3-MANIFOLD KNOT GENUS is NP-complete *

Ian Agol†
Department of Mathematics
University of Illinois at Chicago
Chicago, IL 60607-7045
agol@math.uic.edu

Joel Hass‡
Department of Mathematics
University of California
Davis, CA 95616
hass@math.ucdavis.edu

William Thurston§
Department of Mathematics
University of California
Davis, CA 95616
wpt@math.ucdavis.edu

ABSTRACT

We show that the problem of deciding whether a polygonal knot in a closed three-dimensional manifold bounds a surface of genus at most g is **NP**-complete.

Categories and Subject Descriptors

F.2.2 [**Analysis of algorithms and problem complexity**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Algorithms

Keywords

Computational complexity, knot, 3-manifold, **NP**-complete, normal surface, genus

1. INTRODUCTION

Determining whether a given knot is trivial or not is historically one of the central questions in topology. The problem of finding an algorithm to determine knot triviality was posed by Dehn [2]. Dehn's investigations into this area led to the formulation of the word and isomorphism problems, which played an important role in the development of the theory of algorithms. The first algorithm for UNKNOTTING was given by Haken [4]. Haken's procedure is based on normal surface theory, a method of representing surfaces originated by Kneser [9]. Analysis of the computational complexity of this algorithm is more recent. Hass, Lagarias and Pippenger showed that Haken's unknotting algorithm runs in time at most c^t, where the knot K is embedded in the 1-skeleton of a triangulated manifold M with t tetrahedra, and c is a constant independent of M or K [5]. It was also shown in [5] that UNKNOTTING is **NP**.

Any oriented surface without boundary can be obtained from a sphere by adding "handles". The number of handles is the *genus* of the surface: a sphere has genus 0, a torus has genus 1, etc. The genus is also equal to the number of curves along which a surface can be cut while leaving it connected. If a surface has boundary, its genus is defined as the genus of the surface gotten by gluing a disk to each component of the boundary. The genus of a surface can be computed using any subdivision into polygons, from the formula

$$\text{Euler characteristic} =$$
$$\#\text{vertices} - \#\text{edges} + \#\text{faces} =$$
$$2 - 2 \times \text{genus} - \#\text{boundary curves}$$

UNKNOTTING is a special case of the more general problem 3-MANIFOLD KNOT GENUS. This problem asks for a procedure to determine the *knot genus*, the minimal genus of an orientable spanning surface for a knot in a 3-dimensional manifold. A knot is trivial, or unknotted, precisely when its genus is zero. We will show that this more general problem is **NP**-complete. The only previous results obtained on this problem were given in [5], where it was shown to lie in **PSPACE**. No lower bounds on the running time were previously known.

We establish that 3-MANIFOLD KNOT GENUS is **NP**-hard by giving a new topological construction that transforms an instance of ONE-IN-THREE SAT, a known **NP**-complete problem, to an instance of 3-MANIFOLD KNOT GENUS. To show it is **NP**, we give a certificate that describes, using normal surface coordinates, a surface of minimal genus among the surfaces spanning the knot. In many ways this argument follows that of [5]. However previous methods were unable to effectively check whether a normal surface of given Euler characteristic was connected. Connectedness, rather surprisingly, turns out to be the most difficult feature to certify. The problem was avoided, in the unknotting algorithm looked at in [5], by applying a result of Jaco-Oertel that shows that unknotting disks can be found on the extremal rays of the Hilbert basis for the space of normal surfaces [6]. Connectedness can be quickly certified for extremal solutions. However the existence of extremal solutions corresponding to minimal genus spanning surfaces is not known. This problem is overcome through the intro-

*A full version of this paper will be available as *The Complexity of Knot Genus.*

†Partially supported by ARC grant 420998.

‡This work was carried out during a visit to the Institute for Advanced Study. Partially supported by NSF grant DMS-0072348.

§Partially supported by NSF grant DMS-9704286.

duction of a new algorithm that counts the number of components of a normal surface representing a minimal genus spanning surface in polynomial time. This algorithm is developed in a much more general context, as an algorithm to count the number of orbits of a collection of pairings acting on an interval, and seems likely to find other applications.

A *knot* in a 3-manifold M is a connected simple (non self-intersecting) closed curve in the 1-skeleton of M. Complexity of the input is measured by the number of tetrahedra used to construct the triangulation of M. Knot genus was defined by Seifert [13] in 1935 for knots in the 3-sphere, and extends directly to knots in all 3-manifolds. Given a knot K, consider the class $\mathcal{S}(K)$ of all orientable spanning surfaces for K. These are surfaces embedded in M with a single boundary component that coincides with K. The genus $g(K)$ of a knot K is the minimum genus of a surface in $\mathcal{S}(K)$, or ∞ if $\mathcal{S}(K) = \emptyset$. Genus measures one aspect of the degree of "knottedness" of a curve.

We formulate the problem of computing the genus as a language-recognition problem in the usual way, see [3]. Schubert [12] extended Haken's work to show the decidability of the problem:

Problem: 3-MANIFOLD KNOT GENUS
INSTANCE: A triangulated 3-dimensional manifold M,
a knot K in the 1-skeleton of M,
and a natural number g.
QUESTION: Does the knot K have $g(K) \leq g$?

The size of an instance is given by the sum of the number of tetrahedra t in M and $\log g$. Our main result is the following.

THEOREM 1.1. *3-MANIFOLD KNOT GENUS is* **NP**-*complete.*

Knots are often studied in $\mathbb{R}^3$ or S^3 rather than in a general manifold. Our methods show that determining knot genus in $\mathbb{R}^3$ or S^3, or any fixed manifold, is **NP**. It is not clear whether the problem remains **NP**-hard in $\mathbb{R}^3$ or S^3 or any fixed 3-manifold.

2. NORMAL SURFACES

General surfaces in 3-manifolds can wind and twist around the manifold in complicated ways. But Kneser showed that surfaces can be "pulled taut" until they take a simple and rigid position [9]. In this *normal* position, they have very succinct algebraic descriptions. We use an approach to normal surfaces based on work of Jaco and Rubinstein [7] and Jaco and Tollefson [8]. A *normal surface* S in a triangulated compact 3-manifold M is a PL-surface whose intersection with each tetrahedron in M consists of a disjoint set of *elementary disks*. These are either triangles or quadrilaterals. See Figure 1.

Within each tetrahedron of M there are four possible triangles and three possible quadrilaterals, up to isotopies that map the tetrahedron to itself. The triangles separate one vertex of the tetrahedron from the other three, and the quadrilaterals separate the the vertices into pairs. W. Haken observed that a normal surface is determined up to such isotopies by the number of pieces of each of the seven kinds of elementary disks that occur in each tetrahedron, or a vector in $\mathbb{Z}_+^{7t}$. A normal surface is described by a non-negative integer vector $\mathbf{v} = \mathbf{v}(S) \in \mathbb{Z}^{7t}$, that gives the *normal coordinates* of S. The set of allowable values $\mathbf{v}(S)$ for a normal surface S lie in a certain homogeneous rational cone $\mathcal{C}_M$ in $\mathbb{R}^{7t}$, called the *Haken normal cone.* If $\mathbf{v} = (v_1, v_2, \dots, v_{7t}) \in \mathbb{R}^{7t}$, then the Haken normal cone is specified by linear equations and inequalities of the form

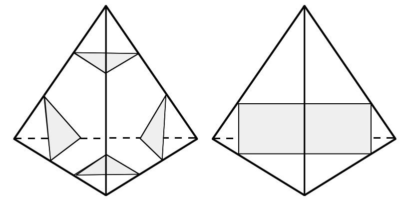

Figure 1: Elementary disks of a normal surface in a tetrahedron of a triangulated manifold.

$$v_{i_1} + v_{i_2} = v_{i_3} + v_{i_4} \quad \text{(up to } 6t \text{ equations)} ,$$

$$v_i \geq 0 , \quad \text{for } \ 1 \leq i \leq 7t .$$

The first set of equations expresses *matching conditions,* which say that the number of edges on a common triangular face of two adjacent tetrahedra, coming from a collection of elementary disks in each of the tetrahedra, must match. For each triangular face there are three types of edges (specified by a pair of edges on the triangle), which yield 3 matching conditions per face. Triangular faces in the boundary ∂M give no matching equations. The cone $\mathcal{C}_M$ is *rational,* because the above equations have integer coefficients. We let $\mathcal{C}_M(\mathbb{Z}) = \mathcal{C}_M \cap \mathbb{Z}^{7t}$ denote the set of integral vectors in the cone $\mathcal{C}_M$. An additional set of conditions, the *quadrilateral conditions,* is required for an integral vector in the cone to correspond to the normal coordinates of an embedded surface. This condition states that of the three types of quadrilateral found in each tetrahedron, only one can occur in the vector with non-zero coefficient. The quadrilateral conditions are required because two distinct types of quadrilateral in a single tetrahedron necessarily intersect, and we are interested in embedded, non-self-intersecting surfaces. A vector in the Haken normal cone that satisfies the quadrilateral conditions corresponds to an embedded normal surface. This surface is unique up to an isotopy fixing each tetrahedron (called a *normal isotopy*). However it is important to note that the surface may not be connected. A *fundamental normal surface* is a normal surface S such that

$$\mathbf{v}(S) \neq \mathbf{v}_1 + \mathbf{v}_2 , \quad \text{with } \mathbf{v}_1, \mathbf{v}_2 \in \mathcal{C}_M(\mathbb{Z}) \backslash \{\mathbf{0}\} .$$

In the terminology of integer programming, such a vector $\mathbf{v}(S)$ is an element of the *minimal Hilbert basis* $\mathbb{H}(\mathcal{C}_M)$ of $\mathcal{C}_M$, see Schrijver [11, Theorem 16.4]. A fundamental normal surface is always connected, but connected normal surfaces need not be fundamental. Hass, Lagarias and Pippenger [5] gave a bound for the complexity of any fundamental surface.

THEOREM 2.1. *Let M be a triangulated compact 3-manifold, possibly with boundary, that contains t tetrahedra.*

- *Any vertex minimal solution $\mathbf{v} \in \mathbb{Z}^{7t}$ of the Haken normal cone C_M in $\mathbb{R}^{7t}$ has* $\max_{1 \leq i \leq 7t} (v_i) \leq 2^{7t-1}$.
- *Any minimal Hilbert basis element $\mathbf{v} \in \mathbb{Z}^{7t}$ of the Haken fundamental cone C_M has* $\max_{1 \leq i \leq 7t} (v_i) < t \cdot 2^{7t+2}$.

Schubert [12] showed that a surface of smallest genus spanning K can be found among the fundamental surfaces.

3. 3-MANIFOLD KNOT GENUS IS NP-HARD

In this section we indicate how to reduce an instance of ONE-IN-THREE SAT to an instance of 3-MANIFOLD KNOT GENUS.

Problem: ONE-IN-THREE SAT
INSTANCE: A set U of variables and a collection C of clauses over U such that each clause $c \in C$ has $|c| = 3$.
QUESTION: Is there a truth assignment for U such that each clause in C has exactly one true literal?

Schaefer [10] established that ONE-IN-THREE SAT is **NP**-complete.

THEOREM 3.1. *3-MANIFOLD KNOT GENUS is* **NP**-*hard.*

PROOF. Let $U = \{u_1, u_2, \ldots u_n\}$ be a set of n variables and $C = \{c_1, c_2, \ldots c_m\}$ be a set of m clauses in an arbitrary instance of ONE-IN-THREE SAT. We construct a knot K in a compact 3-dimensional manifold (with no boundary) and an integer g such that K bounds a surface of genus smaller or equal to g if and only if C is satisfiable with each clause in C having exactly one true literal. We construct the 3-manifold M in stages. First we construct a 2-dimensional simplicial complex, then we thicken this complex, replacing triangles with subdivided triangular prisms, getting a triangulated, 3-dimensional manifold with boundary, as indicated in Figure 2. Finally we use a doubling construction, taking two copies of the manifold with boundary and gluing their boundaries together, to obtain a closed 3-manifold.

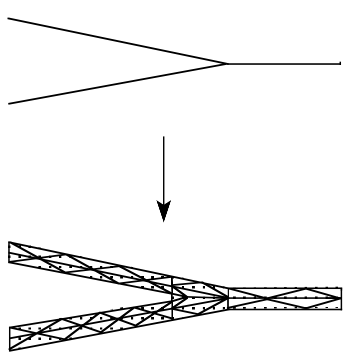

Figure 2: The construction begins with a branched surface B. This surface is then thickened to produce a triangulated 3-manifold with boundary. A cross section of this thickening is shown.

To begin, we construct a "branched" surface B by identifying boundary curves of a collection of $2n+1$ surfaces with boundary. Let k_i be the number of times that the variable u_i appears in the collection of clauses, and $\bar{k}_i$ be the number of times that $\bar{u}_i$ appears. For $i = 1, \ldots, n$, define the surfaces F_{u_i} and $F_{\bar{u}_i}$ to be genus one surfaces with $k_i + 1$ and $\bar{k}_i + 1$ boundary curves respectively. Set F_0 to be a planar surface with $n + m + 1$ boundary curves, one of which will be the knot K. The branched surface B is constructed by identifying these surfaces along appropriate boundary components as indicated in Figure 3. Branching occurs when more than two boundary curves are identified along a single curve.

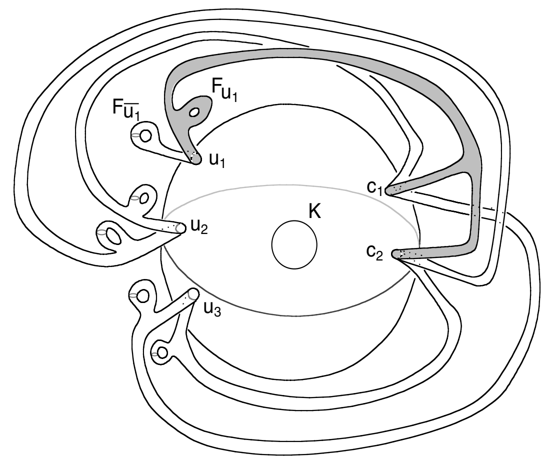

Figure 3: A branched surface with boundary curve K corresponding to the boolean expression $(u_1 \vee u_2 \vee u_3) \wedge (u_1 \vee \bar{u}_2 \vee \bar{u}_3)$. The shaded surface F_{u_1} indicates the occurrence of u_1 in each of the clauses c_1 and c_2.

Label by $K, u_1, \ldots u_n, c_1, \ldots c_m$ the $1 + n + m$ boundary components of F_0. The boundary component K of F_0 has nothing identified to it. For each i, $1 \leq i \leq n$, one boundary curve from surface F_{u_i} is identified to u_i. The remaining k_i boundary components of F_{u_i} are identified with k_i of the curves $c_1, \ldots c_m$ on ∂F_0, with one component of ∂F_{u_i} identified with c_j for each occurrence of the literal u_i in the j^{th} clause in C. Similarly, one curve of $\partial F_{\bar{u}_i}$ is identified to u_i, and the remaining $\bar{k}_i$ boundary components are glued to $c_1, \ldots c_m$, with a component glued to c_j for each occurrence of the literal $\bar{u}_i$ in the j^{th} clause of C. A total of three surface boundaries are identified along each of $u_1, \ldots, u_n$, and exactly four surface boundaries are identified along each of $c_1, \ldots, c_m$, as in Figure 3.

LEMMA 3.1. *There is a truth assignment for U such that each clause in C has exactly one true literal if and only if there is a surface of genus at most $m + n$ mapped continuously into B with boundary K.*

PROOF. Suppose there is a truth assignment for U such that each clause in C contains exactly one true literal. Form a surface S inside B by taking the union of F_0 and either F_{u_i} if u_i is true, or $F_{\bar{u}_i}$ if u_i is false. Then exactly two boundary components will be identified along each of the boundary components of F_0 other than K itself, and K becomes the boundary of the resulting embedded surface S. There is a contribution of one to the genus from each of the literals, since F_{u_i} and $F_{\bar{u}_i}$ have genus one, and a contribution of one to the genus from each handle formed when a boundary component of a surface F_{u_i} or $F_{\bar{u}_i}$ is glued to F_0 along a curve c_j. The genus of S is therefore equal to $m + n$.

Now suppose there is a surface S_1 of genus $\leq n + m$ mapped continuously into B that has a single boundary component mapped homeomorphically to the boundary curve K. We will show in this case that there is a surface S_3 with the same boundary having genus precisely $n+m$, and containing, along each curve c_j, exactly one of the three pieces of surface joining F_0.

The image of S_1 may be quite complicated, winding back and forth across B, but by standard transversality arguments we can homotop S_1 so that it is formed from a union of surface pieces, each piece being mapped homeomorphically to one of the subsurfaces F_0, F_{u_i}, $F_{\bar{u}_i}$ forming B. The image of S_1 in each subsurface has a degree, which is the number of times that an interior point of that subsurface is hit by the pieces of S_1. The degree of S_1 on F_0 is odd, since K is the boundary of S_1 and therefore gets hit an odd number of times. The sum of the degrees along each of the surfaces meeting a curve u_i or c_j is even, since these are not boundary curves of S_1. So for each $1 \leq i \leq n$, exactly one of F_{u_i}, $F_{\bar{u}_i}$ has odd multiplicity in S_1. Form a new surface S_2 by taking the union of F_0 and each of the surfaces F_{u_i}, $F_{\bar{u}_i}$ which have odd multiplicity in S_1. Since S_2 is obtained from S_1 by discarding a collection of surfaces with non-positive Euler characteristic, we have $\chi(S_2) \geq \chi(S_1)$. The union of the surfaces F_{u_i}, $F_{\bar{u}_i}$ contained in S_2 can be attached to F_0 along any common boundary curves among $\{u_1, \ldots u_n\}$, giving a connected surface of genus n, whose Euler characteristic equals that of S_2, since gluing along circles preserves Euler characteristic. Finally, an embedded connected surface S_3 with boundary K is obtained by identifying pairs of curves that are mapped to the boundary components $c_1, \ldots c_m$. The Euler characteristic of S_3 is the same as that of S_2, and therefore greater or equal to that of S_1. So the genus of S_3 is at most that of S_1, $\text{genus}(S_3) \leq n+m$. Each identification of a pair of surface boundaries along $c_1, \ldots c_m$ contributes one to the genus of S_3. There is at least one such identification along each of the m curves $\{c_j\}$ and there may be two, if each of the three surfaces meeting F_0 along c_j has multiplicity one in S_2. So identification of curves along $c_1, \ldots c_m$ contributes at least m to the genus of S_3. It follows that $\text{genus}(S_3) \geq n+m$. Since we have seen that $\text{genus}(S_3) \leq n+m$, equality must hold. Equality holds when exactly one of the three surface pieces meeting F_0 along c_j has odd multiplicity for each $1 \leq j \leq m$. We then assign the value "TRUE" to a literal u_i if F_{u_i} is used in S_3, and the value "FALSE" to u_i if $F_{\bar{u}_i}$ is used in S_3. This gives a truth assignment to U in which each clause in C has exactly one true literal, as claimed. □

We now show that 3-MANIFOLD KNOT GENUS is **NP**-hard.

PROOF. Given an instance of ONE-IN-THREE SAT, form B as in Lemma 3.1. Form a 3-manifold M_1 by thickening B so that it is embedded inside a 3-dimensional manifold with boundary N. The curve K remains on the boundary of M_1. Form a closed manifold M by doubling N along its boundary, gluing two copies of the manifold along their boundaries to obtain a closed 3-manifold.

It is straightforward to check that if K bounds a surface of genus $m+n$ with interior in $M-K$ then it bounds a surface of the same genus mapped into B. The constructions of B, N and M described above each requires only a linear number of steps in the size of the initial instance of ONE-IN-THREE SAT, so that the reduction can be done in polynomial time. □

4. ORBITS OF INTERVAL ISOMETRIES

In this section we develop a combinatorial procedure that computes the number of orbits of a collection of k isometries between subintervals of an interval $[1, N] \subset \mathbb{Z}$ in time polynomial in $k \log N$. By lining up the intersections of a normal surface with the edges of a triangulation, we obtain such subintervals. Arcs of the normal surface on the faces of the triangulation give rise to correspondences of these intersection points which are subinterval isometries. We then apply the algorithm we develop here to count the number of components of a normal surface.

Assume that we have a set of integers $\{1, 2, \ldots N\}$ and a collection of monotone bijections, $g_i : [a_i, b_i] \to [c_i, d_i]$, $1 \leq i \leq k$, $a_i \leq c_i$, either increasing or decreasing, that are called *pairings*. We are interested in the orbit structure of the collection, since with appropriate interpretation an orbit corresponds to a connected component of a normal surface. If a pairing identifies two intervals $[a, b]$ and $[c, d]$ of equal width by sending a to c and b to d, we call it an *orientation preserving* pairing, and if it sends a to d and b to c, we call it *orientation reversing*. We can compose two pairings if the range of the first lies in the domain of the second. The collection of pairings generates a pseudo-group, under the operations of composition where defined, inverses, and restriction to subintervals.

We introduce several simplification processes on the set of pairings. Details will appear in [1].

A pairing $g : [a, b] \to [c, d]$ with overlapping domain and range is called *periodic*, with *period* or *translation distance* $t = c - a$. The combined interval $[a, d]$ is called a *periodic interval* of period t.

LEMMA 4.1. *When an orientation preserving pairing $g : [a, b] \to [c, d]$ has $a < c \leq b+1$, then the pairing has $t = c-a$ orbits on $[a, d]$.*

An operation called a *periodic merger* replaces two periodic pairings g_1 and g_2 that have sufficiently overlapping domains and ranges by a single periodic action on the union of their domains and ranges.

A more detailed description of when a merger operation can be made is given in the following.

LEMMA 4.2. *Let g_1, g_2 be periodic pairings with periods $t_1 = c_1 - a_1, t_2 = c_2 - a_2$. Suppose that J_1, J, J_2 are intervals with $g_1(J_1) = J$ and $g_2(J) = J_2$ and that J is contained in the union $J_1 \cup J_2$. Then the orbits of $g_1 \cup g_2$ are the same as those of a single periodic pairing of period* $\text{GCD}(t_1, t_2)$, *acting on the union of the periodic intervals of g_1, g_2.*

The proof consists of an examination of the orbit structure of overlapping periodic intervals.

The operation of *contraction* is performed on an interval $[r, s]$ whose points are identified to no other points by pairings. The interval is removed and pairings containing it are adjusted accordingly.

The next operation simplifies an orientation reversing pairing whose domain and range overlap. If $g : [a, b] \to [c, d]$ is a pairing with $g(a) = d, g(b) = c$ and $b \geq c$, define a new pairing $g' : [a, (a+d)/2) \to ((a+d)/2, d]$ by restricting the domain and range of g, and say that g' is obtained from g by *trimming*. The domain and range of a trimmed pairing are disjoint.

If there is an interval in the domain and range of exactly one pairing, then the interval can be peeled off without changing the orbit structure. We restrict this operation to strip off points from the right of the interval $[1, N]$. When there is a pairing $g : [a, b] \to [c, N]$ and a value N' with $c \leq N'+1 \leq N$, such that all points in the interval $[N'+1, N]$ are in the domain or range of no pairing other than g, then we can perform an operation called *truncation* of g. Truncation shortens the interval $[1, N]$ to the interval $[1, N']$, and similarly truncates the domain and range of g.

Define a *maximal pairing* to be one whose range contains N' and the range of any other pairing containing N'. More precisely, we use a linear order on pairings coming from lexicographical order of $(d_i, -c_i, -a_i, -\text{orientation})$, so that the maximal pairing has the highest upper endpoint, and among those the widest range, and among those the biggest translation distance (if orientable).

Our final operation is called *transmission*. In a transmission two pairings are composed. A maximal pairing g_1 is used to shift down, as much as possible, the domain and range of a second pairing g_2.

The precise construction of a transmission is the following. If the domain of g_2 is not contained in the range of g_1, form the composite map $g_2' = g_1^{-r} \circ g_2$, where $r = 1$ if g_1 is orientation reversing and otherwise $r \geq 1$ is the largest integer such that $g_1^{-r+1}([c_2, d_2])$ is contained in the range of g_1. If the domain of g_2 is also contained in the range of g_1 then form the composite map $g_2' = g_1^{-r} \circ g_2 \circ g_1^s : g_1^{-s}([a_2, b_2]) \to g_1^{-r}([c_2, d_2])$, where r is as above, $s = 1$ if g_1 is orientation reversing, and otherwise $s \geq 1$ is the largest integer such that $g_1^{-s+1}([a_2, b_2])$ is contained in the range of g_1.

Lemma 4.3. *The operations of periodic merger, trimming, truncation and transmission preserve the number of orbits of a collection of pairings. A contraction decreases the number of orbits by the width of the contracted interval.*

We now sketch the algorithm to count orbits. We initially set a counter for the number of orbits to zero, and denote by N' the current interval size. The orbit counting algorithm reduces N' and the number of bijections until none remain. Details will appear in [1], along with a more general version of the algorithm that counts weighted orbits.

Orbit counting algorithm:

1. Delete any pairings that are restrictions of the identity.
2. Make any possible contractions. Increment the orbit count.
3. Trim all orientation reversing pairings whose domain and range overlap.
4. Find a maximal pairing g_i. Recall that its range contains N' and the range of any other pairing containing N'. If g_i is periodic, search for other periodic pairings g_j that have range $[c_j, N']$ containing N'. If such a g_j exists with sufficient overlap, perform a merger, replacing g_i and g_j by a single periodic pairing with period $\mathrm{GCD}(t_i, t_j)$ acting on the union of the domains and ranges of g_i and g_j.
5. Find a maximal g_i. For each g_j with $j \neq i$, if the range of g_j is contained in $[c_i, N']$, transmit g_j by g_i.
6. Find the smallest value of c such that the interval $[c, N']$ intersects the domain and range of at most one pairing. Truncate the pairing whose range contains the interval $[c, N']$.
7. If the interval size N' has decreased to zero, output the number of orbits and stop. Otherwise start again with Step (1).

Theorem 4.1. *The orbit counting algorithm gives the number of orbits of the action of the pairings $\{g_i\}_{i=1}^k$ on $[1, \ldots, N]$ in time bounded by a polynomial in $k \log N$.*

The proof proceeds by showing that each cycle through the steps of the algorithm reduces a certain integer complexity X. Set $w_i = b_i - a_i + 1$ to be the width of a pairing $g_i : [a_i, b_i] \to [c_i, d_i]$, $1 \leq i \leq k$, where k is the number of pairings. The complexity X is defined to be

$$X = 4^k \prod_{i=1}^{k} w_i$$

An analysis shows that X decreases by a factor of at least two after $3k$ cycles. Since each cycle runs in time polynomial in $k \log N$, so does the whole algorithm [1].

Corollary 4.1. *Let M be a 3-manifold with a triangulation T containing t tetrahedra and let F be a normal surface in M whose normal coordinates are each at most W. There is a procedure for counting the number of components of F that runs in time polynomial in $t \log W$.*

Proof. The 2-skeleton $T^{(2)}$ of T contains at most $4t$ faces and the 1-skeleton $T^{(1)}$ contains $e \leq 6t$ edges. Fix once and for all an ordering of the edges. A normal surface F intersects each edge in a finite number of points. Let N be the total number of intersections points. Then $N \leq eW \leq 6tW$. Label the intersections of F and the first edge of the 1-skeleton by $1, 2, \ldots, i_1$, and the intersection of F and the j^{th} edge of F by $i_{j-1}+1, i_{j-1}+2, \ldots, i_j$, with $i_0 = 0$ and $i_e = N$. To a pair of edges on a triangular face of the 2-skeleton we associate a pairing between the intervals at either end of the corresponding set of arcs. There are at most three pairings for each face, and the number of faces is at most $4t$, so the number of pairings that results is bounded by $12t$. These pairings are determined by the normal coordinates of F, as indicated in Figure 4.

The number of connected components of F is the same as that of $F \cap T^{(2)}$. This is the number of orbits of $F \cap T^{(1)}$, under the equivalence where two points are identified if they are connected by an edge of $F \cap T^{(2)}$ contained in a face of the 2-skeleton, and this is the number returned by the orbit counting algorithm. So in time polynomial in $t \log W$ we can determine the number of components of the normal surface F. $\square$

Proof that 3-MANIFOLD KNOT GENUS **is NP:** Suppose that we have a simplicial complex that triangulates a 3-manifold M. Form the second barycentric subdivision of the triangulation of M, replacing each tetrahedron by 576 tetrahedra. Remove all tetrahedra intersecting K, giving a 3-manifold with a torus boundary component, M_K. The knot K has genus g in M if and only if there is a surface in M_K of genus g that has a single boundary component that

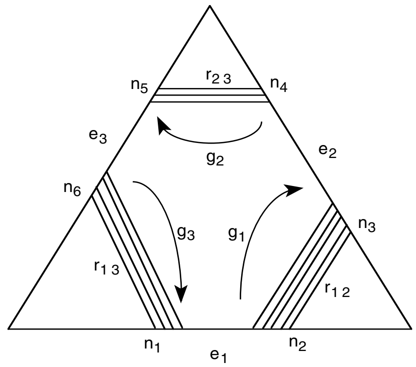

Figure 4: Normal arcs on a triangular face give three pairings of ordered intervals: $g_1 : [n_1 + r_{13}, n_2] \to [n_3, n_3 + r_{12} - 1]$**,** $g_2 : [n_3 + r_{12}, n_4] \to [n_5, n_5 + r_{23} - 1]$**, and** $g_3 : [n_5 + r_{23}, n_6] \to [n_1, n_1 + r_{13} - 1]$ **The normal coordinates associated to this face are** r_{12}, r_{13} **and** r_{23}**.**

is an essential curve on ∂M_K. We can restate our problem in the triangulated manifold M_K as the question of whether there exists in M_K an orientable surface of genus $\leq g$ with a single essential boundary component on ∂M_K. the "peripheral torus" that surrounds the knot K. By Schubert [12], if such a surface exists, then it can be found among the fundamental normal surfaces in M_K. The certificate consists of an integer vector $\mathbf{w}$ in $\mathbb{Z}_+^{7t}$ giving the normal surface coordinates of this surface.

Recall that not all normal vectors correspond to normal surfaces. It is necessary that $\mathbf{w}$ satisfies the quadrilateral conditions. These state that of each triple of coordinates corresponding to three quadrilateral types in each tetrahedron, at least two must equal zero. This condition can be checked in time linear in t for the coordinates of $\mathbf{w}$. So there is a normal surface F with $\mathbf{w} = \mathbf{v}(F)$.

To check that ∂F is essential, we find a non-trivial cycle in the 1-skeleton of ∂M_K that intersects ∂F in an odd number of points. This implies that ∂F represents a non-trivial element in $H_1(\partial M_K; Z_2)$ and is therefore essential, ensuring that ∂F corresponds to a curve running once around K. Let N be the total number of points at which F meets the 1-skeleton of M_K. Using the orbit counting algorithm, we can count the number of components of the boundary of the normal surface F and verify that it is connected in time polynomial in $t \log N$. Similarly, we can apply Corollary 4.1 to verify that F is connected in time polynomial in $t \log N$.

To check that F is orientable, we take the normal vector $2\mathbf{v}(F)$, that doubles each coordinate of F, and apply the orbit counting algorithm to determine if it is connected. Since M is orientable, $2\mathbf{v}(F)$ will be connected precisely when F is non-orientable, and otherwise have two components.

The Euler Characteristic of F is determined by the number of vertices, edges and faces of F, which are easily derivable from its normal coordinates [8]. Since the genus of a connected orientable surface with one boundary component and Euler characteristic χ is equal to $(1-\chi)/2$, we can determine the genus of F in time polynomial in $t \log N$. Theorem 2.1 gives that the normal coordinates of this surface are at most $t2^{7t+2}$. There are $7t$ normal coordinates, and each represents a triangle or quadrilateral, so that the total number of intersections with the 1-skeleton satisfies $N \leq 28t^2 2^{7t+2}$. In particular, $\log N$ is bounded above by a polynomial in t. So it can be verified in time polynomial in t that F is a spanning surface for K.

5. REFERENCES

[1] I. Agol, J. Hass, and W. Thurston. The complexity of knot genus. *in preparation*, 2002.

[2] M. Dehn. Uber die topologie des dreidimensional raumes. *Math. Annalen*, 69:137–168, 1910.

[3] M. Garey and D. Johnson. *Computers and intractability. A guide to the theory of* **NP**-*completeness*. W. H. Freeman and Co., San Francisco, 1979.

[4] W. Haken. Theorie der normalflachen: Ein isotopiekriterium fur den kreisknoten. *Acta Math.*, 105:245–375, 1961.

[5] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *Journal of the ACM*, 46:185–211, 1999.

[6] W. Jaco and U. Oertel. An algorithm to decide if a 3-manifold is a haken manifold. *Topology*, 23:195–209, 1984.

[7] W. Jaco and J. H. Rubinstein. Pl equivariant surgery and invariant decompositions of 3-manifolds. *Advances in Math.*, 73:149–191, 1989.

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

[9] H. Kneser. Geschlossene flachen in dreidimensionalen mannigfaltigkeiten. *Jahresbericht Math. Verein.*, 28:248–260, 1929.

[10] T. Schaefer. The complexity of satisfiability problems. In *Proc 10th Ann. ACM. Symp. on Theory of Computing*, pages 216–226. ACM, 1978.

[11] A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley and Sons, Chichester, 1986.

[12] H. Schubert. Bestimmung der primfaktorzerlegung von verkettungen. *Math. Zeitschr.*, 76:116–148, 1961.

[13] H. Seifert. Uber das geschlecht von knoten. *Math. Annalen*, 76:571–592, 1961.