1$, then the gluing is locally the immersion of a manifold at p, while if $n \neq 1$ then the immersion has an isolated singularity at the point p. Clearly a gluing is regular if and only if it contains no branch points.

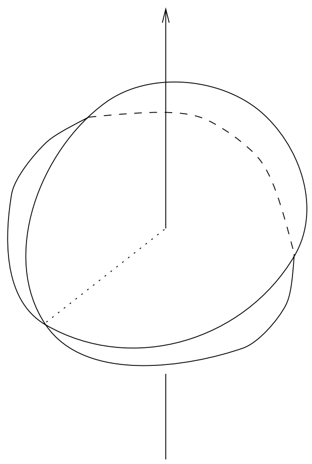

FIGURE 4. A branch point of order 2.

A regular gluing has a natural decomposition into quadrilaterals and triangles, given by the normal disks. The dual subdivision divides the gluing into a set of polygons, centered on the points of intersection of the gluing with the 1-skeleton of Δ. We call these polygons *edge disks*. A typical example of an edge disk is shown in Figure 5. A gluing is regular if and only if it meets the 1-skeleton of Δ always in edge disks. Edge disks will be used extensively in [Matsumoto and Rannard 1997].

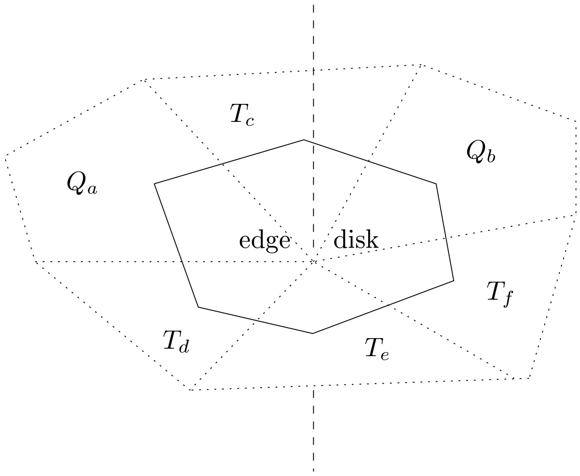

FIGURE 5. An edge disk. The vertical dashed line represents an edge of Δ.

Embedded normal surfaces may also be understood in terms of regular gluings; each regular vector represents at most one embedded regular gluing. In the numbering of the normal disks given above, each embedded normal surface corresponds to the identity of the corresponding gluing group. Such a class may have immersed regular gluings in addition to the embedded one.

3C. Normal Surfaces are Really Equivalence Classes of Gluings

Let H be a solution vector with some coordinate greater than 1. Then H represents a set of normal disks, and two of these normal disks will be of the same type; call them δ_1 and δ_2. Let g be a gluing in the gluing group G_H, and form the gluing g' by deleting δ_1 and δ_2 from g (to give a gluing with two boundary components) then replacing δ_2 in place of δ_1 and vice versa. We have swapped δ_1 and δ_2 in the gluing g, as shown in Figure 6. Clearly g and g' represent normally isotopic immersions of the same surface, however they are different elements of the gluing group. We have shown:

Proposition 3.4. *Up to normal isotopy, a normal surface corresponds to an equivalence class of gluings under the following equivalence relation: two gluings are equivalent if and only if one can be obtained from the other by exchanging normal disks of the same type.*

We denote the normal surface containing the gluing g by $[g]$.

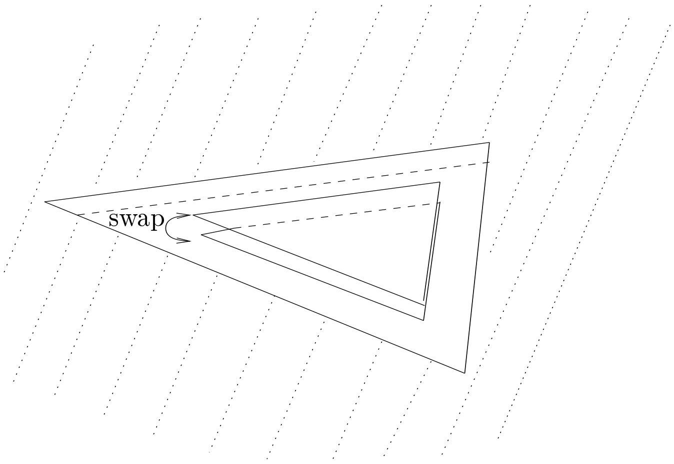

FIGURE 6. Swapping two normal disks of the same type gives an identical normal surface up to normal isotopy.

The set of regular normal surfaces has a natural algebraic description. Swapping δ_1 and δ_2 in g as described above corresponds to multiplying g by an element of the gluing group (as an arbitrary convention we assume multiplication on the right). Moreover, the group element so defined is independent of g. For a given vector H, there is a subset of the gluing group G_H consisting of all elements which swap the edges of a pair of normal disks of the same type; let A_H be the subgroup of G_H generated by this set. An element of A_H is a product of group elements corresponding to swapping normal disks of the same type, so multiplying any gluing by such an element is equivalent to performing a series of such swaps. We call A_H the *anonymity subgroup* of G_H. The normal surfaces of H correspond to the (right) cosets $\{A_H g \mid g \in G_H\}$.

Since swapping normal disks of the same type preserves the regularity of a gluing, the regular normal surfaces of H form a partition of the regular gluings of G_H. Each coset is composed of gluings which are either all regular or all irregular. For more details see [Rannard 1997].

4. EMBEDDED NORMAL SURFACE THEORY FOR THE FIGURE-EIGHT KNOT COMPLEMENT

The complement of the figure-eight knot in the 3-sphere is a hyperbolic manifold M_8 of finite volume with a single cusp [Thurston 1997]. It can be triangulated using two ideal tetrahedra, as illustrated in Figure 7, where the faces of the tetrahedra are identified so that the marked arrows agree. Let Δ denote this triangulation; Δ has two edges and four

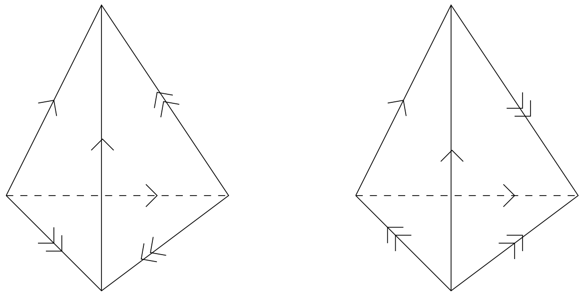

FIGURE 7. The canonical ideal triangulation Δ of the figure-eight knot complement.

faces, and is a pseudotriangulation of M_8. We write the figure-eight knot complement with this triangulation as the pair (M_8, Δ).

The normal disks for (M_8, Δ) are shown in Figure 8. There are 7 normal disks in each tetrahedron, making 14 in total. To each normal surface we associate the vector

$$(t_1, t_2, t_3,\ t_4, q_1, q_2, q_3,\ t_1', t_2', t_3',\ t_4', q_1', q_2', q_3') \in \mathbb{Z}^{14}$$

where t_i and q_i are the numbers of normal disks of type T_i and Q_i, respectively. The normal solution space is the set of vectors with nonnegative integer coordinates in the 4-dimensional subspace of $W(M_8, \Delta) \approx \mathbb{R}^{14}$ spanned by the vectors

$$\begin{aligned}
A &= (1,1,1,1,\ 0,0,0,\ 1,1,1,1,\ 0,0,0),\\
B &= (0,0,0,0,\ 1,1,1,\ 0,0,0,0,\ 1,1,1),\\
C &= (0,0,1,1,\ 1,0,2,\ 0,0,1,1,\ 1,0,2),\\
C' &= (1,1,0,0,\ 1,2,0,\ 1,1,0,0,\ 1,2,0),\\
D &= (1,1,1,1,\ 0,0,0,\ 0,0,0,0,\ 1,1,1),\\
D' &= (0,0,0,0,\ 1,1,1,\ 1,1,1,1,\ 0,0,0).
\end{aligned}$$

The vectors A, B, C, C', D, D' are the fundamental normal solution vectors for (M_8, Δ). The normal solution space is a 4-dimensional cone spanned by

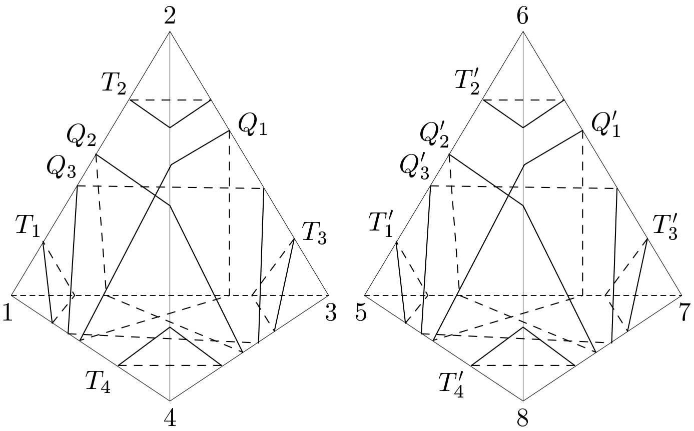

FIGURE 8. Normal disks in the ideal triangulation Δ.

these vectors. Note that the fundamental solution vectors are not linearly independent; for example, $C + C' = A + 2B$ and $D + D' = A + B$.

The fundamental solutions may be interpreted geometrically as follows:

- A consists of one each of all the T-disks. These join together in a unique way to give a link of the vertex, which is a torus. Hence $[A]$ is regular.
- B consists of one each of all the Q-disks. These also join together in a unique way to give a regular gluing, so $[B]$ is regular.
- C and C' are related to the links of the two edges of Δ. They are interchanged by an involution of M which respects Δ. We shall show later that both $[C]$ and $[C']$ contain regular gluings (see Table 1).
- D consists of one each of the T-disks in the first tetrahedron together with one each of the Q-disks in the second tetrahedron. D and D' are related by a second involution of M which respects Δ. Neither $[D]$ nor $[D']$, nor any multiple of these classes, is regular. Each vector nD, for $n = 1, 2, \ldots$, has only T-disks in one tetrahedron and only Q-disks in the other. These normal disks can never be joined together to give an edge disk.

The projective solution space, which we denote $\mathcal{P}$, is the intersection of the normal solution space with the codimension 1 hyperplane with equation

$$x_1 + \cdots + x_{14} = 1.$$

Since the solution space has dimension 4, the projective solution space has dimension 3 and is a solid polyhedron (Figure 9). It is easily seen to be an octahedron; this was first demonstrated in [Jaco 1988]. Since the sum of two points in the projective solution space lies on the line segment joining those points, we see that the line segments joining the pairs A, B and D, D' intersect, as do the line segments joining the pairs A, B and C, C'. Its six vertices are the projective images of the six fundamental solutions. The paper [Matsumoto and Rannard 1997] is concerned with identifying what parts of the projective solution space contain regular surfaces.

The symmetry group of $\mathcal{P}$ is isomorphic to $\mathbb{Z}_2 \oplus \mathbb{Z}_2$, with one generator flipping C and C', the other flipping D and D'. These isometries come from isometries of the figure-eight knot which preserve the triangulation; see [Rolfsen 1976]. There is another symmetry of the knot which maps each solution vector onto itself. These symmetries imply that we need only determine the regularity of vectors which are sums of A, B, C and D, the regularity of the other vectors being apparent by these symmetries. We may also use the fact that the fundamental normal solution vectors are not linearly independent to determine the regularity of other solutions. For instance, $C + C'$ is equal to $A + 2B + C$ and hence is regular.

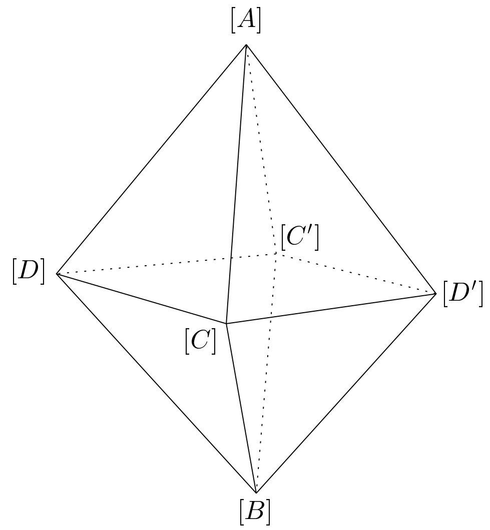

FIGURE 9. The projective solution space of (M_8, Δ).

5. STRATEGIES FOR ENUMERATION OF REGULAR SURFACES

Given a solution vector V, we would like to know:

1. Is V regular? That is, does V represent normal disks which can be joined to give an immersed surface?
2. If V is regular, how many regular surfaces are in $[V]$? Of these surfaces, how many are connected? How many are orientable?
3. If V is regular, does $[V]$ contain surfaces with a given combinatorial property?

We now present various computational methods for answering these questions. The size of the gluing group rises exponentially with the number of normal disks, making the problem nontrivial. We present various different strategies aimed to answer the different questions above.

All the strategies rest on a procedure to determine if a given gluing is regular. This was achieved in our implementation by checking there was an edge disk at each point where the gluing meets the 1-skeleton.

The method is to "rotate" a normal arc about the 1-skeleton at that point, imagining one endpoint fixed and the other sliding across the normal disks, and check that it returns to its original position after passing through the same number of tetrahedra as the order of the edge (in this particular case the order was always 6). This is done for all normal arcs. If a branch point is found at any stage, then the gluing is known to be irregular and checking is terminated.

The strategies below have been encoded in a program written in Magma [Bosma et al. 1997]. This computational algebra system has its own internal programming language, which supports groups and other mathematical objects. The program has been run mainly on a DEC AlphaServer 2100 4/275, having 640MB of memory and 3 CPUs running in parallel at 275MHz, and also on an IBM RISC platform which runs at about half the speed. A good late-model desktop PC can achieve comparable performance.

The program runs in two stages. In the first stage the program uses a particular strategy to search for regular gluings. In the second stage, the program processes the regular gluings in the following ways. First, the set of regular surfaces is extracted from the set of regular gluings by calculating coset representatives of the anonymity subgroup. The orientability and connected components of these representatives are calculated using the obvious procedure, implemented recursively — essentially the same as [Jaco and Tollefson 1995, Algorithm 9.4]. The set of normal disks corresponding to a gluing is traversed, with each being assigned an orientation, and the program checks that this assignment can be made consistently.

5A. First Strategy: Check Every Gluing for Regularity

The simplest and most naïve approach to finding regular surfaces corresponding to a given solution vector is simply to test each gluing in the gluing group in turn for regularity. This gives an exact answer with no uncertainty. This strategy is the only one in the set which is a true algorithm, since it is certain to terminate after a finite time. Moreover, it is the only strategy in the set capable of determining whether a vector is *not* regular.

The practicality of this strategy is limited to very small classes, since the number of gluings rises as the product of the factorial of the number of normal disks of each type. The limit with the computing power available to us is classes containing less than about 10^9 gluings; see Table 1. This strategy is faster the fewer regular gluings are present, since we do not have to fully check irregular gluings.

Optimization is possible in some special cases by factoring the gluing group by the anonymity subgroup before starting the first stage. However, this operation takes nontrivial time, and the number of cosets is still exponential in the norm of the solution vector.

5B. Second Strategy: Opinion Polling

The second strategy is to use a polling method to estimate the number of regular gluings in a given class. The idea is to choose a number of elements at random from the gluing group, and see what proportion of them are regular; we assume roughly the same proportion of the whole group is regular.

The estimate is calculated in the following way. The program samples G gluings at random, finding R of them to be regular. Knowing G_{tot}, the total number of gluings in the class, and A, the size of the anonymity subgroup (the number of gluings which represent each surface), the estimated number of surfaces S is given by

$$S = \frac{R\,G_{\mathrm{tot}}}{A\,G}.$$

A 95% confidence interval for the true number of regular surfaces is $(S - C,\, S + C)$, where C is calculated using the formula

$$C = \frac{D\sqrt{R}}{\sqrt{G}},$$

with

$$D = \frac{1.960(G_{\mathrm{tot}}/A)}{\sqrt{G}};$$

see [Noether 1990].

The formula for D is evaluated separately to reduce roundoff error. The constant 1.960 is determined by our choice of 95% as the confidence level.

The numbers of connected and orientable surfaces are calculated using the same formulas as regular surfaces. Due to the smaller values for R involved, the relative error C/S is usually larger for these numbers.

The polling strategy effectively makes use of the redundancy induced by the way many gluings corresponded to a single surface. It is most efficient when applied to a gluing group with a large anonymity subgroup compared to the size of the gluing group. Examples are the groups corresponding to $3A$, $4A$ and $2C$. The first strategy is not practical for these classes, but accurate results can be obtained efficiently using the opinion-poll strategy. However, for large classes the determination of accurate results takes a prohibitively long time. In addition this strategy is not suitable for checking whether a vector is regular.

5C. Graph-Based Search Strategies

In some situations one needs only to determine if a given class is regular. An important example is given in [Matsumoto and Rannard 1997]. In this case we are searching the set of gluings for one with the desired property of regularity. The graph-based search strategies do this by constructing a weighted graph from the gluing group, and searching it using some known algorithms.

The graph-based search strategies are based on a graph derived from the gluing group in a similar way to the Cayley graph. The vertices of the graph are in one-to-one correspondence with the elements of the gluing group (the gluings for that solution vector). A pair of vertices are joined by an edge if and only if the two gluings differ by a transposition in the gluing group, corresponding to interchanging two parallel normal arcs. Each vertex is given a weight, called its *score*, which measures how far it is from being regular. We choose the score to have the following properties:

- Two vertices connected by an edge have scores differing by at most a fixed small amount;
- The minimum score is zero;
- A gluing is regular if and only if its score is zero.

We now have the task of searching for vertices with minimum weight in a weighted graph, a well-studied problem. We use two related nondeterministic strategies.

We use as the score of a gluing the following nonnegative integer. For each edge of the triangulation Δ, we calculate the number of arc swaps necessary to convert each branch point about that edge to a set of edge disks. The score is the sum of these values over each edge of Δ. This is roughly proportional to the sum of the orders of all the branch points.

Note that when searching all the gluings for a given vector, the time taken is determined by how fast the score of a gluing can be calculated. We want to minimize the number of gluings which we have to examine.

All these strategies are helped by the fact that one surface is represented by many different gluings.

5D. Third Strategy: Hill Climbing

In this strategy we search the graph with a simple nondeterministic depth-first search, keeping track of a "current" vertex of the graph. The process is as follows:

1. Choose a vertex at random to be the initial current vertex.
2. Calculate the scores of all the neighbors of the current vertex.
3. If all these scores are no less than the score of the current vertex, go to Step 1.
4. Otherwise, change the current vertex to be the neighbor with the least score. Stop if this gluing is regular, otherwise go to Step 2.

The process will stop either when a score of 0 is found (meaning the corresponding gluing is regular) or when every neighbor of the current vertex has at least as high a score. In the latter case the routine aborts and begins again from another randomly-selected point; this is the usual result. Note that this strategy is not guaranteed to find a regular gluing should one exist.

This strategy is much faster than the previous two for finding a regular gluing where these exist. Its performance depended on the class, possibly because the proportion of false minima in the graph varied between classes.

The hill-climbing strategy appeared to be effective for the vector $[2B+4D]$, but not for $[2C+4D]$. For the latter, 537 invocations of this search routine (taking up a whole weekend of computer time) failed to uncover a regular surface, although the class was found to be regular using the strategy below. This may be because of "plateau" effects in the graph [Cormen et al. 1990]; there may be regions in the graph where the score is close to constant (the "plateau"), dropping to zero only in small regions about the regular gluings (narrow "gorges"). Unless the (randomly-chosen) starting point is on the side of a gorge, this strategy will be unable to locate the regular gluings.

5E. Fourth Strategy: Bounded Nonincreasing Random Walk

To overcome the problems cause by plateau effects noted in the previous section, we modify our search procedure for the gluing graph. The difference is that we do not necessarily change the current vertex to be the neighbor with the least score at each step. Instead, the procedure sends out a "feeler", exploring a sequence of vertices obtained from the current one by a random walk, and moves to the vertex in this sequence of the lowest score if this is less than the score of the current vertex. The length of the sequence is allowed to grow with the number of tests for regularity since a relocation took place, up to a maximum value. The process is:

1. Choose an vertex at random to be the initial current vertex.
2. Let a variable called the feeler-tip be the current vertex, and a variable called the feeler length be zero.
3. Move the feeler-tip to a neighbouring gluing chosen at random, and increment the feeler length by 1.
4. If the score of the feeler-tip is:
   - zero, then we have found a regular gluing, so abort the search.
   - no greater than the score of the current vertex, let the current vertex be the feeler-tip.
   - greater than the score of the current vertex, do nothing.
5. If the feeler length is greater than a preset value, go to Step 1, else go to Step 2.

The idea behind this strategy is that the feelers will be able to wander over the plateau region of the graph to detect gorges containing regular gluings, even if the current vertex is caught in a local minimum.

Note that no gluing with score greater than that of the initial vertex will ever be reached using this strategy, hence only a restricted set of gluings need be checked for regularity.

6. RESULTS FOR SOME SOLUTION VECTORS

Table 1 summarizes the results obtained from the implementation of the strategies above. Results for other vectors can be obtained from these by applying the symmetries noted in Section 4. The data is invariant under these symmetries, as they give a bijection between the gluing groups associated with the vectors. For example, the number of regular surfaces in $[A + C']$ is the same as the number in $[A+C]$.

We have an example of a regular class which appears to contain no connected immersed surfaces. The class $[2A+2D]$ was searched a total of 4 times. In each case, the regular surface found was disconnected. We conjecture that there is no connected regular surface in this class.

Consider the gluing group G_H associated with a vector

$$H = aA + bB + cC + dD.$$

The vector D is irregular while A, B and C are all regular, so we would expect, very roughly, the proportion of regular gluings in the gluing group of H to decrease as the ratio

$$d : a+b+c$$

increases. The data in Table 1 is consistent with this conjecture, but only when the values of a, b, c, d are "sufficiently large". So, for instance, the gluing group corresponding to $A+A+D$ contains no regular gluings, while $A+A+A$ and $A+D+D$ do; yet doubling these vectors, we can write $4A+2D$ as $(A+2D)+3A$ hence this vector is regular. The moral is that small vectors are atypical; we see a more ordered behaviour when we consider multiples of these vectors. Note that a multiple of a vector projects to the same point of the projective solution space as the original vector.

We have found several examples of solution vectors that are not regular themselves, but some multiple of them is regular, as shown in Table 1.

Definition 6.1. A solution vector is *virtually regular* if some multiple of it is regular. The *regular projective solution space* is the subset of the projective solution space corresponding to all regular and virtually regular solution vectors.

Considering the regular projective solution space instead of the projective solution space "smooths" the data, hiding the anomalies arising from vectors of small norm to reveal the broader distribution of immersed normal surfaces. In [Matsumoto and Rannard 1997] we determine exactly the regular projective solution space for (M_8, Δ), and show that twice any virtually regular vector is regular.

Class	$\log_{10}$ of number of surfaces	gluings	# of regular surfaces		# of connected regular surfaces		# of orientable regular surfaces	
$[A]$	0	0	1		1		1	
$[B]$	0	0	1		1		0	
$[C]$	2.8	2.2	13		13		4	
$[D]$	0	0	0		0		0	
$[A+A]$	3.6	1.2	4		3		3	
$[A+B]$	3.6	3.6	29		28		8	
$[A+C]$	7.1	5.3	59		46		22	
$[A+D]$	3.6	2.4	0		0		0	
$[B+B]$	3.6	1.8	16		15		1	
$[B+C]$	7.1	4.9	54		41		1	
$[B+D]$	3.6	2.7	2		2		0	
$[C+C]$	12.4	7.9	683	(597–770)	586	(505–666)	138	(99–177)
$[C+D]$	7.1	5.1	12		12		0	
$[A+A+A]$	9.3	3.1	3	(2.8–3.2)	1	(1.2–1.4)	3	(2.8–3.2)
$[A+A+B]$	9.3	6.9	141		109		48	
$[A+A+D]$	9.3	5.0	0		0		0	
$[A+B+B]$	9.3	7.5	486		273		62	
$[A+B+D]$	9.3	7.2	40		38		0	
$[A+D+D]$	9.3	5.3	1		1		0	
$[B+B+D]$	9.3	6.1	12		10		0	
$[B+D+D]$	9.3	5.8	0		0		0	
$[D+D+D]$	9.3	3.9	0		0		0	
$[4A]$	16.6	5.5	5	(4.3–5.5)	2	(1.3–2.0)	5	(4.3–5.5)
$[4B]$	16.6	8.3	1427	(1294–1559)	1324	(1197–1452)	16	(2.0–30.1)
$[2B+4D]$	34.3	19.3	some found					
$[2C+4D]$	45.7	26.1	some found					

TABLE 1. Numbers of regular surfaces in various classes. Ranges in parentheses were obtained using the opinion-poll strategy, the range indicating the 95% confidence interval. The other numbers were obtained using the first strategy. On the last two (incomplete) rows, we have found regular surfaces using the third and fourth strategies, but cannot estimate their total number.

7. AN UNUSUAL SURFACE

We conclude by considering one particularly interesting immersed normal surface. A search using the first strategy revealed that there is only one immersed normal surface corresponding to the vector $A+2D$, out of nearly two billion possibilities. A simple Euler characteristic argument shows that this is an orientable surface of genus 2. Moreover, examination of the data revealed that this surface has an interesting combinatorial property: each edge disk of this surface is of the same type. At every point where the surface crosses an edge of Δ, exactly two Q-disks meet, with a single T-disk between them on one side and three T-disks between them on the other; this allows one to continue the development pattern shown in Figure 10.

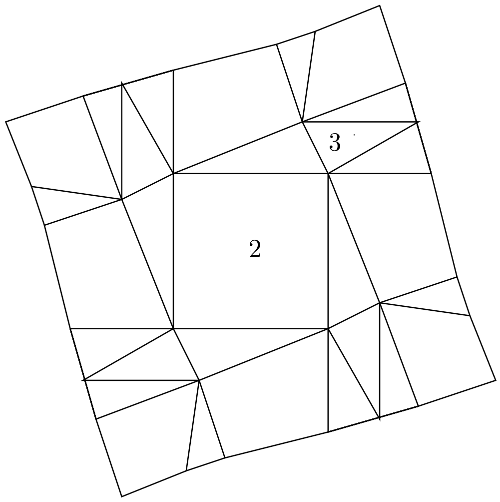

FIGURE 10. Thurston's surface.

This surface has been known to a select few for some years, although it has not yet appeared in the literature. I was told about it by J. H. Rubinstein and I. Aitchison, who credited W. Thurston with its discovery, and who later showed it to be incompressible [Rubinstein 1998].

In [Matsumoto and Rannard 1997] the authors show that this exceptional surface is a vertex of the regular projective solution space.

ACKNOWLEDGMENTS

I thank J. Hyam Rubinstein for his excellent advice and support. Thanks also to Saburo Matsumoto for a very enjoyable collaboration while researching this material, and for allowing me to use Figure 8, which was drawn by him.

REFERENCES

[Aitchison et al. 1997] I. Aitchison, S. Matsumoto, and J. H. Rubinstein, "Surfaces in the Figure 8 knot complement", preprint, University of Melbourne, 1997.

[Bosma et al. 1997] W. Bosma, J. Cannon, and C. Playoust, "The Magma algebra system, I: The user language", *J. Symbolic Comput.* **24**:3-4 (1997), 235–265. See http://www.maths.usyd.edu.au:8000/comp/magma/Overview.html.

[Cormen et al. 1990] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to algorithms*, MIT Press, Cambridge, MA, 1990.

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

[Hemion 1992] G. Hemion, *The classification of knots and* 3*-dimensional spaces*, Oxford University Press, New York, 1992.

[Jaco 1988] W. Jaco, 1988. Unpublished personal notes.

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

[Jaco and Rubinstein 1989] W. Jaco and J. H. Rubinstein, "PL equivariant surgery and invariant decompositions of 3-manifolds", *Adv. in Math.* **73**:2 (1989), 149–191.

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

[Letscher 1997] D. Letscher, Ph.D. thesis, University of Michigan, 1997.

[Matsumoto and Rannard 1997] S. Matsumoto and R. Rannard, "Immersed normal surfaces in the figure 8 knot complement II", preprint, 1997. Submitted to *Experimental Mathematics* under the title "On the regular projective solution space of the figure-eight knot complement".

[Noether 1990] G. E. Noether, *Introduction to statistics: the non-parametric way*, Springer series in statistics, Springer, New York, 1990.

[Rannard 1997] R. Rannard, *Normal surface theory and the topology of* 3*-manifolds*, Ph.D. thesis, University of Sydney, 1997.

[Rolfsen 1976] D. Rolfsen, *Knots and links*, Mathematics Lecture Series **7**, Publish or Perish, Berkeley, 1976. Reprinted with corrections, Publish or Perish, Houston, 1990.

[Rubinstein 1998] J. H. Rubinstein, June 1998. Personal communication.

[Thurston 1997] W. P. Thurston, *Three-dimensional geometry and topology,* 1, Princeton Math. Series **35**, Princeton University Press, Princeton, NJ, 1997.

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

Richard Rannard, School of Mathematics and Statistics, University of Sydney, Sydney, New South Wales 2006, Australia (richardr@maths.usyd.edu.au)
Current address: P.O. Box Q311, QVB, Sydney 1230 Australia (richard.rannard@etc.com.au)

Received January 16, 1997; accepted in revised form June 16, 1998