

INJECTIVELY IMMERSED TORI IN BRANCHED COVERS OVER THE FIGURE EIGHT KNOT

KERRY N. JONES

Section 0 - Introduction.

The purpose of this paper is to prove the following theorem and to describe the algorithm in question:

Theorem 6.1. *Let M be a branched cover of S^3 , branched over the figure eight knot, such that all branching indices of M are ≥ 3 . Then, there exists an effective algorithm for deciding whether or not M admits any π_1 -injectively immersed tori, and for constructing any that exist.*

The importance of this result lies primarily in the fact [HLM] that the figure eight knot is *universal*. That is, that all closed, orientable 3-manifolds are obtainable as branched covers over S^3 , branched over the knot. We also obtain results about incompressible (embedded) tori in some cases. Existence or nonexistence of injectively immersed tori is critical to the understanding of a 3-manifold in light of the Jaco-Shalen and Johannson torus decomposition [J-S,Jh] which plays such a pivotal role in Thurston's geometrization conjecture and the recent result of Gabai and Casson (independently, see [Ga]) which, when combined with earlier results of Mess and Scott (see [M],[Sc]) shows that a 3-manifold which admits an injectively immersed torus, but no incompressible surfaces, must be Seifert-fibered.

Our approach to developing this algorithm will be partly geometric and partly combinatorial in nature. We will first describe a fixed geometric structure on S^3 which lifts in a particularly nice way to a geometric structure on any branched cover of S^3 , branched over the figure eight knot, which has the required minimum branching index. We will then use this geometric structure to gradually reduce the problem of finding injectively immersed tori in this branched cover to a problem of finding paths in a certain graph which satisfy certain easily verified conditions. We will also obtain an *a priori* upper bound on the length of paths which must be considered, so that our algorithm can not only find tori when they exist, but also ascertain when they do not.

The overall organization of this paper is as follows: the geometric structures that we will need (Euclidean cone manifold structures) are described in Section 1. Section 2 discusses the particular cone manifold structure on S^3 which will be lifted to the branched covers of S^3 . Section 3 describes the first dimensional reduction that we will need, reducing questions about tori in 3-manifolds to questions about geodesics in 2-manifolds. Section 4 describes how to make the actual calculations needed in section 3 from combinatorial data

Typeset by $\mathcal{A}\mathcal{M}\mathcal{S}$ -TEX

about the particular branched cover on which we are working. In section 5, we discuss the final dimensional reduction, to a particular kind of graph, and we describe the actual algorithm for finding tori in section 6. Section 7 applies the algorithm to two different examples of tenfold branched covers, one of which contains an incompressible torus and one of which is atoroidal.

We will begin by fixing some notation: let M be a closed, orientable 3-manifold, $K \subset S^3$ the figure eight knot, $\rho : M \rightarrow S^3$ a branched covering map of degree d with downstairs branching locus K . Let $L = \rho^{-1}(K)$ have q components L_1, L_2, \dots, L_q . Then, $\rho | (M - L) : \rightarrow (S^3 - K)$ is a covering map. We associate to this covering map its *monodromy*, defined as follows: let x_0 be a base point in $S^3 - K$ and $\rho^{-1}(x_0) = \{\hat{x}_0, \hat{x}_1, \dots, \hat{x}_{(d-1)}\}$. Then, the monodromy of the covering map is the homomorphism $\varphi : \pi_1(S^3 - K, x_0) \rightarrow S_d$ defined by (allowing permutations to act on the right) $\hat{x}_{i\varphi(\alpha)}$ = terminal point of the lift of a loop representing α with initial point \hat{x}_i . This gives a 1-1 correspondence between degree- d covering maps with connected domain and conjugacy classes of transitive representations of $\pi_1(S^3 - K)$ to S_d (note that $\rho_*(\pi_1(M - L, \hat{x}_i)) = \varphi^{-1}(\text{Stab}(i))$). The *branching indices* of ρ are the integers n_1, n_2, \dots, n_q which are the degrees of the branched covering restricted to disks transverse to the components of L . We note that these are the cycle lengths of $\varphi(\mu)$, where μ is a (fixed) meridional homotopy class of $S^3 - K$, except that the cycles will be duplicated precisely when a component of L covers K nontrivially (longitudinal wrapping).

Section 1 - Cone Manifolds.

The process by which we find the tori in these manifolds is principally geometric in nature and makes extensive use of cone manifold structures. These may be defined quite generally for any constant-curvature model (and even more generally for variable curvature models, see [Ho]), but for our purposes it will suffice to make the following

Definition. *A Euclidean cone manifold is a metric space obtained as the quotient space of a disjoint union of a collection of geodesic n -simplices in \mathbb{E}^n by an isometric pairing of codimension-one faces in such a combinatorial fashion that the underlying topological space is a manifold.*

Such a space possesses a flat Riemannian metric on the union of the top-dimensional cells and the codimension-1 cells. On each codimension-2 cell, the structure is completely described by an angle, which is the sum of the dihedral angles around all of the codimension-2 simplicial faces which are identified to give the cell. The *singular locus* of a cone manifold is the closure of all of the codimension-2 cells for which this angle is not 2π (the Riemannian metric may be extended smoothly over all cells whose angle is 2π). We are principally interested in the 3-dimensional case in which the singular locus is a link with constant cone angle on each component and we make this assumption throughout the remainder of this paper.

There is a strong relationship between cone manifolds on the one hand and branched covers on the other. Specifically, if $\rho : M \rightarrow N$ is a branched covering map, branched over $L \subset N$, then any cone manifold structure on N with the singular locus contained in L lifts to a cone manifold structure on M with the cone locus contained in $\rho^{-1}(L)$. This construction is accomplished by lifting the triangulation given by the cone manifold structure, and the cone angles on M are the obvious ones: if K is a component of $\rho^{-1}(L)$

with branching index n and $\rho(K)$ has a cone angle of θ (possibly 2π), then K has a cone angle of $n\theta$. Note, in particular, that a non-branched cover (all indices = 1) has a cone manifold structure with the same cone angles as the base cone manifold. Note also that some components of the singular locus of the branched covering may actually have cone angle 2π and thus are not singular in the induced cone manifold structure. To avoid confusion, we will refer to the singular locus of the branched covering as the “singular locus” and to the subset of this that has cone angle $\neq 2\pi$ as the “cone locus.”

Next, we consider the geodesics in a cone manifold. A geodesic is extendible until it hits the cone locus. When a geodesic hits a component of the cone locus with cone angle less than 2π , it admits no extension as a geodesic. When a geodesic hits a component of the cone locus with cone angle greater than 2π , it has a pencil of possible geodesic extensions, resulting in the impossibility of extending an exponential map (without allowing singularities in the exponential map itself) past that point, even though geodesic extensions exist. It is fairly easy to see that this does not result in any points of the manifold being inaccessible by uniquely extendible geodesics from any other point unless some cone angle is greater than 2π . Thus, the exponential map is surjective when all cone angles are less than 2π (there are conjugate points even if the manifold is simply-connected, thus implying that cone angles less than 2π “act like” positive curvature). Now, consider a simply-connected Euclidean cone manifold in which all cone angles are greater than 2π . The exponential map is no longer surjective, but it is one-to-one and is in fact a diffeomorphism onto its image. It should be noted that in any case, the domain of the exponential map is \mathbb{E}^3 with a collection of “generalized half-planes” deleted, where a “generalized half-plane” is determined by a line, ray or line segment L in \mathbb{E}^3 and consists of the union of rays originating at points of L which emanate directly away from the origin (the generating segment of each generalized half-plane corresponds to a portion of a component of the cone locus that is in the domain of the exponential map). We note here also that the domain of the exponential map is a subset of \mathbb{E}^3 only when the exponential map is based at a point which is not a cone point. We may, nevertheless define the exponential map based at a cone point, but its domain of definition is a subset, not of \mathbb{E}^3 , but of the Euclidean cone manifold which is homeomorphic to \mathbb{R}^3 and has a single line of cone locus with cone angle equal to the cone angle at the base point. We will, in fact, need this construction subsequently.

It was observed above that cone angles less than 2π “act like” positive curvature. The reverse is true in the following sense:

Lemma 1.1. *If M is a Euclidean cone manifold with all cone angles $> 2\pi$, then M admits a metric of nonpositive sectional curvature, which is flat on the complement of an arbitrarily small neighborhood of the cone locus. Furthermore, any surface that intersects the cone locus transversely has negative sectional curvature relative to this metric in a neighborhood of the intersection.*

Proof. Working in Fermi coordinates in a solid torus neighborhood of a given component of the cone locus, we see that the metric is of the form

$$ds^2 = dt^2 + dr^2 + g^2(r)d\theta^2. \quad \text{where } 0 < r < r_0$$

Where $g(r) = r\psi/2\pi$, and ψ is the cone angle at the given component of the cone locus

Considering all metrics of the given form (allowing g to vary), we observe that such a metric is smooth at the core geodesic if and only if $g(0) = 0$ and $g'(0) = 1$. A straightforward but tedious calculation shows that the principal curvatures are $0, 0$, and $-g''(r)/g(r)$. Our method of smoothing the cone metrics will be to remove a (metrically) regular neighborhood of each component of the cone locus and replace each of these solid tori by another (smooth) solid torus whose boundary metric and second fundamental form agree with the original solid torus. Let us choose the smooth metric to also be of the given symmetric form, say $ds^2 = dt^2 + dr^2 + \hat{g}^2(r)d\theta^2$ for $0 < r < \hat{r}_0$. Then, the two metrics may be glued together smoothly whenever the values and derivatives of g at r_0 agree with the values and derivatives of \hat{g} at \hat{r}_0 to whatever degree of smoothness is desired.

Now, referring to the sectional curvature computation made earlier, it is apparent that if \hat{g} is non-concave and positive, the resulting metric will have nonpositive sectional curvature. Thus, it suffices to construct a non-concave function \hat{g} such that $\hat{g}(0) = 0$, $\hat{g}'(0) = 1$, $\hat{g}(\hat{r}_0) = g(r_0)$, and $\hat{g}'(\hat{r}_0) = g'(r_0)$ for some $\hat{r}_0 > 0$. But, this can be done precisely when $g'(r_0) > 1 \Leftrightarrow \psi > 2\pi$. The second claim is clear from the fact that it is the (r, θ) plane that picks up the negative sectional curvature in this construction. It should be noted that the resulting metric is the product of a negatively curved disk and a circle. \square

Next, we make a few observations about injectively immersed tori in Euclidean cone manifolds with all cone angles $> 2\pi$. The main results we will need involve the relationship between totally geodesic tori in the cone metric and totally geodesic tori in the smooth metrics approximating the cone metric.

Lemma 1.2. *If M is a Euclidean cone manifold with all cone angles $> 2\pi$ then in any of the above smoothings of that metric, any injectively immersed torus is homotopic to a totally geodesic torus. Furthermore, M is irreducible.*

Proof. First, we smooth the cone metric to a metric of nonpositive sectional curvature via Lemma 1.1. We now note that the 2-plane that actually picks up negative curvature in the construction is the 2-plane spanned by dr and $d\theta$. Thus, by the Gauss equation, any minimal surface intersecting the old cone locus transversely will inherit negative curvature in a neighborhood of the intersection. Since there is no place in the manifold to pick up canceling positive curvature, a minimal torus must not intersect the cone locus transversely. A similar Gauss equation argument shows that any minimal torus in a 3-manifold of nonpositive sectional curvature must be totally geodesic.

To complete the proof of the first part of the lemma, we appeal to the standard result that injectively immersed surfaces in 3-manifolds are homotopic to minimal surfaces [SY].

To show irreducibility, we use a similar argument to show that there are no essential 2-spheres, and use the Cartan-Hadamard theorem [Sp] to assert that the universal cover of M is \mathbb{R}^3 and thus that there are no fake 3-cells in M . \square

The next result is not needed in this paper, but illustrates clearly the structure of totally geodesic tori (cone metric and smooth metric) as well as the techniques we will be using subsequently and so is included here as an intuition-building exercise for the reader.

Lemma 1.3. *If M is a Euclidean cone manifold with all cone angles $> 2\pi$ and S is a totally geodesic (cone metric) immersed torus, then there is a smooth metric (as in Lemma*

1.1) in which S is homotopic to a totally geodesic torus which has the same intersection pattern with the smoothing neighborhood that S does with the cone locus.

Proof. If S may be homotoped off of the cone locus, then we're done (take a sufficiently tight smoothing so that S is disjoint from the smoothing neighborhood, then S itself is totally geodesic), so assume that S intersects the cone locus. In order to be totally geodesic in the cone metric, it must contain any component of the cone locus that it intersects and the dihedral angles on each side of the intersection with the cone locus must be $\geq \pi$. By our blanket assumption of "no vertices," the intersections of S with the cone locus must cut S up into immersed annuli of uniform width.

Let w be the minimum width of one of these annuli and ϵ be the maximum distance such that the uniform neighborhood of S of radius ϵ is a regular neighborhood. Take a smoothing of radius $r = \min(w/4, \epsilon/2)$. Now, let U be the uniform neighborhood of S of radius $2r$ (in the cone metric). U is the image of an immersion of $T = S^1 \times S^1 \times I$ into M (by regularity). Call this immersion $\phi : T \rightarrow M$. Thus, we may pull back both metrics (cone and smooth) from M to T , yielding a smooth and a cone metric on T . S being totally geodesic implies that T is convex in the cone metric and hence convex in the smooth metric (they are the same in a neighborhood of ∂T). Furthermore, the action of S^1 on T by translation in the direction of the components of the cone locus is an isometric action with respect to both metrics. Thus, we have induced cone and smooth metrics on $R = S^1 \times I$ (quotient of T by this action - the projection, which we denote by ψ , is a fibration. In fact, it is a fiber bundle projection map with structure group S^1). R is also convex in both of these metrics. In particular, there is a closed geodesic, λ , (with respect to the smooth metric) in the homotopy class of a generator of $\pi_1(R)$. Note that λ cannot "back up" (that is, its pattern of intersections with the smoothing neighborhood is the same as the intersection pattern of $\psi(\phi^{-1}(S))$ with the smoothing neighborhood) since the smooth metric on R has nonpositive curvature.

$\psi^{-1}(\lambda)$ is a totally geodesic (smooth metric) torus homotopic to $\phi^{-1}(S)$ and thus $\phi(\psi^{-1}(\lambda))$ is a totally geodesic (smooth metric) torus homotopic to S via a homotopy in U . Clearly the intersection pattern did not change (by the remark at the end of the preceding paragraph). \square

Using similar techniques, we obtain the following lemma, which we *will* need.

Lemma 1.4. *If M is a Euclidean cone manifold with all cone angles $> 2\pi$ and S is an injectively immersed torus, then S is homotopic to a totally geodesic (cone metric) immersed torus.*

Proof. First, some preliminary remarks about our method of proof: Lemma 1.2 allows us to construct totally geodesic tori homotopic to S in any smoothing of the cone metric. Once we have such a torus, we may "straighten it out" to a piecewise totally geodesic torus in the cone metric by essentially reversing the procedure of Lemma 1.3. However, it need not be the case that all of the dihedral angles at intersections along components of the cone locus are $\geq \pi$. Our game plan is to show that if we take a sufficiently tight smoothing to begin with, this construction does yield a totally geodesic (cone metric) torus.

To begin, we take any smoothing of the cone metric. By Lemma 1.2, S is homotopic to a totally geodesic torus in this metric. If this torus can be taken to be disjoint from

the smoothing neighborhood, then we're done. So, assume that this totally geodesic torus cannot be homotoped off of the smoothing neighborhood (and stay totally geodesic). Then, by the symmetry and structure of the smooth metric in the smoothing neighborhood we must have the torus divided up into annuli that are alternately contained in and disjoint from the smoothing neighborhood. Let λ be a curve in this torus intersecting each of the boundary curves of these annuli exactly once. Let $\hat{\lambda}$ be a closed geodesic (cone metric) in M freely homotopic to λ . It is easy to see that $\hat{\lambda}$ must intersect the cone locus (otherwise the torus could be homotoped off of the cone locus). Let r be the maximum radius such that the uniform neighborhood of the union of $\hat{\lambda}$ and the cone locus is a regular neighborhood and consider now the smoothing of radius $r/2$.

In this smoothing, any geodesic freely homotopic to $\hat{\lambda}$ must have the same intersection pattern with the smoothing neighborhood that $\hat{\lambda}$ has with the cone locus. This follows from a convexity argument similar to the one used in the proof of Lemma 1.3. One obtains a regular neighborhood of $\hat{\lambda}$ which is an immersion of a convex (in either metric) solid torus.

Thus, the totally geodesic (smooth metric) torus homotopic to S has an intersection pattern around its longitude that is the same intersection pattern possessed by the cone geodesic $\hat{\lambda}$. In particular, if we "straighten out" this torus, the dihedral angles at each intersection of two annuli are $\geq \pi$ on both sides. Thus, the straightened torus is in fact totally geodesic in the cone metric. \square

Section 2 - the Figure Eight Knot.

The relevance of Euclidean cone manifolds to the figure eight knot becomes clear when we observe that

Lemma 2.1. *S^3 admits a Euclidean cone manifold structure with cone locus the figure eight knot and cone angle $2\pi/3$.*

Proof. Consider the tessellation of \mathbb{E}^3 by rhombic dodecahedra (see Fig. 2.1).

Figure 2.1

To see that these do indeed tessellate \mathbb{E}^3 , observe that a rhombic dodecahedron may be constructed from two cubes of equal size by cutting one cube into six isometric pyramids with apex at the center of the cube, and then gluing these pyramids onto the faces of the other cube. The tessellation is then clear - take the usual tessellation of \mathbb{E}^3 by cubes, color them red and blue alternately, and dissect all the blue cubes, attaching the pyramids to the adjacent red cubes. Next, consider the Euclidean 3-manifold whose gluing diagram is given by Fig. 2.2 (all identifications are isometric since all faces of a rhombic dodecahedron are congruent).

Figure 2.2

We assert that this manifold is the 3-fold cyclic branched cover of S^3 branched over the figure eight knot. To see this, observe that the edges labeled 3 and 6, together with an arc from the vertex at the center of the diagram to the vertex at infinity are identified into a closed loop. Furthermore, observe that rotation through $2\pi/3$ about the arc gives a \mathbb{Z}_3 action on the manifold with fixed-point set equal to the loop just described (one needs merely to check that this rotation commutes with the identifications on the polyhedron). Also note that, due to the identifications on the polyhedron, this rotation is in fact the same as a rotation about edge 3 or edge 6 (which is as it should be since they are in the same component of the fixed point set as the axis of rotation). The quotient of the Euclidean manifold by this group action gives a Euclidean cone manifold (actually a Euclidean orbifold) whose fundamental domain and identifications are given, along with some fundamental group calculations, in Fig. 2.3.

To see that this manifold is topologically S^3 and that the cone locus is the figure eight knot, observe the topological realization of this diagram in Fig. 2.4. \square

Section 3 - the Associated 2-manifold.

By combining the results of the previous two sections, we see that any branched cover of S^3 , branched over the figure eight knot, with minimum branching index 3, has a Euclidean cone manifold structure with all cone angles greater than 2π , and thus Lemma 1.4 tells us that all of the injectively immersed tori are in fact homotopic to totally geodesic tori.

\parallel = cone locus

a = loop through A ,
back through A' .
Similarly for b, c, d .

$\pi_1(M - \text{cone locus})$
= $\langle a, b, c, d : d^{-1}b^{-1}c,$
 $b^{-1}aba^{-1}c, a^{-1}d^{-1}c \rangle$

Figure 2.3

all identifications
are on ∂B^3 .
Heavy solid arcs are
in the interior.

glue C, C'
to yield S^3 .

Figure 2.4

In particular, if we base the exponential map at a point on the lift of such a torus to the

universal cover of M (which we denote by \hat{M}), we have that the lift of the torus is in the image of the exponential map (being totally geodesic) and that such a torus corresponds to a 2-plane in \mathbb{E}^3 which

- (a) misses (at least with regard to transverse intersections) all of the “bad” line segments (that map to the cone locus of \hat{M}) and
- (b) has compact image in M (via the composition of the exponential map with the covering map).

At this point, we have two distinct cases to consider, depending on whether or not there is any cone locus in our branched cover M . If there is no cone locus, then any plane satisfying (b) will automatically satisfy (a). In the case of branched covers over the figure eight knot with all branching indices equal to 3, (b) is always satisfied by any plane having a normal with rational coordinates relative to the coordinatization given above (since translation by 2 in the principal directions is an orbifold deck transformation of the figure eight knot orbifold structure, any such plane has compact image in S^3 . Since the branched covering is a proper map, this plane must have compact image in M , also). Thus, when all branching indices are equal to 3, there are always injectively immersed tori in M .

Now, consider the case in which the cone locus of M is nonempty (the remainder of this paper will be occupied with this case). The following lemma allows us to restrict our search to tori that contain a component of the singular locus.

Lemma 3.1. *Let M be a Euclidean cone manifold with nonempty, compact cone locus which has all cone angles $> 2\pi$ and let S be an injectively immersed torus in M . Then, S is homotopic to a torus T which contains some component of the cone locus of M .*

Proof. First, by Lemma 1.4 we may assume that S is totally geodesic and flat, and thus contains any component of the cone locus that it intersects. So, if S actually intersects the cone locus of M , we’re done. Therefore, assume that S is disjoint from the cone locus of M . Let r be the minimum distance between S and the cone locus of M . Lift S to a plane P in \hat{M} , the universal cover of M . Then, the uniform r -neighborhood of P contains no cone locus and hence is a flat “slice” of \hat{M} . Clearly, S is homotopic to the image in M of either boundary component of this neighborhood. Furthermore, these contain a component of the cone locus of \hat{M} (if they intersect transversely, the cone locus would penetrate the r -neighborhood, in violation of the definition of r). Let T be the image in M of either of these boundary components. \square

Thus, to find a “canonical form” for injectively immersed tori in M (where M is a branched cover over the figure eight knot with minimum branching index 3) we will homotope the torus to a totally geodesic torus containing some component of the cone locus, which will cut the torus up into totally geodesic annuli. We should note that such a “canonical form” is not, strictly speaking, canonical. When we have a totally geodesic torus which is disjoint from the cone locus, there are two possible directions in which it may be translated, and these give rise to distinct (but homotopic) tori which are cut up into totally geodesic annuli by intersections with the cone locus. We will discuss the resolution of this ambiguity in Section 6. Our next question, however, is when such annuli exist, or equivalently, when two components of the cone locus in the universal cover of

M may be joined by an infinite flat strip (the universal cover of an annulus). To answer this question, we will work in \mathbb{E}^3 , or rather, the domain of the exponential map of \hat{M} , the universal cover of M .

Suppose that we are given a component C_0 of the cone locus of \hat{M} , and we wish to determine which other components of the cone locus cobound, together with C_0 , an infinite flat strip which does not intersect the cone locus except at its boundary (we will call this a *simple strip*). Let us base the exponential map at a point on C_0 . Recall that the domain of definition of the exponential map based at a cone point is a subset of the cone manifold homeomorphic to \mathbb{R}^3 with a single line of cone locus. Let us call this manifold N . Note that, generally speaking, N is no more difficult to work with than \mathbb{E}^3 , because there is no interaction between various components of the cone locus. Working in N , we see that any component of the singular locus which cobounds a simple strip with C_0 has a preimage in N parallel to C_0 . Consider now the projection of N to a plane P orthogonal to the preimage of C_0 (and to the other component, as well). Note that this plane will be Euclidean with a single cone point. Let this projection be denoted $\eta : N \rightarrow P$. Recall that the complement of the domain of the exponential map consists of a number of generalized half-planes. Denote by S the subset of P consisting of the projected images of the generating segments of these generalized half-planes. S consists of line segments (corresponding to portions of the cone locus not parallel to C_0) and isolated points (corresponding to components of the cone locus parallel to C_0) and is a subset of the set consisting of the edges and centers of the hexagons in the tessellation of P by equilateral hexagons and triangles (there is also one equilateral $2n$ -gon, where $2\pi n/3$ is the cone angle at C_0 – see Figure 3.1). Let S_0 denote the union of the isolated points of S . Then, the other components of the singular locus of \hat{M} which cobound a simple strip with C_0 are the preimages under η of the points of S_0 such that the interior of the geodesic segment from $\psi(C_0)$ is disjoint from S . Two simple strips with a common boundary component can be combined into a geodesic strip precisely when the projection at this common component of the cone locus consists of geodesic segments which meet with angles $\geq \pi$ on each side (measured in the normal cone 2-manifold).

It is convenient at this point to introduce the notion of an *elemental strip*. An elemental strip is an infinite flat totally geodesic strip in \hat{M} cobounded by two adjacent components of the preimage in \hat{M} of the figure eight knot (that is, the lifts to \hat{M} of the singular locus of M). By *adjacent*, we mean that the distance between them is the minimum distance possible for parallel components of the preimage of the figure eight knot. An elemental strip corresponds to a segment from the center of a hexagon to the center of an adjacent hexagon in Fig. 3.1.

Next, we observe that we may construct a (disconnected) 2-dimensional Euclidean cone manifold with boundary that encapsulates all the information necessary to decide whether or not there exists a geodesic strip cobounded by any two given components of the cone locus of \hat{M} . We will refer to this as the *associated 2-manifold* of M and construct it as follows: For each component C of the preimage in \hat{M} of the figure-eight knot, take an equilateral polygon with $2n$ sides, where $2\pi n/3$ is the cone angle at C (may be 2π). This equilateral polygon will be metrically a Euclidean cone manifold with (at most) a single cone point with cone angle $2\pi n/3$ at the center, and corresponds to the central polygon of the normal cone manifold in the domain of the exponential map, based at some point

Note: only the left picture is conformally correct

Figure 3.1

on C . We then glue together (at their *vertices*, not their edges) polygons corresponding to components of the figure eight knot which cobound elemental strips, in a fashion which respects the obvious ordering of the elemental strips around a component of the cone locus and respects the orientation of each polygon by the normal direction it inherits by lifting an orientation of the figure eight knot. At this stage, we have the “hexagons” in our tessellation taken care of, and just need to add the triangles (which are *really* triangles since there is no cone point in the interior of a triangle in Fig. 3.1). We do this by gluing in an equilateral triangle on “both sides” of any vertex along which two polygons are glued. That is, at any vertex where two polygons are glued, the ordering of the elemental strips and the orientation of the polygons fixes two pairs of edges (one edge from each polygon in each pair) which are on the same side of the vertex, and we will glue in a triangle at each such pair of edges. We will identify any two such triangles that share a common edge. This construction produces a Euclidean cone 2-manifold with boundary. The boundary consists of edges of the polygons, neither of whose endpoints correspond to an elemental strip, together with the free edges of any triangles which are incident on only two polygons.

By construction and our remarks above, two components of the cone locus of \hat{M} cobound a geodesic strip if and only if there is a geodesic in the associated 2-manifold joining the centers of their corresponding polygons. In fact, as the following lemma will show, they cobound a geodesic strip if and only if their polygons are in the same component of the associated 2-manifold.

Lemma 3.2. *The associated 2-manifold of M , a branched cover over the figure eight knot with all branching indices > 3 is convex and each component is simply-connected.*

Proof. Convexity first: Since the boundary of the associated 2-manifold is piecewise geodesic, it suffices to verify that, when measured from “inside” the associated 2-manifold, all corners have angle $\leq \pi$. Since the boundary consists of two distinct kinds of edges (free edges of polygons and free edges of triangles), we have three kinds of corners to consider.

For the moment, let us assume that polygons share a vertex *only* when there is an elemental strip between them (by construction, if there is an elemental strip then they share a vertex, but we have not shown the converse yet).

Polygonal-polygonal corners: any two polygonal boundary edges that meet at a corner must be consecutive edges of the same polygon, since, if they belonged to different polygons, two triangles would have been glued in at that vertex and neither of the edges in question would have been on the boundary. Consecutive edges of the same polygon meet at an angle of $2\pi/3$ when measured from the center of that polygon.

Polygonal-triangular corners: inspection of Fig. 3.1 shows that these corners all meet at an angle of π .

Triangular-triangular corners: this is the only nontrivial case. It suffices to show that if a polygon, P_0 , is joined to two other polygons, P_1 and P_2 , at vertices which are separated by only one vertex, then P_0 is also joined to a polygon P_3 at the intervening vertex. This will show that there are no triangular-triangular corners, since any such corner would necessarily involve two different triangles (by construction, a triangle has at most one free edge). So, suppose that P_0 , P_1 , and P_2 are configured in this way and there is no elemental strip from the center of P_0 through the vertex between the two vertices where P_0 is joined to the other two polygons. Consider the exponential map based at a point on the component of the singular locus of \hat{M} corresponding to the center of P_0 . The domain of this map contains an infinite prism isometric to $P_0 \times \mathbb{R}$. The fact that there is no elemental strip in the direction of the vertex in question (call it v_3), implies that there is some component of the cone locus of \hat{M} which intersects the corresponding line $v_3 \times \mathbb{R}$. Now, a component of cone locus remains “visible” (in the domain of the exponential map) at least until it passes “behind” another component of the cone locus. In particular, any components of the cone locus that intersect $v_3 \times \mathbb{R}$ traverse one or the other (depending on their slope) of the faces of $P_0 \times \mathbb{R}$ that is incident on $v_3 \times \mathbb{R}$. So, one of the adjacent $v_i \times \mathbb{R}$ intersects this same component of cone locus. But there are elemental strips containing each of these lines, so they can’t intersect the cone locus at all. Thus, there must be an elemental strip through $v_3 \times \mathbb{R}$ as well.

Now, this argument also shows that no two vertices of polygons are “accidentally” identified. That is, if any two vertices of polygons are identified because they are identified to the same vertex of a triangle, then they are also identified because there is an elemental strip through that vertex. This is because the nonexistence of an elemental strip implies the existence of some cone locus which would be positioned so as to necessarily rule out the existence of one or the other of the identifications at the other two vertices of the triangle (which are *both* necessary to force such an “accidental” identification).

With convexity proven, it is an easy matter to show simple connectivity, since we need only show that there are no geodesic loops in the 2-manifold. But this is clear, since a geodesic loop in the 2-manifold would imply the existence of a geodesic segment in \hat{M} whose endpoints are on the same component of the cone locus, but its interior is disjoint from that component of the cone locus. This would imply, in a suitably tight smoothing of \hat{M} , the existence of two distinct geodesic segments with the same endpoints, contradicting the Cartan-Hadamard theorem. \square

The connection between tori in M and geodesics in the associated 2-manifold is as

follows: to each polygon of the associated 2-manifold there corresponds a component of the singular locus of \hat{M} . Each component of the singular locus of \hat{M} covers a component of the singular locus of M . Label the polygons of the associated 2-manifold with numbers from 1 to k , where k is the number of components of the singular locus of M , using this correspondence. Then an injectively immersed torus in M corresponds to an infinite geodesic in the associated 2-manifold of M which “repeats” in the sense that the pattern of labels of the polygons which are intersected by the geodesic repeats.

It should be noted here that *most* totally geodesic tori can be homotoped to one of their canonical forms (the tori containing components of the cone locus) simply by parallel translating the corresponding geodesics in the associated 2-manifold. This is not universally true, however: it is possible to encounter a geodesic boundary component of the 2-manifold before encountering any cone points in the 2-manifold. In such a case, the boundary geodesic corresponds to a geodesic in another component of the 2-manifold that *does* pass through a cone point. If we wanted to be able to realize any homotopies of totally geodesic tori inside the 2-manifold, we would need to identify these boundary components to the corresponding geodesic in the other component of the 2-manifold, creating a 2-complex. We will not take this approach, but we will have more to say about this subject near the end of Section 5.

Thus, we have succeeded in reducing questions about tori in 3-manifolds to questions about geodesics in 2-manifolds. Our next task is to answer the 2-dimensional questions from the monodromy of M .

Section 4 - the Elemental Strip Condition.

In this section, we will derive a condition for determining, from the monodromy of M , whether or not an elemental strip exists at any point in \hat{M} .


Consider first \mathbb{E}^3 in its tiling by fundamental domains of the orbifold structure on the figure eight knot. It is clear that a plane containing an elemental strip must either contain face A (of some fundamental domain) or bisect the angle between face A and face A' . So, let us consider all the ways such a plane can intersect other fundamental domains of the tiling. Note that an elemental strip only intersects the cone locus on its boundary, so if an elemental strip exists, its intersection pattern with these fundamental domains is in fact exactly like the intersections in \mathbb{E}^3 where there is no cone locus.

There are eight combinatorially distinct ways that a plane containing face A or bisecting faces A and A' can intersect other fundamental domains. These are pictured in Fig. 4.1.

To see that these are, in fact, all of them, we observe that these regions glue together as described in Fig. 4.2 to give a four-punctured annulus in $S^3 - K$. Thus, we may extend our plane in any direction without leaving this set of intersection patterns (note, however, that we have suppressed one of each pair of patterns which are faces of the fundamental domain, e.g., face A is listed, but not face A').

Perhaps a few words of explanation are in order about Fig. 4.2: this graph is the dual graph to the tiling of the annulus by the regions given in Fig. 4.1. The reason why this is the object of interest is that the monodromy can be used, together with this diagram, to calculate *which* fundamental domains in M contain each piece of the tiling of a given strip, and thus what the cone angles are at each intersection with the singular locus.

Since each of the eight planes occurs exactly once in our four-punctured annulus, we

**Figure 4.1**

bold lines denote non-transverse intersection with the singular locus

* denotes a transverse intersection with the singular locus

dashed lines indicate homotopic reference paths

Figure 4.2

may take any one of them as our reference point in stating our conditions for existence of elemental strips. We will consider plane 1 (see Fig. 4.1) as the reference point. Let us fix as our basepoint in $S^3 - K$ a point x_0 on plane 1 which is not in any of the other 7 planes. Then, if we denote by $\hat{x}_0, \hat{x}_1, \dots, \hat{x}_{d-1}$ the d points of M in the pre-image of x_0 , we can consider the problem of whether or not our annulus in S^3 can be lifted to an annulus in M at any of these d points. This is equivalent to deciding which elemental strips exist in \hat{M} .

Let us say that a lift of plane 1 at \hat{x}_k is an *acceptable lift* if the interior of the entire par-

allelogram region depicted in Fig. 4.2 can be lifted at that basepoint without intersecting the cone locus (in other words, the four punctures all correspond to the locus of threefold branching). By examining Fig. 4.2, we see that there is an acceptable lift at \hat{x}_k if and only if (recall that φ is the monodromy of M and that we are writing permutation actions on the right)

$$k \in F(M) = \begin{array}{l} \text{fix}(\varphi(d^3))\varphi(a^{-1}cab^{-1}) \cap \text{fix}(\varphi(d^3))\varphi(b^{-1}) \\ \cap \text{fix}(\varphi(a^3))\varphi(b^{-1}) \quad \cap \text{fix}(\varphi(a^3))\varphi(cab^{-1}) \end{array}$$

Now, we consider the problem of whether or not a sequence of such acceptable lifts will actually close up into an annulus. This will happen for a lift at x_k if and only if

$$k \in G(M) = \{j \mid \text{orbit}(\langle \varphi(ba^{-1}c^{-1}ad) \rangle, j) \subset F(M)\}$$

Furthermore, the set of “other boundary components” (taking plane 5 as the reference point for the other side) of such annuli is $G(M)\varphi(bd^{-1})$.

We will summarize the results of this section in a lemma, for which we need the following

Definition. *a potential elemental strip is an ordered pair (C, V) where C is a component of the singular locus of \hat{M} , and V is a normal vector to C in one of the $2k$ possible directions for an elemental strip (where $2k\pi/3$ is the cone angle at C). Each potential elemental strip is either a primary or secondary potential elemental strip depending on whether the vector V lies on or bisects (respectively) the A, A' faces in the tessellation of \hat{M} by fundamental domains for $S^3 - K$. Note that an elemental strip always has one side corresponding to a primary potential elemental strip and the other corresponding to a secondary elemental strip and that each potential elemental strip can be labelled with a subset of $\{0, 1, \dots, d-1\}$ by taking the labels of the set of “plane 1” (primary) or “plane 5” (secondary) planar pieces in M which are covered by the potential elemental strip.*

With this definition, we can state the following

Lemma 4.1. *If \hat{M} is the universal cover of M , a branched cover of S^3 , branched over the figure eight knot, with all branching indices ≥ 3 , then a potential elemental strip in \hat{M} corresponds to an actual elemental strip if and only if its label set is a subset of $G(M)$ in the case of primary potential elemental strips or $G(M)\varphi(bd^{-1})$ in the case of secondary elemental strips.*

Section 5 - Flat Graphs.

So, we now have a condition, in terms of the monodromy, for deciding on the existence of an elemental strip. Let us now consider the problem of finding the “repeating” geodesics in the associated 2-manifold that correspond to tori in M . Our plan of attack on this problem will again be a dimensional reduction. We will define a graph and a map taking paths in the associated 2-manifold to paths in the graph such that geodesics in the associated 2-manifold are taken into an algorithmically recognizable class of paths in the graph. Furthermore, “repeatable” geodesic segments (i.e. segments that are the basic period for a repeating geodesic) are readily recognizable in the graph and there is a universal upper bound to the length of a nonrepeatable geodesic image segment in the graph. This enables

us to obtain a terminating algorithm for finding any tori that are present by enumerating all paths in the graph of length less than this upper bound and checking whether there are any that are repeatable, thus corresponding to repeating geodesics in the associated 2-manifold and to tori in M . In fact, we can do better than this in terms of efficiency: the class of paths that are possible images of geodesics in the associated 2-manifold is not only algorithmically recognizable, it is algorithmically enumerable, so we can limit our search significantly by simply enumerating all of the paths of this class which have length less than the upper bound.

We will begin by making the following

Definition. *a flat graph is a graph whose vertices are oriented polygons and whose edges are line segments joining two polygonal faces (see Fig. 5.1) such that each face of a vertex is incident on at most one edge.*

Figure 5.1

Note that a flat graph is much like a “fat graph” except that there are “stumps” in a flat graph, that is, unattached edges of the polygonal vertices. It will be clear momentarily why we need to know not only in what order the edges occur around a vertex (which would be encoded in a fat graph) but also how many “stumps” occur between edges.

Now, in addition to the 2-manifold associated to M (and closely related to it), we will construct an *associated flat graph*, which will be compact, and which will also, as we will see, encode all of the information necessary to find tori in M or to determine atoroidality.

To construct the associated flat graph of M , take one polygon with $2k$ sides for each k -cycle in $\varphi(a)$. If this cycle is $(a_1 a_2 \dots a_k)$ then label the sides of this $2k$ -gon $-a_1, +a_1, -a_2, +a_2, \dots, -a_k, +a_k$ in clockwise order around the polygon. Take the disjoint union of these polygons over all cycles in $\varphi(a)$. Then, for every edge labelled $+i$ where $i \in G(M)$ (see section 4), join $+i$ to $-i\varphi(bd^{-1})$.

Now, there appears to be an obvious mapping from paths in the associated 2-manifold to paths in the flat graph (although some care must be taken in making this well-defined

for nongeneric paths that intersect some polygon only in one of its vertices – we will make the convention that we perturb such paths slightly to the right before mapping it to the flat graph). That is almost true, but not quite! In fact, it is a one-to-many mapping, since there may be more than one cycle in $\varphi(a)$ for each component of the singular locus in \hat{M} . This is because of the possibility of “longitudinal wrapping” in which a component of the singular locus of M nontrivially covers the figure eight knot in S^3 . However, all of the properties in which we are interested are equivalently true for any of these mappings, so if we allow this “mapping” to be a multi-valued function, it will in fact cause us no difficulties. Let us denote this mapping by η . Denote by \mathbb{G} the set of geodesic segments in the 2-manifold with endpoints at centers of the polygons in the 2-manifold. We would like to be able to recognize $\eta(\mathbb{G})$ algorithmically. We will not, however, do exactly that. Instead, we will describe an algorithmically recognizable class, \mathbb{F} , of paths in the flat graph which contains $\eta(\mathbb{G})$. Furthermore, we will show that paths in \mathbb{F} give rise to paths in \mathbb{G} .

To describe \mathbb{F} , we first observe that, in a flat graph, a path may be described completely by specifying an initial (oriented) edge and a sequence of “turns” where each turn specifies how many faces apart the entering and exiting edges are at each intermediate vertex. Each turn may be measured in either the positive (clockwise with respect to the orientation of the vertices) or negative direction (see Fig. 5.2). If the particular vertex under consideration has k faces, then $+j$ is the same turn as $-(k - j)$.

Figure 5.2

To describe the class of paths \mathbb{F} , we will describe a finite-state automaton, which, when presented with paths in the form just described (initial edge, followed by a sequence of turns) “accepts” some paths and “rejects” others. The class of paths accepted by this finite-state automaton will be defined to be \mathbb{F} .

The finite-state automaton has $6r$ states where r is the number of edges in the flat graph. Each state consists of an ordered pair (e, m) where e is an oriented edge in the flat graph and $m \in \{-1, 0, +1\}$. If the initial oriented edge of the path is e_0 , then the initial state of the automaton is $(e_0, 0)$. Thereafter, the following rules determine the state transitions

from state (e_k, m_k) to (e_{k+1}, m_{k+1}) under a turn t_k where t_k is an integer modulo n_k , the number of edges at the terminal vertex of e_k :

- (1) e_{k+1} is the obvious edge obtained by making the required turn. If this edge does not exist in the graph, then REJECT.
- (2) if $t_k = 0, \pm 1$, then REJECT.
- (3) if $m_k = 0$ and $t_k = \pm 2$, then $m_{k+1} = \pm 1$.
- (4) if $m_k = \pm 1$ and $t_k = \pm 2$ (same sign), then REJECT.
- (5) if $m_k = \pm 1$ and $t_k = \mp 2$, then $m_{k+1} = \mp 1$.
- (6) if $m_k = \pm 1$ and $t_k = \pm 3$ (same sign), then $m_{k+1} = \pm 1$.
- (7) if none of (2) - (6) apply, then $m_{k+1} = 0$.

If the end of the path is reached without a rejection, then the path has been accepted by the finite-state automaton and is in \mathbb{F} .

Hereafter, we will, by a slight abuse of terminology, also refer to the second coordinate of the state pair as the state of a path at a vertex. It will be clear from the context whether we mean for “state” to include both coordinates or just the second.

Lemma 5.1. $\eta(\mathbb{G}) \subset \mathbb{F}$

Proof. First, note that η of any path in the associated 2-manifold will never be rejected on account of rule 1 above – that rule is just for rejecting turn sequences that don’t even correspond to paths, much less geodesics. Now, a geodesic segment in \mathbb{G} has the following structure: it consists of a number of Euclidean geodesic segments between cone points (that is, segments that have a neighborhood in which there are no cone points except for the cone points at the ends) which meet in such a way that the angle on each side is $\geq \pi$. Now, each Euclidean geodesic segment has the property that η maps it to a path in the flat graph that is identical to the path of a line segment in the “no cone point” tiling on the left in Fig. 5.1. These segments are all accepted by our finite-state automaton since each turn in these segments is either a ± 3 (whichever it takes to remain on the proper side of the cone point centers of the polygons) or a ± 2 , where the $+2$ and -2 turns alternate (possibly with some intervening ± 3 turns). The state variable thus has the geometric significance of remembering the sense of the last change of direction.

Now, when our path goes through a cone point, the last edge and the state determine (together) a $\pi/3$ sector in which the initial point of this Euclidean segment might lie (to an observer at the terminal cone point). If this cone point has cone angle $2k\pi/3$, then the next polygon along a geodesic can be any one of $2k - 4$ polygons. All of these but 2 are completely unconstrained, that is, *any* Euclidean segment that starts out into that polygon will necessarily make an angle of π or more on each side of the cone point. This corresponds to “state 0.” The other two polygons are such that only Euclidean segments that lie to one side or the other of their center correspond to geodesic continuations of the previous Euclidean segment, and these are precisely what will be enforced by the finite-state automaton. See Figure 5.3 for an example of this where $k = 4$. In this figure, the shaded areas are the sectors in which the previous cone points must lie, which are determined by the state: state $+1$ has the shaded sector to the left of the edge, and state -1 has the shaded sector to the right of the edge. In state 0, all previous cone points are accessible only by first retracing the last edge (they are all “hidden” behind the most

Figure 5.3

recent cone point). In state 0, there are $2k - 3$ possible continuations of the path which lie in \mathbb{F} .

Thus, we see that any path in \mathbb{G} is mapped by η into \mathbb{F} . Further, it does not matter which “value” of η we use in this argument (since the graph is defined in a manner which is equivariant with respect to longitudinal translation), and we get our first confirmation of the assertion made earlier that the fact that η is multi-valued will cause us no problems. \square

It is tempting to hope that our finite-state automaton is a “geodesic recognition” automaton, but this is not the case: there are many paths in \mathbb{F} that are *not* the images under η of any geodesic in \mathbb{G} . To see this, consider Figure 5.4

Figure 5.4

The heavy path in Fig. 5.4 is in $\eta(\mathbb{G})$. However, the light paths also map under η to \mathbb{F} .

This is a situation that happens quite generally, as the following lemma suggests:

Lemma 5.2. *In the flat graph associated to a branched cover over the figure eight knot, if S_1 and S_2 are any sequences of turns, then, with any initial edge whatsoever,*

$$\langle S_1, j, -2, k, S_2 \rangle \in \mathbb{F}$$

if and only if

$$\langle S_1, j + 1, +2, k - 1, S_2 \rangle \in \mathbb{F}$$

Proof. The two paths are situated as pictured in Figure 5.5, where the “upper” path is $\langle S_1, j, -2, k, S_2 \rangle$.

Figure 5.5

Now, it is clear from a routine case-by-case verification, based on the states before and after the j (resp. $j + 1$) turn (the state after the ± 2 is always ± 1) that the only possible problem would be that one path or the other would simply fail to exist in the flat graph, i.e., only rule 1 is a problem. And, indeed, in a general flat graph, the existence of one of the two paths does *not* imply the existence of the other. However, the flat graphs that arise from our construction from a branched cover are special; the proof of Lemma 3.2 shows us that two special conditions are true in these flat graphs:

- (1) At a vertex V_0 , if e_1 and e_2 are two adjacent edges with other vertices V_1 and V_2 respectively, then V_1 and V_2 are joined by an edge e_3 which is mutually adjacent to e_1 and e_2 at V_1 and V_2 , respectively (this is equivalent to the observation that no two polygons are “accidentally” identified in the 2-manifold).
- (2) At a vertex V_0 , if e_1 and e_2 are two edges with only one possible intervening edge, then there is an actual edge e_3 between e_1 and e_2 (this is equivalent to the observation that there are no triangular-triangular corners on the boundary of the 2-manifold).

Thus, the existence of the edges of the upper path implies the existence of the edge that bisects the parallelogram (condition 2) which then implies the existence of both edges of the lower path (condition 1 used twice). Similarly, the existence of the edges of the lower path implies the existence of the edges of the upper path. \square

This lemma tells us that any path in \mathbb{F} which contains a ± 2 turn gives rise to a path in \mathbb{F} which is not in $\eta(\mathbb{G})$ (since the two paths have the same endpoints and thus at most one of them can be in $\eta(\mathbb{G})$). However, (as the following lemma will show) it is true that repeatable paths in \mathbb{F} give rise to repeatable geodesic segments in \mathbb{G} . Our method of proving this will be to alter an arbitrary repeatable path in \mathbb{F} to a path in $\eta(\mathbb{G})$ by using moves of the type described in Lemma 5.2, then to show that this path “lifts” to a repeatable geodesic segment in the 2-manifold. Let us first formalize the definition of “repeatable”:

Definition. *a subpath of a path in \mathbb{F} is said to be repeatable if its initial and terminal oriented edges are the same and are traversed in the same state. A geodesic segment g in \mathbb{G} is said to be repeatable if g passes through a cone point, $\eta(g)$ is closed and there is a geodesic segment h in \mathbb{G} such that $\eta(h) = \eta(g)\eta(g)$.*

The reason for taking subpaths in \mathbb{F} is that all paths begin in state 0, whereas this is not true for subpaths of these paths. This enables us to have more control over where a subpath is allowed to go (e.g., any lift of a subpath that starts in state +1 must stay “to the right” of its initial edge) in a technically convenient way. The cone point condition for repeatable segments in \mathbb{G} only rules out the geodesics that do not correspond to “canonical” tori in the sense of Lemma 3.1.

Lemma 5.3. *There exists a repeatable subpath of a path in \mathbb{F} if and only if there exists a repeatable geodesic segment in \mathbb{G} .*

Proof. The “if” portion of this lemma is clear from Lemma 5.1 together with the following observation: if $g \in \mathbb{G}$ is repeatable, then there is also a geodesic segment $k \in \mathbb{G}$ such that $\eta(k) = \eta(g)\eta(g)\eta(g)$. Then, the subpath of $\eta(k)$ consisting of the middle $\eta(g)$ together with the initial edge of the last $\eta(g)$ is a repeatable subpath in \mathbb{F} .

For the “only if” part, we will first need a number of technical definitions and observations.

First, note that any path in \mathbb{F} has the property that it is also in \mathbb{F} when traversed backwards (the automaton essentially only rules out paths that make 0 or ± 1 turns or consecutive ± 2 turns of the same sign without an intervening turn larger than ± 3 – in this form it is clearly symmetric). Let us denote this path by \bar{p} .

Now, given a path p in \mathbb{F} , we will define its *inflections* to be the vertices at which either p or \bar{p} re-enters state 0. A subpath between two inflections will be called an *elementary subpath*. Note that an elementary subpath consists entirely of ± 2 and ± 3 turns. It will be technically convenient to always include the endpoints of a path (or subpath) in the set of inflections.

Given an elementary subpath q of a path p in \mathbb{F} , we can define a turn sequence $t(q)$ as follows: in the flat tessellation (no cone points) given in Fig. 5.1, trace out a path given by the turn sequence of q (using only $\pm 2, \pm 3$ turns), then straighten it to a geodesic g . $t(q)$ is the turn sequence of this geodesic g (again using only $\pm 2, \pm 3$ turns). Note that it is

possible that the initial edge of g is different from the initial edge of the “lift” of q to the flat tessellation. Because of this, we will prepend to this turn sequence a turn describing the difference between these two edges (this turn will be a 0 or ± 1 turn but will really be a way of specifying the initial edge of the geodesic turn sequence). Note that $t(q)$ need not be in \mathbb{F} simply because it need not correspond to a path in the 2-manifold – it is simply an idealized geodesic turn sequence. We make here the observation (which we will need subsequently) that if we consider both paths to start in the only nonzero state which will allow the first ± 2 turn to take place, then each turn sequence (q and $t(q)$) traverses the same number of edges in each nonzero state (this is true of any two paths in the flat tessellation which follow the rules of \mathbb{F} and have the same endpoints).

Similarly, given an elementary subpath q , we can define its *geodesic agreement number* $n(q)$ to be the sum of the lengths of the maximal initial and terminal subpaths of q which agree with $t(q)$ (or its reversal). In the case where q agrees completely with $t(q)$, however, we will set $n(q) = \ell(q)$ where $\ell(q)$ is the length of q (note that the definition without this proviso would have been $2\ell(q)$). Note that for the purposes of determining “agreement” we will consider the ± 3 turns in $t(q)$ to have whichever signs maximize the agreement.

Now, define the *complexity* of a subpath q of a path p in \mathbb{F} to be $c(q) = (s, t)$ where s = the number of non-inflection vertices of q and t = the sum of $\ell(q') - n(q')$ over all elementary subpaths q' of q . We will order complexities lexicographically.

Now, let q be a repeatable subpath of a path $p \in \mathbb{F}$. Suppose that $q \neq \eta(g)$ for g any geodesic segment in \mathbb{G} . Then, the second coordinate of $c(q)$ is nonzero. We will show that we can alter the subpath q by a sequence of moves of the type described in Lemma 5.2 to reduce $c(q)$. Thus, there is some elementary subpath q' of q and a vertex v_0 of q' where q' departs from $t(q')$. There are four distinct cases to consider – let t_1 be the turn in q' at v_0 and t_2 be the turn in $t(q')$ at v_0 . Then, the cases are:

- (1) $t_1 = \pm 3, t_2 = -2$
- (2) $t_1 = \pm 3, t_2 = +2$
- (3) $t_1 = -2, t_2 = \pm 3$
- (4) $t_1 = +2, t_2 = \pm 3$

We will do case (1) in detail; the other cases are completely analogous. Since, as observed earlier, q' and $t(q')$ traverse the same number of edges in each state (and we must be in state $+1$ for the $t(q')$ turn to be legal) that there must be some vertex v_1 in q' , later than v_0 , where q' makes a -2 turn. All intervening turns between v_0 and v_1 must be ± 3 since q' is an elementary subpath, so, before and at v_1 we have the turn sequence $\pm 3, -2$ to which we may apply the move described in Lemma 5.2 to obtain either $-2, +2$ or $+4, +2$. In the former case, we can repeat the move, until we have either moved the -2 turn all the way back to v_0 or introduced an additional inflection into q . In the first case, we have increased the geodesic agreement number of the subpath by at least 1 (note that we could not have decreased the geodesic agreement of the other end of the subpath – this is easily seen by tracing out both turn sequences in the flat tessellation and is left to the reader). In the second case, the introduction of an additional inflection also lowers the complexity (although in this case the total geodesic agreement will often be reduced – it is essential that the coordinates of the complexity not be reversed from the order in which we have defined it). Note that our new subpath is repeatable if and only if the old subpath was (although we may have to alter the path in which it is imbedded also – any moves that

altered the initial or terminal edges or states could also have been done in the immediate vicinity of the opposite end to maintain repeatability).

Thus, we have a new repeatable subpath \hat{q} which is $\eta(g)$ for some geodesic segment g between two singular points in the 2-manifold. Now, suppose that \hat{q} contains a *real* inflection (i.e., not just an endpoint). Then, by taking a circular shift of \hat{q} (possible because it is repeatable) to place the endpoints of \hat{q} on this inflection, and then repeating the complexity reduction process just described, we see that $\hat{g} \in \eta(\mathbb{G})$ and that the incidence angles when we juxtapose two lifts of \hat{g} are both $\geq \pi$ (recall the definition of an inflection and examine Fig. 5.3). Thus, in this case we have a repeatable geodesic segment in \mathbb{G} .

Now, when \hat{q} contains no real inflections, the geodesic g only intersects cone points incidentally, that is, in such a way that one of the two angles of incidence is π . For such geodesics the angles between the geodesic and the incoming (resp. outgoing) edge at any singular point it encounters are constant. Thus, the positioning within the $\pi/3$ -width described in the proof of Lemma 5.1 is fixed for these geodesics. Let us now consider the situation when two lifts of \hat{q} are juxtaposed: clearly the incoming and outgoing angles are such that the incidence angles are equal to 0 modulo $\pi/3$, and then the repeatability condition on \hat{q} implies that one angle or the other is equal to π (since they are always forced to be $> 2\pi/3$ by the conditions on \mathbb{F} and if they were *both* greater than π then that vertex would be a real inflection). Thus, we are done, if this geodesic passes through a cone point. If not, we may translate this geodesic in a normal direction until we encounter a cone point or the boundary of the 2-manifold. If we encounter a cone point, then we're done. If we encounter the boundary of the 2-manifold, then by convexity of the 2-manifold, we must have a geodesic in the boundary of the 2-manifold. This corresponds to a torus in M that is parallel to two distinct parallel families of singular locus (one family corresponds to the component in which we have been translating our geodesic, the other corresponds to the family of parallel components of cone locus which prevent the 2-manifold from being extended past this component of boundary). This torus, like all the others, has a canonical form in which it contains a component of the singular locus, but its canonical form corresponds to a geodesic in *another* component of the 2-manifold! (see Fig. 5.6 – the two components of the 2-manifold are depicted as locally embedded in \hat{M} normal to the singular locus and tiled by hexagons. The plane corresponding to the geodesic g may be homotoped to the plane corresponding to the geodesic h by the indicated translation). This is the only case in which a torus cannot be homotoped to its canonical form by using moves corresponding to the moves in Lemma 5.2. If we wanted to eliminate this source of ambiguity, we would need to replace the 2-manifold by a 2-complex by joining the geodesic boundary components to the geodesics in the other component to which they correspond (this seems an unnecessary complication! we will discuss later several possible ways for resolving this ambiguity and finding each homotopy class of tori exactly once). However, none of this causes us any difficulties in constructing a repeatable geodesic segment in \mathbb{G} in the case where \hat{q} contains no real inflections.

Thus, whether \hat{q} contains real inflections or not, we have a repeatable geodesic segment in \mathbb{G} . \square

At this point it is clear what our algorithm must do: it must find all the repeatable subpaths of paths in \mathbb{F} . Note that to find such subpaths, it suffices to examine paths in \mathbb{F}

Figure 5.6

of length less than $6k + 1$ where k is the number of edges in the flat graph, since that is the number of distinct states of the automaton. If our automaton ever repeats a state, it has found an injectively immersed torus in M . In particular, we have an *a priori* upper bound on the number of paths in \mathbb{F} that we need to consider, which is a crucial part of Theorem 6.1, which we have now proven and will proceed to discuss.

Section 6 - Description of the Algorithm.

At this point, we have proven

Theorem 6.1. *Let M be a branched cover of S^3 , branched over the figure eight knot, such that all branching indices of M are ≥ 3 . Then, there exists an effective algorithm for deciding whether or not M admits any π_1 -injectively immersed tori, and for constructing any that exist.*

A few words of summary are in order here to describe precisely how the algorithm is implemented. The algorithm is presented with the branched cover in the form of a monodromy representations from $\pi_1(S^3 - K) \rightarrow S_d$ (specifically, two permutations corresponding to the two fundamental group generators a and c in the presentation derived in Fig. 2.3). From these two permutations, the set $G(M)$ is calculated using the elemental strip condition derived in section 4. The flat graph is then constructed from $G(M)$ and the monodromy and the path length bound is calculated. Then, the paths in \mathbb{F} of length less than the bound are enumerated and the repeatable subpaths are listed. Each one of these corresponds to an injectively immersed torus in M and all such tori correspond to some repeatable subpath (if we allow these “short” subpaths to be combined in the obvious way into “long” repeatable subpaths which exceed the path length bound). If there are no repeatable subpaths in \mathbb{F} , then M is atoroidal.

It is of course of interest to decide which of these tori are actually embedded (although this is of somewhat reduced importance in light of the results of Gabai and Casson mentioned in section 0). It is known that a smooth minimal torus homotopic to an embedded torus is either embedded or double covers an embedded one-sided surface [FHS], so that

the “embeddedness” of a given homotopy class of tori can be checked from its minimal representative. We will show an example of this in section 7. Note, however, that we must be allowed to adjust our totally geodesic maps in the vicinity of the cone locus (essentially we just take a smoothing and use the lemmas of section 3) since these tori are only piecewise smooth. We remark that whether or not a given repeatable path corresponds to an embedded torus is also algorithmically decidable - one first decomposes the torus into pieces of the form detailed in Fig. 4.1, then checks to see that any that intersect do so only along their edges, and finally that these edges may be “pulled apart” by checking the types of turns involved at an edge.

The algorithm as we have described it so far is not perhaps as sharp as one might like, since the correspondence between tori and repeatable subpaths is many-to-many. However, the ambiguities are readily resolvable in the sense that the algorithm may be “post-processed” to obtain a unique torus in each homotopy class. We may do this by declaring two repeatable subpaths to be equivalent if they differ by a “Lemma 5.2 move”, a circular shift, the action of $\varphi(ba^{-1}c^{-1}ad)$ or a translation (where this includes the jump between components illustrated in Fig. 5.6). Then, each homotopy class of tori corresponds to an equivalence class of repeatable subpaths under this equivalence relation.

It should be mentioned that in [Jo2], another algorithm was described for finding these tori when they exist, but that algorithm did not terminate when the branched cover was atoroidal. The present algorithm not only terminates for any branched cover, but is much more efficient than the other algorithm in any event – this algorithm can frequently be applied “by hand” as we will see in the next section, whereas the other algorithm was really only suitable for machines and gave no real insight as to the overall structure of M in the way that the associated 2-manifold and flat graph do.

Section 7 - Examples.

We will conclude with two examples of covers to which this method is applicable, taken from the list (exhaustive up to degree ten) in [He]. Hempel’s presentation of the figure eight knot group, $\langle \mu, x, y : \mu x \mu^{-1} = x y^{-1}, \mu y \mu^{-1} = y^2 x^{-1} \rangle$, is quite similar to the one obtained in Fig. 2.3, the isomorphism being given by

$$\begin{aligned} a &\mapsto \mu \\ b &\mapsto x \mu x^{-1} \\ c &\mapsto x \\ d &\mapsto x \mu^{-1} \end{aligned}$$

Our two examples are numbers 43 and 37 in [He] which have branching indices (3,3,4) and (3,7) respectively (call them M_1 and M_2). The monodromy homomorphisms are as follows (the group is generated by x and μ):

$$\begin{aligned} \varphi_1 : x &\mapsto (052497)(1)(3)(68) & \mu &\mapsto (021)(347)(5698) \\ \varphi_2 : x &\mapsto (0356)(1)(2487)(9) & \mu &\mapsto (021)(3758496) \end{aligned}$$

Using our criteria, we compute that M_1 has an injectively immersed torus, while M_2 is atoroidal. Doing these examples explicitly will help clarify the torus-finding algorithm.

First, we must find determine which elemental strips exist. For the first example, we compute that $F(M_1) = \{6, 8\}$, so that acceptable lifts of the parallelogram in Fig. 4.2 exist at \hat{x}_6 and \hat{x}_8 . Now, $\varphi(ba^{-1}c^{-1}ad) = (0\ 1\ 2)(3\ 7\ 4)(5\ 9)(6\ 8)$ and so $G(M_1) = F(M_1)$. Therefore, our graph will have two edges, with heads marked 6 and 8 and tails marked $6\varphi(bd^{-1}) = 5$ and $8\varphi(bd^{-1}) = 9$. Examining $\varphi(a)$ tells us how to label the polygons (vertices) and the resulting flat graph is given in Fig. 7.1. Traversing either edge repeatedly in state 0 gives a repeating geodesic and both of these correspond to the same torus (since $\varphi(bd^{-1})$ interchanges these two edges). Furthermore, these repeating geodesics may be pushed off the cone locus to be disjoint from each other as described in section 6, and thus the self-intersection along the cone locus is not essential.

Figure 7.1

In the second example, atoroidality is clear, since $F(M_2) = \emptyset$, so the graph has no edges (and hence no closed geodesics!).

By doing a bit more work with Fig. 4.2, one can show that M_1 is a Haken manifold – the torus found consists of two pieces of each type in Fig. 4.1, and they occur in the following fundamental domains: domain 0 contains pieces 2, 3, and 6; domain 1 contains pieces 4 and 8; domain 2 contains piece 7; domain 3 contains pieces 4 and 8; domain 4 contains pieces 2, 3, and 6; domain 5 contains piece 5; domain 6 contains piece 1; domain 7 contains piece 7; domain 8 contains piece 1; domain 9 contains piece 9. In particular, no domain contains intersecting pieces (except for intersections along the cone locus which we have observed are inessential), thus the torus is homotopic to an incompressible surface. So, this manifold is Haken (irreducibility comes from Lemma 1.2). Hempel asserts that its fundamental group is isomorphic to that of the union of a trefoil complement and the twisted I -bundle over the Klein bottle glued along their boundary, thus (since Haken manifolds of the same homotopy type are homeomorphic [Wa]) this manifold is homeomorphic to the union of two Seifert-fibered spaces glued along an incompressible torus. In [Jo2], conditions are derived which predict that if this manifold decomposes into two Seifert-fibered spaces that they must have \mathbb{E}^3 geometry and $\mathbb{H}^2 \times \mathbb{R}$ geometry and this is indeed the case for these two.

The second manifold, M_2 is atoroidal and irreducible, and thus is believed (but not asserted) to be hyperbolic.

Note also that these methods apply to any link in S^3 which is the singular set of a Euclidean orbifold, the only real differences being in the details of the construction of the associated 2-manifold and flat graph as well as differences in the elemental strip condition. There are 9 such links [Du] (not counting the figure eight knot) and these will be treated in a subsequent paper [Jo1]. In some sense, the results obtained for the other links are stronger than for the figure eight knot since most of the other links have 2-fold branching rather than 3-fold branching in their orbifold structure (thus, the requirements on the branching indices are less restrictive).

REFERENCES

- [Du] William D. Dunbar, *Geometric orbifolds*, preprint.
- [Ga] David Gabai, *Convergence groups are Fuchsian groups*, preprint.
- [FHS] M. Freedman, J. Hass and P. Scott, *Least area incompressible surfaces in 3-manifolds*, Invent. Math. **71** (1983), 609–642.
- [He] John Hempel, *The lattice of branched covers over the figure eight knot*, Topology and its Applications **34** (1990), 183–201.
- [HLM] H.M. Hilden, M.T. Lozano, and J.M. Montesinos, *On knots that are universal*, Topology **24** (1985), 499–504.
- [Ho] Craig Hodgson, *Geometric structures on 3-dimensional orbifolds: notes on Thurston's proof*, preprint.
- [J-S] William H. Jaco and Peter B. Shalen, *Seifert fibered spaces in 3-manifolds*, Mem. AMS **21** (1979).
- [Jh] Klaus I. Johannson, *Homotopy equivalences of 3-manifolds with boundary*, Lect. Notes Math. **761** (1979).
- [Jo1] Kerry N. Jones, *Injectively immersed tori in branched covers over Euclidean links*, in preparation.
- [Jo2] ———, *Cone manifolds in 3-dimensional topology: applications to branched covers*, thesis, Rice University (1990).
- [M] Geoff Mess, *Centers of 3-manifold groups and groups which are coarse quasi-isometric to planes*, preprint.
- [Sc] Peter Scott, *A new proof of the annulus and torus theorems*, American J. Math. **2** (1978), 241–277.
- [Sp] Michael Spivak, *A Comprehensive Introduction to Differential Geometry, Vol. IV*, Publish or Perish, 1970.
- [SY] R. Schoen and Shing-Tung Yau, *Existence of incompressible minimal surfaces and the topology of three dimensional manifolds with non-negative scalar curvature*, Ann. of Math. **110** (1979), 127–142.
- [Wa] Friedhelm Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Ann. of Math. **87** (1968), 56–88.