

STUDYING LINKS VIA CLOSED BRAIDS III: CLASSIFYING LINKS WHICH ARE CLOSED 3-BRAIDS

JOAN S. BIRMAN AND WILLIAM W. MENASCO

A complete solution is given to the classification problem for oriented links which are closed three-braids. The Classification Theorem asserts that, up to a finite list of exceptional cases, links which can be represented by closed 3-braids are represented by a unique conjugacy class in the group of 3-braids. The exceptional cases are the expected ones (links of braid index 1 and 2) and an unexpected infinite family of invertible links, each member of which has two 3-braid axes. The two axes correspond to diagrams which are related by “braid-preserving flypes”.

An algorithm is given which begins with an arbitrary closed 3-braid (or alternatively any link diagram with 3 Seifert circles), and converts it into a normal form which characterizes its oriented link type in oriented 3-space. One can decide from the normal form whether the link is prime or composite, split or irreducible, amphicheiral and or invertible. One can decide if the braid index is 3, 2 or 1. Using related results of P. J. Xu, one may determine the genus and construct a surface of maximum Euler characteristic with boundary the given link.

It is proved that the stabilization index of a link which is represented by a closed 3-braid is ≤ 1 , i.e. any two 3-braid representatives of the same link type become conjugate after a single stabilization to B_4 .

1. Introduction. This manuscript is the third in a series of six papers in which the authors have been studying the closed braid representatives of links. The other five are referenced here as [B-M I, II, IV-VI] In [B-M, I] the authors prove that, with an appropriate definition of the complexity of a closed braid representative of a link, there are finitely many conjugacy classes of braids of minimum complexity. In this paper we work out in detail the first non-trivial example of this “finiteness theorem”: we prove that for links of braid index $n \leq 3$ there are at most two conjugacy classes of minimum complexity. We then use this result to give a complete numerical invariant of link type, and an algorithm to find it, starting with an arbitrary closed 3-braid.

The complexity function which was used in [B-M, I] relates to incompressible surfaces which are bounded by the link. These surfaces have maximum Euler characteristic, among all oriented surfaces with

boundary the given oriented link. It was shown in [Be] and in [B-M, II] that when the link has braid index ≤ 3 these surfaces are especially simple. As a result, the combinatorial data which was introduced in [B-M, I] to describe the surfaces reduces to a cyclic word in a special set of generators for the braid group B_3 . The algorithm which we give then not only solves the link problem for links of braid index 3, it also (with the help of results in [X]) provides a surface of maximal Euler characteristic with boundary the given link. The long range goal of the series of manuscripts [B-M, I-VI] is to generalize these results to arbitrary braid index n .

Let \mathcal{L} be an oriented link type in oriented S^3 , let \mathbf{L} be a representative of \mathcal{L} , and let \mathbf{A} be an oriented circle in $S^3 - \mathbf{L}$. We say that \mathbf{L} is a *closed n -braid* with braid axis \mathbf{A} if there is a choice $\mathbf{H} = \{\mathbf{H}_\theta; \theta \in [0, 2\pi]\}$ of a *fibration* \mathbf{H} of the open solid torus $S^3 - \mathbf{A}$ by meridian discs \mathbf{H}_θ such that \mathbf{L} intersects each \mathbf{H}_θ transversally. The number n is the number of points in $\mathbf{L} \cap \mathbf{H}_\theta$. In 1923 James Alexander proved in [A1] that every link could be so-represented, in many ways and for many different values of n . The connection between braids and links was further developed by Emil Artin. In his 1925 manuscript [Ar] he introduced the *n -string braid group* B_n ($n = 1, 2, 3, \dots$) and uncovered its structure. In 1935 A. A. Markov studied the disjoint union B_∞ of all the B_n 's, describing in [Ma] the equivalence relation in B_∞ which relates all of the open braids whose closures have the same oriented link type. Two elementary equivalences are needed: conjugacy in B_n and a stabilization move which changes n . In 1968 Garside [Ga] solved the conjugacy problem in B_n . This was a highly non-trivial step in the translation of Markov's equivalence relation into an algorithmic solution to the link problem.

The role of "stabilization" in Markov's equivalence relation (and indeed in many other areas of low dimensional topology such as Heegaard theory, bridge representations of links, Kirby's Calculus, and Waldhausen's study of the Heegaard splittings of S^3) has been somewhat mysterious. The present investigation began with an attempt to understand what is accomplished when one stabilizes, in the situation of Markov's theorem. In the papers [B-M, I-VI] and subsequent ones we hope to establish a version of Markov's theorem which does not involve stabilization, i.e. in which one studies the braid representatives of a link in B_n , for fixed n , rather than in the disjoint union B_∞ of all of the braid groups. In this manuscript we do exactly that for the first non-trivial case, $n = 3$. The classification of links which

are closed 3-braids follows immediately, as there are several known algorithmic solutions to the conjugacy problem in B_3 ([SC], [G], [X]), each yielding unique normal forms for conjugacy classes.

By the work of Yamada [Y] links which are defined by closed n -braids have an alternative interpretation as links which are defined by diagrams with n Seifert circles. A computer program is now available [F] which begins with an arbitrary link diagram with n Seifert circles and produces from it a closed n -braid representative (in fact, for $n = 3$ the program gives a 3-braid which is defined by the special generators a_1, a_2, a_3 which are introduced in this paper). We also have a second program which yields normal forms for conjugacy classes in B_3 . Thus our results are very concrete.

Two other special cases of Markov's theorem without stabilization are proved in [B-M, IV] (which is a study of the n -braid representatives of split links and composite links) and [B-M, V] (where the n -braid representatives of the r -component unlink are studied for fixed $n \geq r$). The pathologies which are investigated in the three papers [B-M, III], [B-M, IV] and [B-M, V], all come into play if one attempts to generalize the results of this paper to higher braid index.

To state our result, let σ_1 and σ_2 be the standard elementary braids which generate the 3-string braid group B_3 . The central result in this paper may be stated in the following way:

THE CLASSIFICATION THEOREM (Version 1). *A link \mathcal{L} which is represented by a closed 3-braid admits a unique conjugacy class of 3-braid representatives, with the following exceptions:*

- (i) \mathcal{L} is the unknot, which has three conjugacy classes of 3-braid representatives, namely the classes of $\sigma_1\sigma_2$, $\sigma_1^{-1}\sigma_2^{-1}$ and $\sigma_1\sigma_2^{-1}$.
- (ii) \mathcal{L} is a type $(2, k)$ torus link, $k \neq \pm 1$, which has two conjugacy classes of 3-braid representatives, namely the classes of $\sigma_1^k\sigma_2$ and $\sigma_1^k\sigma_2^{-1}$.
- (iii) \mathcal{L} is one of a special class of links of braid index 3 which have 3-braid representatives which admit "braid-preserving flypes". These links have at most two conjugacy classes of 3-braid representatives, namely the classes of $\sigma_1^p\sigma_2^q\sigma_1^r\sigma_2^\delta$ and $\sigma_1^p\sigma_2^\delta\sigma_1^r\sigma_2^q$, where p, q, r are distinct integers having absolute value at least 2 and where $\delta = \pm 1$.

The exceptional cases (i) and (ii) are the expected ones, but (iii) was not previously recognized except in special cases (see [Mu]). It is




FIGURE 1.1: A flype

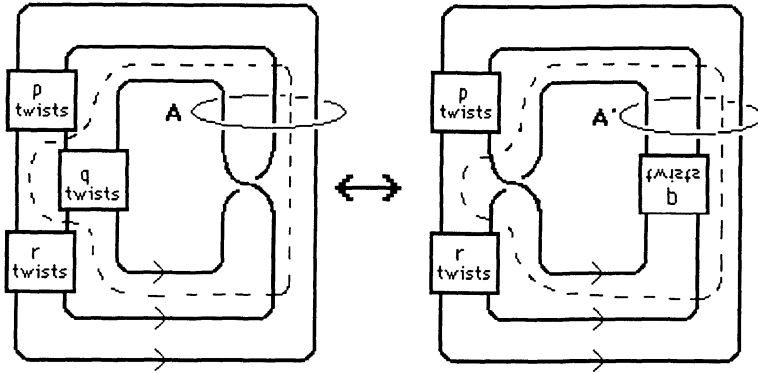



FIGURE 1.2: A braid-preserving flype

an infinite family of counterexamples to the long-standing conjecture that the map from conjugacy classes in B_3 to oriented link types was injective on links of braid index 3. The examples in (iii) are especially important because they generalize to related examples for every braid index $n \geq 3$, so we discuss them briefly.

The left picture in Figure 1.1 shows part of a link diagram (enclosed in a dotted oval). The right picture shows a different projection of the same part of the diagram. The left and right pictures differ by a *flype*. It is easy to see that diagrams which contain regions which are related by flypes, but are identical everywhere else, determine the same oriented link type. Figure 1.2 shows a pair of closed 3-braid projections of the same link, related by a *braid-preserving flype*. The flype occurs inside the disc which is bounded by the dotted curve. During the flype the link L cuts through the braid axis A . Notice that the defining open braids have the form of the examples given in (iii) above.

Figure 1.3 shows the same phenomena from a different point of view, exhibiting the braid axis A' as it would appear threaded into the




Inequivalent axes for
 n -braid structure on
 $L, n=n_1+n_2+n_3$

FIGURE 1.3

projection shown in the left picture of Figure 1.2. At the same time, Figure 1.3 shows how flypes generalize to produce similar examples for every n . The strands are labeled with integers n_1, n_2, n_3 , indicating that the single strand which is illustrated is to be replaced by $n_i \geq 1$ parallel strands. Thus the braid which is labeled P in Figure 1.3 is a braid on $(n_1 + n_2)$ strands, and so-forth. The generalized braid in Figure 1.3 has a very special form, even though P, Q and R are arbitrary braids on the appropriate numbers of strands.

We wish to answer the question: when can \mathcal{L} be represented by more than one conjugacy class in B_3 ? We choose a representative L of \mathcal{L} , and study the related geometric question: when can L have two axes, A and A' for 3-braid structures, with A' not isotopic to A in $S^3 - L$? We wish to show that this only happens under very special circumstances. Since it is known exactly how many axes are possible in the special cases when the braid index of \mathcal{L} is 2 or 1, we may assume (and shall do so from now on) that \mathcal{L} has braid index 3. It will also be convenient to make the assumption that \mathcal{L} is prime and non-split, since composite and split links complicate the proof even though it is well known (via the work of Murasugi in [Mu] and Morton in [Mo, 1], or see [B-M, IV]) that composite and split links of braid index 3 have unique conjugacy classes of 3-braid representatives. Therefore, from now on, we assume that \mathcal{L} is prime, nonsplit and has braid index 3.

The first step in our work is to attach extra structure to our braid representatives. Let S be an oriented surface, with boundary an



Bennequin surfaces

FIGURE 1.4

oriented link \mathbf{K} of braid index n . We call \mathbf{S} a *Bennequin surface* (see Figure 1.4) if the following conditions hold:

(i) \mathbf{S} has maximal Euler characteristic $\chi(\mathbf{S})$, among all connected oriented surfaces with boundary the oriented link \mathbf{K} .

(ii) $\mathbf{K} = \partial\mathbf{S}$ is a closed n -braid relative to the braid axis \mathbf{A} and to some choice of a fibration \mathbf{H} of the open solid torus $S^3 - \mathbf{A}$ by meridian discs.

(iii) \mathbf{S} has a decomposition as a union of n discs, each pierced once by \mathbf{A} , joined up by half-twisted bands, each of which has exactly one

saddlepoint tangency with a fiber of \mathbf{H} . Each disc is radially foliated by its intersections with fibers of \mathbf{H} .

Bennequin surfaces were introduced (implicitly) in Proposition 3 of [Be]. They differ from the obvious Seifert surface \mathbf{S}' spanned by a closed n -braid in the following way: Both \mathbf{S} and \mathbf{S}' have \mathbf{K} as boundary, and both are a union of n discs $\mathbf{D}_1, \dots, \mathbf{D}_n$ joined up by half-twisted bands; however the bands in \mathbf{S}' are usually limited to ones which join \mathbf{D}_i to \mathbf{D}_{i+1} for each $i = 1, 2, \dots, n-1$, whereas in a Bennequin surface \mathbf{S} one may have half-twisted bands joining any pair of discs \mathbf{D}_j and \mathbf{D}_k .

An arbitrary link \mathcal{L} of braid index n may or may not have a Bennequin spanning surface, but links of braid index 3 always do. Even more it is proved in [B-M, II] that the following holds:

BENNEQUIN'S THEOREM (see [B-M, II]). *Let \mathcal{L} be prime and non-split and have braid index 3. Let \mathbf{L} be a representative which is a closed 3-braid relative to the axis \mathbf{A} . Let $\mathbf{H} = \{H_\theta; \theta \in [0, 2\pi]\}$ be a fibration of the open solid torus $S^3 - \mathbf{A}$ by meridian discs. Then any connected, oriented surface \mathbf{F}' of maximal Euler characteristic with boundary \mathbf{L} is isotopic to a Bennequin surface \mathbf{F} with the same boundary, relative to \mathbf{H} .*

A Bennequin surface for a closed 3-braid has a natural decomposition as a union of three horizontal discs joined up by half-twisted bands. We label the discs $\mathbf{D}_1, \mathbf{D}_2, \mathbf{D}_3$ and let a_i denote a half-twisted band which joins \mathbf{D}_i to $\mathbf{D}_{i+1 \pmod{3}}$, in the standard disc-band decomposition of \mathbf{F} relative to \mathbf{H} . See Figure 1.5 (next page). Then \mathbf{F} and $\mathbf{L} = \partial\mathbf{F}$ can be described by a cyclic word:

$$\mathbf{W} = (a_{\mu_1})^{p_1} (a_{\mu_2})^{p_2} \cdots (a_{\mu_k})^{p_k}$$

in the symbols a_1, a_2, a_3 and their inverses, which records the sequence of half-twisted bands encountered as the polar angle varies over the interval $[0, 2\pi]$. This word describes the surface \mathbf{F} as an embedded surface in $S^3 - \mathbf{A}$. It is a shortest such word because $\chi(\mathbf{F})$ is maximal. This is one of the very nice features of Bennequin surfaces—the braid word which describes \mathbf{L} also describes a very nice spanning surface \mathbf{F} .

By hypothesis, our link \mathcal{L} has *two* axes \mathbf{A} and \mathbf{A}' for 3-braid structures, with associated fibrations $\mathbf{H} = \{\mathbf{H}_\theta; \theta \in [0, 2\pi]\}$ and $\mathbf{H}' = \{\mathbf{H}'_\phi; \phi \in [0, 2\pi]\}$. If we arrange things so that $\mathbf{L} = \partial\mathbf{F}$, where

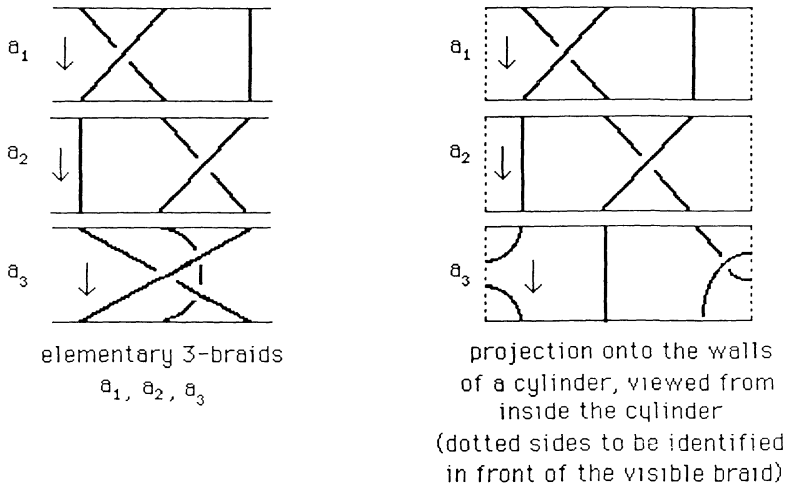


FIGURE 1.5

F is a Bennequin surface relative to H , then Bennequin's Theorem tells us there is an isotopy h_s of S^3 , $s \in [0, 1]$, with $h_0 = \text{identity}$ such that $h_1(F)$ is a Bennequin surface relative to H' . Therefore, by modifying the fibers of H as we modify F , replacing each h_θ by $h_1(H_\theta)$, we obtain as an immediate corollary:

COROLLARY 1.1. *We may assume that F is simultaneously a Bennequin surface relative to both H and H' . Moreover, there is a second shortest word $W' = W'(L)$, necessarily of the same length as $W = W(L)$, in generators b_1, b_2, b_3 which describes the Bennequin disc-band decomposition of F relative to H' .*

We now describe the strategy of our proof of the Classification Theorem, which is contained in §§2–6. Motivated by the results in an early version of this paper, P. J. Xu has studied cyclically symmetric presentations of B_3 using the generators a_1, a_2, a_3 . She gives in [X] an algorithm for finding a shortest representative for each conjugacy class, and she also investigates the collection of all shortest words in each conjugacy class. Using one of her results (see Lemma 2.2) we show in §2 that we may choose among all Bennequin surfaces with a given 3-braid as boundary ones with additional structure. Their defining words W and W' have two properties: first, they have the shortest possible syllable length, and second they use all three generators. We call them “admissible” words.

In §3 we introduce additional geometric structure in the form of a family Ω of essential discs which we call *digons* which “plug up the

holes” in the standard disc-band decomposition of F . The foliation of the digons by their arcs of intersection with fibers of H and H' will be the key to our proof. The standard axis A is disjoint from Ω and the foliation of each digon which is induced by the standard fibration H is free of singularities in the interiors of the digons. The non-standard axis A' , however, may pierce Ω repeatedly, and we need to understand the relationship of A' and the fibers of H' to Ω . We prove that S^3 split along $F \cup \Omega$ is a union of three 3-balls B_1, B_2, B_3 and a solid torus T . We show that L has a natural “projection” onto ∂T . We show that W can be determined from the projection of L on ∂T . We then prove that A' , like A , may be assumed to be disjoint from the union of all the digons.

In §4 we study the foliation of the digons which is induced by the fibration H' of $S^3 - A'$. We show that its singularities are very special, and begin to show (see Corollary 4.9) how knowledge of the foliation enables us to draw conclusions about W .

In §5 we show that we may assume that A' is inside T .

In §6 we use all of this structure to prove the Classification Theorem.

In §7 the Classification Theorem is translated into an algorithmic solution to the link problem for links which are defined by closed 3-braids. The algorithm depends directly upon known solutions to the conjugacy problem in B_3 . In addition to giving a unique and complete invariant for links defined by closed 3-braids, the algorithm answers other questions: Is \mathcal{L} invertible or amphicheiral or both? Is the braid index 3, 2 or 1? Is \mathcal{L} split or composite? With the help of [X], one can also answer the question: what is the genus of \mathcal{L} ? We also prove that the stabilization index for two 3-braid representatives of the same link is one.

Knowing a complete invariant for a class of links such as links of braid index ≤ 3 , it is natural to ask how this invariant compares to “Jones invariants” i.e. the Jones polynomial [J] and its generalizations. We call a link type invariant a *Jones invariant* if it is constructed from traces in matrix representations of the braid group. The key property is that the traces in question change in a very simple way when one changes the defining braid by the moves which generate Markov’s equivalence relation. It is to be expected that Markov equivalence will imply the existence of trace identities in any representation of B_3 which admits a Jones invariant. We discuss this question briefly at the end of §7. We prove that, for any $N > 0$, there exist infinitely many sets, each of which contains N distinct links of braid index 3 which

share the same 2-variable Jones polynomial. It would be important to know whether any finite collection of Jones invariants suffices to classify links of braid index 3.

Section 8 contains a brief discussion of the closed braid representatives of composite links of braid index 3. They have unique conjugacy classes of 3-braid representatives; however our proof would have been more complicated if we had included them in the main body of the proof. That was a hint to us of pathology which needed to be understood. That pathology will be investigated in detail in the next paper in this series, [B-M, IV], which studies the closed braid representatives of split and composite links.

REMARKS. (1) While links defined by closed 3-braids are special, they are also in many ways a rich class of links. Their “crossing numbers” are unbounded, unlike the links of ≤ 13 crossings which have served as a principal source of examples for knot theorists up to now (e.g. see the Appendix to [Ro]). “Most” of them are non-amphicheiral and non-invertible (we will make this precise in §7). Their geometry includes many examples of hyperbolic link spaces. Their Alexander and Jones and Kauffman polynomials are special (see [J], [Ta], [Bi, 2]), yet they are not fully understood at this writing, suggesting interesting new questions for investigation.

(2) The Classification Theorem is striking for its simplicity, but our proof is long. This raised a question which has troubled us repeatedly during the four years when the work reported on here was in progress: Knowing the correct statement of the Classification Theorem, it seems possible that there is a very easy proof. For, if one has a finite collection of link type invariants which sufficed to distinguish link type for the special class of links which are defined by closed 3-braids, then all one would need to do would be to show that braids in distinct conjugacy classes defined links with distinct invariants. Unfortunately we do not know of any such collection. It would be interesting to know if there is such a collection, and in particular if there is one which is adapted to the normal form in this paper.

Acknowledgment. We thank the Mathematics Department at Rutgers University Newark for its hospitality in hosting an extended seminar in which we presented some of the details in this paper. In particular, we thank Lee Mosher, Ulrich Oertel, Mark Feighn, Jane Gilman, Kay Tatsuoka, Rollie Trapp and Pei-Jun Xu. Their insights and questions were invaluable to us. We also wish to thank the referee, who read our manuscript with great care and expertise.

2. Admissible words. In this section we begin the new work needed for our proof of the Classification Theorem. Let \mathcal{L} , \mathbf{L} , \mathbf{F} , \mathbf{A} , \mathbf{A}' , \mathbf{W} , \mathbf{W}' be as in §1. We say that \mathbf{W} is *admissible* if \mathbf{W} is a shortest representative of its conjugacy class in B_3 , also if among all shortest words it has shortest *syllable length* k , and finally if it involves all three generators a_1 , a_2 , a_3 non-trivially.

LEMMA 2.1. *Given any closed 3-braid \mathbf{L} which represents a prime, non-split link of braid index 3, there is a Bennequin surface \mathbf{F} with $\partial\mathbf{F} = \mathbf{L}$ such that $\mathbf{W} = W(\mathbf{L})$ is admissible.*

Proof. We may assume that \mathbf{L} is the boundary of a Bennequin surface \mathbf{F} with defining word \mathbf{W} . From this it follows that \mathbf{W} is a shortest representative of the conjugacy class of \mathbf{L} , for if not the Euler characteristic $\chi(\mathbf{F})$ would not be maximal. Among all Bennequin surfaces with boundary \mathbf{L} , assume we have chosen ours so that the syllable length of \mathbf{W} is also minimal. Then either \mathbf{W} uses all three generators, in which case we are done, or (up to conjugation by a rotation which cyclically permutes the generators) we may assume that $\mathbf{W} = (a_1)^{p_1}(a_2)^{p_2}(a_1)^{p_3} \cdots (a_2)^{p_s}$, where each $p_i \neq 0$ and $s \geq 1$. If $s = 1$ then \mathcal{L} is the connected sum of a type $(2, p_1)$ and a type $(2, p_2)$ torus link, so $s \geq 2$. Thus \mathbf{W} contains a subword

$$\mathbf{U} = (a_1)^{p_1}(a_2)^{p_2}(a_1)^{p_3}(a_2)^{p_4}$$

of syllable length 4.

Since the assertion of Lemma 2.1 can be proved for any choice of orientations on \mathcal{L} and on 3-space, we may if necessary replace \mathbf{W} by its inverse for the purposes of the proof, so we may assume without loss of generality that $p_3 > 0$. Now observe the following relation holds:

$$(2.1) \quad a_2a_1 = a_3a_2 = a_1a_3.$$

But then, we may apply the relation $a_2a_1 = a_1a_3$ repeatedly to the subword $(a_2)^{p_2}(a_1)$ of \mathbf{U} , replacing it by $(a_1)(a_3)^{p_2}$. Thus \mathbf{U} has the equivalent form $(a_1)^{p_1+1}(a_3)^{p_2}(a_1)^{p_3-1}(a_2)^{p_4}$. Since \mathbf{W} is by hypothesis a shortest word and has shortest syllable length, each of the four syllables in the new expression for \mathbf{U} must be non-trivial, and the lemma is proved. \square

For future use we record the following consequences of the relations (2.1):

$$(2.2) \quad (a_i)^\delta (a_{i\pm 1})^\delta (a_i)^r = (a_{i\pm 1})^r (a_i)^\delta (a_{i\pm 1})^\delta, \quad \delta = \pm 1, \text{ any } r \neq 0.$$

$$(2.3) \quad a_i = (a_{i+1})^{-1} (a_{i-1}) (a_{i+1})$$

$$(2.4) \quad (a_i)^p (a_{i-1}) = (a_{i-1}) (a_{i+1})^p, \quad \text{any } p,$$

$$(2.5) \quad a_{i-1} a_i^{-1} = a_i^{-1} a_{i+1}$$

where $i = 1, 2, 3$ and subscripts are always understood to be mod 3.

P. J. Xu has taken a new look at the conjugacy problem in B_3 and has found a constructive procedure to determine all admissible words, starting with an arbitrary representative of the conjugacy class of \mathbf{L} . She has proved the following important fact:

LEMMA 2.2 [X]. *Any two shortest words \mathbf{W}' , \mathbf{W}'' which represent the same conjugacy class in B_3 are related by a finite sequence of shortest words $\mathbf{W}' = \mathbf{W}_0, \dots, \mathbf{W}_t = \mathbf{W}''$ of shortest words, such that each \mathbf{W}_i differs from \mathbf{W}_{i-1} by a single application of relations (2.1).*

By Corollary 1.1 the surface \mathbf{F} is a Bennequin surface with respect to both \mathbf{H} and \mathbf{H}' . Therefore we can assign to it a second defining word \mathbf{W}' in elementary braids b_1, b_2, b_3 which describes the second disc-band decomposition of \mathbf{F} .

COROLLARY 2.3. *We may assume that $\mathbf{L} = \partial\mathbf{F}$, where \mathbf{F} is a Bennequin surface with defining words \mathbf{W} (with respect to \mathbf{H}) and \mathbf{W}' (with respect to \mathbf{H}'), such that both \mathbf{W} and \mathbf{W}' are admissible.*

Proof. By Corollary 1.1 we may assume that \mathbf{F} is a Bennequin surface relative to both \mathbf{H} and \mathbf{H}' , with defining words \mathbf{W} and \mathbf{W}' . By Lemma 2.1, we may assume that \mathbf{W} is admissible. If \mathbf{W}' is too, we are done. If not, we know from Lemma 2.1 that there is a Bennequin surface relative to \mathbf{H}' which has an admissible defining word \mathbf{W}'' . By Xu's Lemma 2.2 the words \mathbf{W}' and \mathbf{W}'' are related via a sequence of shortest words, with adjacent members in the sequence differing by a single application of relations (2.1). Since each of the relations in (2.1) can be realized geometrically by an isotopy of \mathbf{F} (a simple picture will show this), there is an isotopy h_s of S^3 taking \mathbf{F} to a new Bennequin surface $h_1(\mathbf{F})$ which has the defining word \mathbf{W}'' relative to




FIGURE 2.1

H' . As in the proof of Corollary 1.1, we may then replace each fiber H_θ of H by $h_1(H_\theta)$ in order to preserve the relationship of F to H , to obtain a new surface with the desired properties. Thus we could have assumed to begin with that both W and W' were admissible. \square


We record and review notation which will be used throughout this paper. Fibers in the standard (resp. non-standard) fibration will be parametrized by polar angle θ (resp. ϕ). We will depict fibers of H and H' so that their boundaries (the braid axes A and A') are oriented clockwise, as in Figure 2.1. With this convention the polar angle functions increase clockwise. The symbols $H_\theta, p_j, F_{i\theta}, \dots$, (resp. $H'_\phi, p'_j, F_{i\phi}, \dots$) will be used to indicate a non-singular fiber, a point where the axis pierces F , an arc of intersection of the fiber with F , ... in the standard (resp. non-standard) case. See Figure 2.1. Braids will always be oriented top to bottom in our pictures, as in Figure 1.5. The axes A and A' are oriented by the right hand rule, relative to the given orientation on L . The outward-drawn normal to F is also determined by the right hand rule relative to the given orientation on L . Notice that with this convention the axes A and A' pass from F^- to F^+ as they pierce F . A tangency of a fiber of H or

H' is a positive (resp. negative) singularity if the sense of the oriented normal bundle to F at the point of tangency agrees (resp. disagrees) with the direction of increasing θ or ϕ . With this convention the elementary braids depicted in Figure 1.5 are positive.


We will regard the surface F and its boundary, the link L as being defined by two words: a standard word W in the elementary braids a_1, a_2, a_3 and their inverses, and a non-standard word $W' = W'(b_1, b_2, b_3)$ in elementary braids b_1, b_2, b_3 . There is a homeomorphism h of S^3 which takes A' onto A , fibers of H' onto fibers of H , and L onto a closed braid $h(L)$ about the axis A . The braid $h(L)$ will be defined, up to conjugacy, by the word $X = W'(a_1, a_2, a_3)$. Since A' is not isotopic to A in $S^3 - L$, this implies that W and X represent distinct conjugacy classes in B_3 . As we will see, most prime, non-split links of braid index 3 do not have two such representatives, and none has more than two.

3. Digons. In this section we will introduce additional structure into $S^3 - L$, beyond that provided by a Bennequin surface F . The additional structure will be in the form of a set Ω of discs called “digons” which are properly embedded in S^3 split along F . We will prove that S^3 split along $F \cup \Omega$ is a union of four components whose closures are 3-balls B_1, B_2, B_3 and a solid torus T . (See Proposition 3.1.) We show that there is a natural and symmetric way to visualize our 3-braid via a link diagram on the boundary ∂T of the solid torus (Proposition 3.2). We prove that the non-standard axis A' may be assumed to be disjoint from all of the digons (Proposition 3.4).

We focus for the moment on the geometry determined by the standard axis A and fibration H . Consider the closure M of the 3-manifold obtained by splitting S^3 open along F , $\partial M = F^+ \cup F^-$, $F^+ \cap F^- = L$. A disc Δ which is properly embedded in M will be called a *digon* if Δ is essential and if $\partial \Delta$ meets L in exactly two points. We now construct a collection $\Omega = \{\Delta_i(k); 1 \leq k \leq r_i, i = 1, 2, 3\}$ of digons. Let $\{D_i, i = 1, 2, 3\}$ be the discs in the standard disc-band decomposition of F and let $\{b_i(k); k = 1, 2, \dots, r_i\}$ be the bands which join discs D_i and D_{i+1} . The construction begins with a rectangle \square which is a model for a typical digon, its boundary being divided into four arcs labeled e_1, e_2, e_3, e_4 in clockwise order about $\partial \square$. We want to attach \square to F along $\partial \square$. To describe the attaching curves it will be convenient to think of the bands as rectangles, the singular leaves in the standard foliation being parallel to the sides and through the center of the band. In Figure 3.1(b) we have




(a) Bands embedded in 3-space



If the half-twist is positive

If the half-twist is negative

(b) Digon attaching curves on $\mathbf{b}_i(k)$



(c) Digon attaching curves on \mathbf{D}_i

FIGURE 3.1


sketched two arcs c_{i1} and c_{i2} on the band $\mathbf{b}_i(k)$. Our rules are:

Attach e_1 to F^+ along \mathbf{D}_i between $\mathbf{b}_i(k)$ and $\mathbf{b}_i(k+1)$, close to \mathbf{L} .

Attach e_2 to $\mathbf{b}_i(k+1)$ along c_{i2} .


Attach e_3 to F^- along \mathbf{D}_{i+1} , between $\mathbf{b}_i(k+1)$ and $\mathbf{b}_i(k)$, close to \mathbf{L} .

Attach e_4 to $\mathbf{b}_i(k)$ along c_{i1} .



Example of an embedded digon

FIGURE 3.2



The decomposition of F into subsets $1', 1'', 2, 3$

FIGURE 3.3

The intersections of the attaching curves with a typical disc are shown in Figure 3.1(c). The visible side is F^+ . Figure 3.1(b) and (c) use solid and dotted arcs to distinguish places where the digons are attached to F^+ from places where they are attached to F^- .

An example is given in Figure 3.2 of a digon as it would appear when F is imbedded in 3-space. Notice that each digon is properly embedded in M and that the digons are pairwise disjoint. In particular, we have attached $\Delta_i(k)$ to D_i along F^+ and to D_{i+1} along F^- , to insure that $\Delta_i(k) \cap \Delta_{i+1}(k)$ is empty.

Figure 3.3 shows a decomposition of the surface F into subsurfaces which are defined via the structure introduced here. We will need it later (see the discussion before Lemma 3.3 and the proof of the lemma).

On the other hand, if we regard the digons as subsets of S^3 rather than as subsets of M , then distinct digons $\Delta_i(k)$ and $\Delta_j(m)$ may intersect, but only in a very special way:

- (i) $\text{Int}(\Delta_i(k)) \cap \Delta_j(m) = \emptyset$.
- (ii) $\partial\Delta_i(k) \cap \partial\Delta_j(m) = \emptyset$ unless $|k - m| = 1$, and both digons pass through the singular point on the band or bands which they share.
- (iii) $\partial\Delta_i(k)$ intersects $\partial\Delta_{i\pm 1}(m)$ in either 0, 1 or 2 points on the disc D_i .

Each digon $\Delta_i(k)$ crosses L twice, once on $b_i(k)$ and once on $b_i(k + 1)$. The crossover points will be called *crossing dots*. The crossing dots on $b_i(k)$ cut off a subarc of L which we have labeled $L_i(k)$ in Figure 3.1(b). There is one of these arcs on each band, the one on the band $b_j(k)$ (resp. $b_i(k + 1)$) having its endpoints on the digons $\Delta_j(k - 1)$ and $\Delta_i(k)$ (resp. $\Delta_i(k)$ and $\Delta_i(k + 1)$). The two




FIGURE 3.4

crossing dots on the digon $\Delta_i(k)$ divide $\partial\Delta_i(k)$ into two sub-arcs, one of which is in F^+ and the other in F^- . The subarc of L which lies between the two crossing dots on $b_i(k)$ is labeled $L_i(k)$. There is one such subarc of L on each band. We call these r_i arcs the *undercrossing arcs* (for reasons which will become clear shortly).

The fibration H induces a *standard foliation* of each digon $\Delta_i(k)$. The fact that F is in a very nice position relative to the axis A and the fibers in H implies that this foliation may be assumed to be free of singularities in the interior of $\Delta_i(k)$ (see Figure 3.4). There will, however, necessarily be two degenerate leaves in the foliation. Notice that the crossing dots are nearby to but do not coincide with the degenerate leaves. A degenerate leaf can be in F^+ or in F^- . One of the four possibilities has been illustrated in Figure 3.4. The undercrossing arcs are of course transverse to the fibers of both H and H' .

PROPOSITION 3.1. *S^3 split along $F \cup \Omega$ has four open components. Their closures are 3-balls B_1, B_2, B_3 and a solid torus T . The boundary $\partial B_i, i = 1, 2, 3$, is made up of r_i digons, a disc subset of F^+ which contains D_{i+1} and a disc subset of F^- which contains D_i . The boundary ∂T is a union of $r = r_1 + r_2 + r_3$ digons and subsets of F^+ and F^- .*

Proof. We study how $F \cup \Omega$ meets a generic fiber H_θ of the standard fibration H . See the top left picture in Figure 3.6 on p. 43, which is the same as Figure 2.1(a), but labeled to stress different aspects of the




FIGURE 3.5

geometry. We have added the intersections of \mathbf{H}_θ with the digons, i.e. the three arcs $\lambda_{1\theta}$, $\lambda_{2\theta}$, $\lambda_{3\theta}$. By our rules for attaching the digons to \mathbf{F} , the arc $\lambda_{i\theta}$ begins on the plus side of $\mathbf{F}_{i\theta}$ and ends on the minus side of $\mathbf{F}_{i+1,\theta}$. Thus, if θ is non-singular, \mathbf{H}_θ split along $\mathbf{H}_\theta \cap (\mathbf{F} \cup \Omega)$ will be four discs, labeled $\mathbf{B}_{1\theta}$, $\mathbf{B}_{2\theta}$, $\mathbf{B}_{3\theta}$ and \mathbf{T}_θ . The 3-balls \mathbf{B}_1 , \mathbf{B}_2 , \mathbf{B}_3 and the solid torus \mathbf{T} are defined to be the regions swept out by $\mathbf{B}_{1\theta}$, $\mathbf{B}_{2\theta}$, $\mathbf{B}_{3\theta}$ and \mathbf{T}_θ as θ is varied from θ_1 to $\theta_1 + 2\pi$.

The only possible complication in the above description is that as θ is varied we will pass through singular fibers and crossing dots, so we need to be sure our four discs are always well-defined. Some preliminary remarks are in order. See Figure 3.5. Let S_1 be a surface and $S_2(\phi)$ a family of surfaces, with S_1 transverse to $S_2(\phi)$ except at $\phi = \phi_0$, where a saddle point tangency occurs. Then there are arcs x and y in $S_1 \cap S_2(\phi)$ which coalesce as $\phi \rightarrow \phi_0$. We show in Figure 3.5 a pair of *joining arcs* $\alpha(\phi)$ and $\beta(\phi)$ and a disc $\Delta(\phi)$ which the joining arcs cobound. As $\phi \rightarrow \phi_0$ the joining arcs and the disc shrink to points. The joining arcs are a convenient way to describe the future singularity.

Now see the eight snapshots in Figure 3.6. We have chosen to illustrate a positive singularity of type a_2 (which occurs at $\theta = \theta_5$). Our conventions dictate that for a *positive* singularity between the disc \mathbf{D}_2 and \mathbf{D}_3 in the standard disc-band decomposition of \mathbf{F} the joining arc β (illustrated at θ_1 as a dotted arc, but omitted from succeeding pictures) runs between the *negative* sides of $\mathbf{F}_{2\theta}$ and $\mathbf{F}_{3\theta}$. As we approach the instant of the singularity, this arc shrinks to a point. The arc $\lambda_{2\theta}$ will also shrink to a point as we pass through the band, because the attaching curves for the digons run through the singularities of \mathbf{F} . Therefore, just before the instant of degeneracy the endpoint of $\lambda_{2\theta}$ which is on \mathbf{F}^+ must “migrate” from the $+$ to the $-$

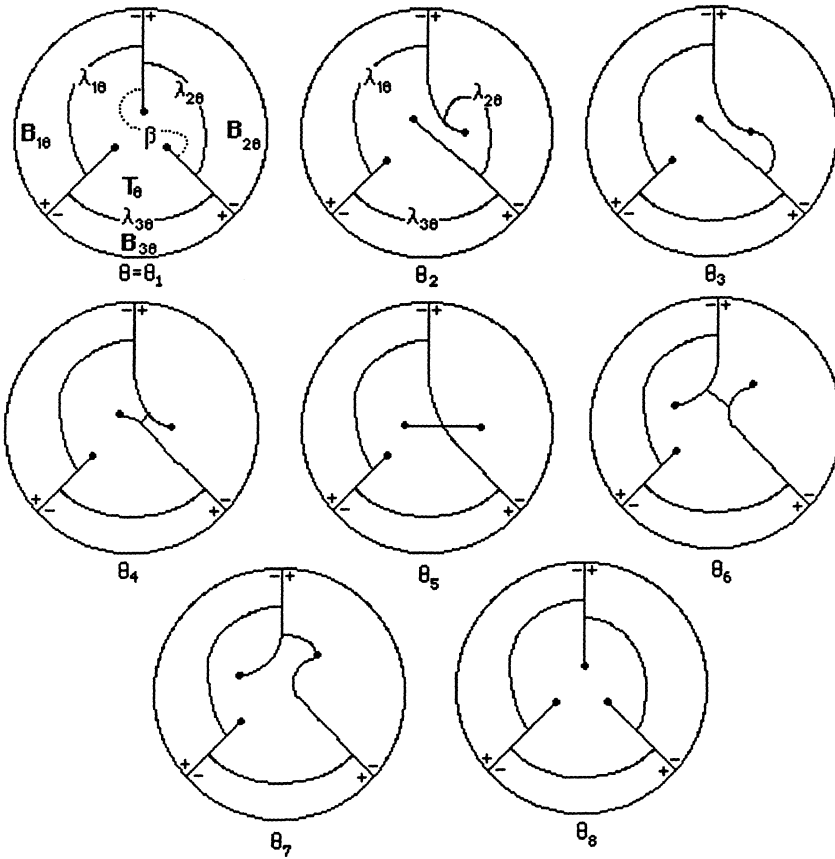


FIGURE 3.6

side of $F_{2\theta}$, as it has done in the passage from θ_2 to θ_4 . A crossing dot lies on the fiber at θ_3 . At the instant of crossover the digon will intersect L , and the intersection point (illustrated earlier in Figures 3.1–3.4) is one of the crossing dots. The F-singularity occurs at θ_5 . At θ_6 the arcs of $F \cap H_\theta$ draw apart again, but in a new way. At θ_7 the arc $\lambda_{2\theta}$ migrates across L a second time, now from the $-$ side to the $+$ side of $F_{3\theta}$, to restore the generic configuration at θ_8 . It is clear from these pictures that the discs $B_{1\theta}$, $B_{2\theta}$, $B_{3\theta}$ and T_θ are well-defined, even at θ_3 , θ_5 and θ_7 . The fact that the region swept out by T_θ is a solid torus results from the fact that $T_\theta \cap T_{\theta'}$ is empty if $\theta \neq \theta'$. On the other hand, if we denote the subarc of A which runs between p_i and p_{i+1} by A_i , $i = 1, 2, 3$, then $B_{i\theta} \cap B_{i\theta'} = A_i$ for all $\theta, \theta' \in [0, 2\pi]$. \square




FIGURE 3.7

PROPOSITION 3.2. *There is a natural diagram for L on the boundary ∂T of the said torus T , visible to an observer who is inside T . The overcrossing arcs are on ∂T . The undercrossing arcs are the arcs $\{L_i(k), k = 1, \dots, r_i, i = 1, 2, 3\}$. They are on $\partial B_1 \cup \partial B_2 \cup \partial B_3$, and thus are not visible to an observer who is inside T .*

Proof. The link L is in ∂T except for small θ -intervals about each F -singularity, e.g. $[\theta_3, \theta_7]$ in Figure 3.6, when L penetrates into B_2 . A few moments reflection should make it clear that the part of L which penetrates B_2 is precisely the subarc we called $L_2(k)$ in Figure 3.1(b). Figure 3.7 shows how these arcs look on a typical 2-sphere ∂B_i . They are the connecting threads which join up the digons in a circular ring. For later use, we have also illustrated in Figure 3.7 how the decomposition of F into subsets $1'$, $1''$, 2 and 3 will appear on ∂B_i .

It follows immediately that if we stand inside T and look out onto ∂T , we will see the link L minus the union of the r subarcs $L_i(k)$, $k = 1, 2, \dots, r_i$, one for each crossing point in the diagram on ∂T .

EXAMPLE. When we defined the planar projection of the elementary braids a_1, a_2, a_3 in Figure 1.5 we showed them as they would appear to a viewer who was standing outside T , between p_1 and p_3 with his feet on the plane of the paper and his head above the paper, and also as they would look to a viewer inside T .




FIGURE 3.8

In the pictures at the left in Figure 3.8 we have augmented the pictures in Figure 1.5 by adding the digons and shading the two sides of F . On the right in Figure 3.8 we show an example of a link diagram on ∂T . The reader is advised to return to Figure 3.8 later, as needed. In future pictures we will omit the shading which we use here to highlight F^+ and F^- . The rule is clear: parts of F which are to the right (left) of the oriented link L , with orientations from the top of the page toward the bottom, are in F^+ (F^-).

One more illustration is in order. Figure 3.9 (next page) shows the region which abuts on two digons, from two points of view: first, as it would appear on ∂T and second as it would appear on ∂B_i . The undercrossing arc is missing from the former and visible in the latter.

We need to use one more feature of the geometry in order to define the complexity function. The union of all of the digon-attaching curves (regarded now as curves on F and not on F^+ or F^-) divide F into a number of connected components. We call these *subsets* #1, 2 and 3 of F . They are defined uniquely by the following conditions:

- (a) If a component is intersected by A it is in subset #3.




FIGURE 3.9

- (b) If a component meets L , it is in subset #1. Also:
 If it meets the undercrossing arcs, it is in subset $1''$
 If not it is in subset $1'$
- (c) If a component has empty intersection with A and L , it is in subset #2.

Figure 3.3 illustrates the subdivision of F into subsets $1'$, $1''$, 2 and 3. Using it, we are almost ready to define our complexity function. Notice that we have already standardized the geometry in the following ways:

1. F is a Bennequin surface with respect to H and H' .
2. W and W' are admissible words.
3. A is disjoint from the digons, by construction.
4. The standard foliation of the interior of Ω is singularity-free.

We now need to standardize the intersections of fibers of H' with the digons:

LEMMA 3.3. *Let Ω be the union of all the digons. We may assume that*

(i) *The intersections of A' with $F \cup \Omega$ are finite in number and transverse.*

(ii) *The non-standard foliation of $F \cup \Omega$ is radial near each pierce point.*

(iii) *All but finitely many fibers H'_ϕ of H' meet $F \cup \Omega$ transversally. Those which do not (the Ω -singular fibers) are each tangent to $F \cup \Omega$ at exactly one point in the interior of Ω .*

(iv) *Each component of intersection of a non-singular fiber \mathbf{H}'_ϕ with \mathbf{F} (respectively with Ω) is an arc. There are no simple closed curves in the foliation of \mathbf{F} (respectively Ω).*

Proof. The arguments are identical with those used to prove Lemma 2 of [B-M, I]. All of the new modifications which are needed can be achieved by isotopy of Ω , so they won't upset earlier standardizations. \square

Finally, we have the long-promised definition of complexity. If S is a set, we write $|S|$ for the cardinality of S . If \mathbf{X} and \mathbf{Y} are families of surfaces in general position, we write $|\mathbf{X} \cdot \mathbf{Y}|$ to denote the cardinality of the set of tangencies between \mathbf{X} and \mathbf{Y} . Our complexity function is a 5-tuple $\mathbf{C} = (c_1, \dots, c_5)$, defined as follows:

- $c_1 =$ the number of syllables in the word $\mathbf{W} = W(a_1, a_2, a_3)$,
- $c_2 =$ the number of syllables in the word $\mathbf{W}' = W'(b_1, b_2, b_3)$,
- $c_3 = |\mathbf{A}' \cap (\mathbf{F} \cup \Omega)|$,
- $c_4 = |\mathbf{A}' \cap (\text{subsets \#1'' and \#2 and \#3 of } \mathbf{F})|$,
- $c_5 = |\mathbf{H}' \cdot (\mathbf{F} \cup (\text{int}(\Omega)))|$.

(Remark: The meaning of c_4 , and of our subsets of \mathbf{F} , will become clear later.)

We assume that we have chosen \mathbf{W} and \mathbf{W}' to minimize c_1 and c_2 . Our next proposition deals with the minimization of c_3 .

PROPOSITION 3.4. *We may assume that the non-standard axis \mathbf{A}' is disjoint from all of the digons.*

Proof. The proof will be based upon a study of the foliation of a single digon Δ by its arcs of intersection with the non-standard fibration \mathbf{H}' . Call this the *non-standard foliation* of Δ . Let \mathbf{H}'_ϕ be a non-singular fiber of \mathbf{H}' . If \mathbf{H}'_ϕ is not one of the exceptional fibers which is tangent to \mathbf{F} , the set $\mathbf{F} \cap \mathbf{H}'_\phi$ will be a union of three arcs, as in Figure 2.1(b). Figure 2.1(c) depicts the fiber \mathbf{H}'_ϕ split open along those arcs. We have labeled the sides of the split-open fiber with the numerals $0, 0', 0'', 1, 2, \dots, 6$ so that sides $0, 0', 0''$ are subarcs of \mathbf{A}' , sides $1, 3$, and 5 are subarcs of $\mathbf{F}^- \cap \mathbf{H}'_\phi$ and sides $2, 4$ and 6 are sub-arcs of $\mathbf{F}^+ \cap \mathbf{H}'_\phi$. There is a natural projection π from the split-open fiber to the closed fiber and a "lifting" π^{-1} . Thus we can describe each leaf in $\Delta \cap \mathbf{H}'_\phi$ by the unordered numbers of the sides to which its endpoints in the split-open fiber belong. We call these

numbers the *name* of the leaf. Define two leaves to be equivalent if they are equivalent under one of the six symmetries of the split-open fiber generated by the rotation which changes $0, 0', 0''$ to $0', 0'', 0$ and the flip which changes $0, 0', 0''$ to $0, 0'', 0'$. Let $[pq]$ denote the equivalence class of the leaf pq in \mathbf{H}'_ϕ .

We want to prove that A' does not intersect the digons. This is equivalent to the assertion that no arc in $\mathbf{H}'_\phi \cap \Omega$ has an endpoint on A' , i.e. leaves with name pq , where p or q is $0, 0'$ or $0''$ do not occur. Up to equivalence, it suffices to prove that leaves of name $0q$ do not occur, for $q = 0, 0', 0'', 1, 2, 3, 4, 5, 6$. As the first step in this direction we prove:

LEMMA 3.5. *We may assume that there are no leaves in the equivalence classes $[00]$ and $[01]$. Also, choosing any fixed fiber \mathbf{H}'_{ϕ_0} we may assume that there are no leaves of type $[11]$ in this fiber.*

Proof of Lemma 3.5. A leaf α (resp. α', β) in the class $[00]$ (resp. $[01], [11]$) cobounds with a subarc of A' (resp. subarcs of $A' \cup F$, subarc μ of $F \cap \mathbf{H}'_\phi$) a disc d (resp. d', d_H) on \mathbf{H}'_ϕ , as illustrated in Figure 3.10. These leaves are in $\Omega \cap \mathbf{H}'_\phi$, so they lie both in some digon Δ and on \mathbf{H}'_ϕ , as illustrated. Figure 3.10 also depicts these




FIGURE 3.10




FIGURE 3.11

arcs and associated discs in 3-space, for types [00] and [01]. The corresponding picture for type [11] is in Figure 3.11.

In the case of leaves of type [00], the only leaves which could intersect the disc d are other leaves of type [00], so we may assume that we have chosen α so that d is innermost. We may then push A' across d to remove α from $\Delta \cap H'_\phi$. This also removes leaves of type [00] from nearby fibers of H' . Since our isotopy of A' reduces c_3 , it reduces complexity. Thus we may assume there are no leaves of type [00].

In the case of leaves of type [01], there could be leaves of type [01] or [11] inside the disc d' , so we begin by removing the latter. Passing to Figure 3.11 we see that leaves of type [11] could intersect a component of $F \cap H'_\phi$ on both the positive and negative sides of F ,

and in fact there are three ways this could happen. If the endpoints of two such arcs β and β' of type [11] do not separate one another, as in the two left sketches, we simply choose β to bound an innermost disc d_H , so that no other leaf of type [11] intersects the arc μ in Figure 3.10. Assume this to be the case. Then, since β is an arc with its endpoints on $\mathbf{F}^\pm \cap \Delta$, it cobounds with a second arc μ' (also in \mathbf{F}) another disc $d_\Delta \subset \Delta$, as illustrated. The two discs d_H and d_Δ fit together to give a disc $d_H \cup d_\Delta$ with boundary $\mu \cup \mu'$ a simple closed curve on \mathbf{F} , as depicted in Figure 3.11. Now, \mathbf{F} is incompressible, so $\partial(d_H \cup d_\Delta)$ must also bound a second disc $d_F \subset \mathbf{F}$. The three discs $d_H \cup d_\Delta \cup d_F$ form a 2-sphere S^2 in $S^3 - L$. Since $S^3 - L$ is irreducible, this 2-sphere must bound a 3-ball \mathbf{B} in $S^3 - L$. We can then remove the intersections by changing the fibration, pushing \mathbf{H}'_ϕ across this 3-ball. After a finite number of repetitions of this argument we will obtain a half-disc on \mathbf{H}'_ϕ whose interior contains no leaves of type [11].

We now examine the same move, when the endpoints of our arc β of type [11] are separated by the endpoints of another arc β'' of type [11]. Notice that the only way this can happen is if β and β'' are on opposite sides of \mathbf{F} , because digons are embedded, and by our construction distinct digons have disjoint interiors. So the situation in 3-space must be as illustrated in the bottom sketch in Figure 3.11. It is clear that our change in fibration will not increase the intersections of fibers of \mathbf{H}' with Δ' , so in every case the complexity is reduced.

Returning to Figure 3.10, we may now proceed to remove leaves of type [01]. We begin with an "innermost" leaf of type [01]; that is, as we proceed along a component of $\mathbf{H}'_\phi \cap \mathbf{F}$, starting at the A' endpoint, we choose the arc α' of type [01] whose \mathbf{F} -endpoint is encountered first. The disc d' in Figure 3.10 then meets no other leaf of type [01]. Pushing A' across d' , we remove the leaf of type [01] from this fiber and also from nearby fibers. Since complexity is reduced, we conclude that leaves of type [01] do not occur. \square

Proof of Proposition 3.4 (continued). In view of Lemma 3.5, we are reduced to showing that $q \neq 0', 0'', 2, 3, 4, 5$. Now, since \mathbf{W}' is admissible, we know that \mathbf{W}' contains the letter $(b_3)^{\pm 1}$. So there is an \mathbf{F} -singularity of type $(b_3)^{\pm 1}$ at, say, ϕ_0 . The joining arcs for a singularity of type $(b_3)^{-1}$ run between the arcs we have labeled 2 and 6, whereas those for b_3 run between sides 1 and 5. Thus $q(\phi_0)$ cannot be $0', 0'', 3$ or 4 in either case, because such a leaf is an obstruction

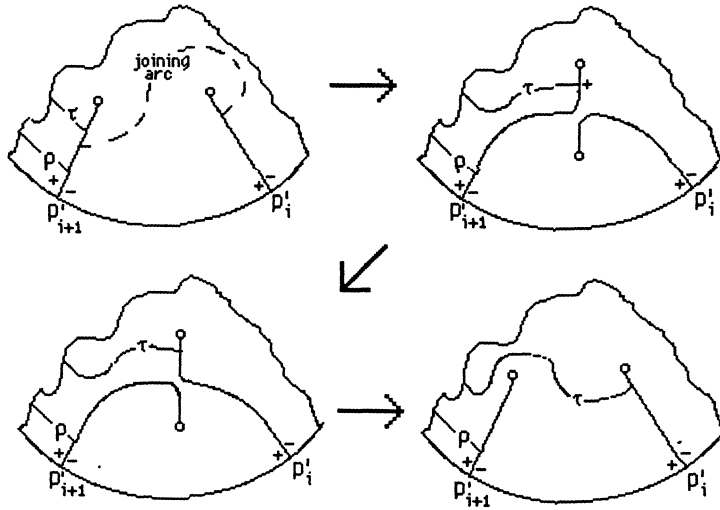
Effect of a type b_i singularity on digon-crossing arcs

FIGURE 3.12

to the joining arc, also if $(b_3)^{-1}$ occurs it cannot be 5, whereas if b_3 occurs it cannot be 2. Assume the former. Then $q(\phi_0)$ must be 2. But now Figure 3.12 shows how a singularity of type $(b_i)^{-1}$ can result in a change in name of other arcs. If $q(\phi) = 2$ before an F-singularity of type b_3^{-1} , then after the singularity there will be a leaf with name 06. But this is impossible by Lemma 3.5 because $[06] = [01]$. Similarly, if an F-singularity of type b_3 occurs at ϕ_0 then the only possibility is that $q(\phi_0) = 5$, but in that case there will be a leaf with name 01 after the singularity, and again we have a contradiction to Lemma 3.5. Thus we conclude that for ϕ -values which are close to ϕ_0 there are no leaves of type $0q$, because all values of q have been ruled out.

Now notice that the points where A' pierces Ω are independent of ϕ . Thus if there is a leaf with name $0q$ in some fiber H'_ϕ , there must be a leaf with name $0\dots$ in every fiber of H' . However we have just shown that no such leaf can occur when ϕ is close to ϕ_0 . Thus it cannot occur at any other ϕ -value either.

But then, by symmetry, there are no leaves of type $0'q$ and $0''q$ either. Thus A' is disjoint from all of the digons. \square

4. The non-standard foliation of the digons. We have shown that A' is disjoint from all of the digons. Thus the entries c_1, c_2, c_3 in our complexity 5-tuple have now been minimized. In this section we will

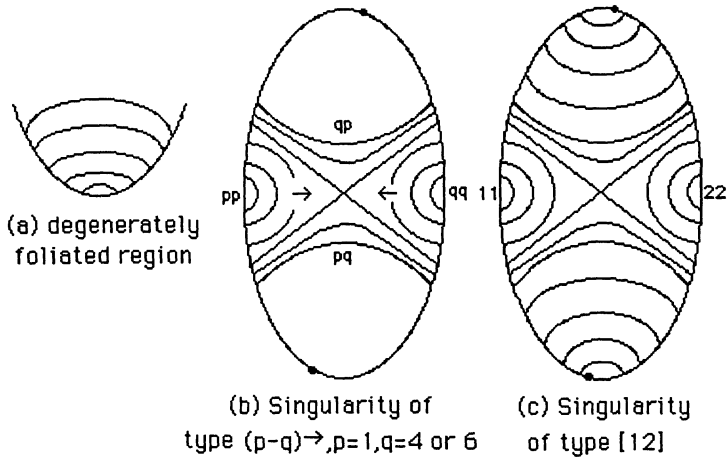


FIGURE 4.1

deal with consequences of the minimization of c_5 . We have not yet dealt with the consequences of the minimization of c_4 . This will not matter. The actual value of c_4 depends upon the location on F of the three points in $A' \cap F$. In this section we will regard A' as rigidly fixed. The minimization of c_5 can thus be carried out, independently, for each of the possible values of c_4 . Later, when we need to move A' in order to minimize c_4 we will be sure that no matter where it is located, subject to earlier restrictions, we will be able to achieve the standardizations of this section.

The central result of this section will be to show that the assumption that c_5 is minimal implies that the possible singularities in the non-standard foliation of the digons are quite limited. See Figure 4.1. The precise restrictions, which can only be stated in terms of the notation set up below, are given in Proposition 4.8. We will end this section with a result (Corollary 4.9) which shows that these singularities are a key to the geometry, because we can use them to deduce an important property of W .

As in §3, a leaf in the non-standard foliation of a digon has name pq , where (by Proposition 3.4) p and q belong to the set $\{1, 2, 3, 4, 5, 6\}$. Its equivalence class under symmetries of H'_ϕ is denoted $[pq]$. There are 5 equivalence classes of leaves, listed in Table 4.1. Sometimes it will be useful to refine this classification by distinguishing equivalence classes with different *signs*, where the signs $++$, $--$ and $+-$ distinguish the cases when the endpoints of a leaf are both in

TABLE 4.1

type	Sign	Members of class
[22]	++	22, 44, 66
	--	11, 33, 55
[24]	++	24, 46, 62
	--	13, 35, 51
[14]	+−	14, 25, 36
[12]	+−	12, 34, 56
[23]	+−	23, 45, 61

$\partial\Delta^+$, both in $\partial\Delta^-$ or one in each. The 5 equivalence classes of leaves divide into 7 subclasses when we consider signs, and these are also listed in Table 4.1.

As the fibers of \mathbf{H}' pass through Ω -singular or \mathbf{F} -singular fibers the names of the leaves will change. If two leaves have names mn and pq before the singularity, and np and qm after it, we can describe the singularity by the symbol “ $mn + pq \rightarrow np + qm$ ”. We choose our notation so that the labels m, n, p, q have that order on $\partial\mathbf{H}'_\phi$. The surgery is then uniquely specified. The equivalence class of a surgery $[mn + pq \rightarrow np + qm]$ means equivalence up to symmetries of \mathbf{H}'_ϕ and (if the symbol “ \rightarrow ” is replaced by “ \leftarrow ”) reversal of the sense of the fibration.

LEMMA 4.1. *Singularities in the class $[mq + qq \leftrightarrow qq + qm]$ do not occur.*

Proof. See Figures 4.2 and 4.3 (next page). The top picture in Figure 4.2 shows the leaves which would be involved in the surgery, as they occur on three fibers in the nonstandard fibration, at $\phi = \phi_0, \phi_1$ and ϕ_2 . The middle picture is the singular fiber and the other two are fibers just before and just after the singularity at s . The point p' is on \mathbf{A}' , and the point $\mathbf{L}(\phi)$ is on the link \mathbf{L} . The heavy arc joining p' to $\mathbf{L}(\phi)$ is on $\mathbf{F} \cap \mathbf{H}'_\phi$. The lighter arcs are on $\Delta \cap \mathbf{H}'_\phi$. Two discs d and δ have been labeled in the middle picture. The next picture in Figure 4.2 shows the foliation of the digon Δ in the region under investigation. Its boundary $\partial\Delta$ is on \mathbf{F} , and we notice that the entire region of interest is attached to either \mathbf{F}^+ or \mathbf{F}^- , i.e. $\partial\Delta$ does not cross \mathbf{L} . We have labeled two regions of Δ as d' and δ' , noting that $\partial d = \partial d'$ and $\partial \delta = \partial \delta'$. The final picture in Figure 4.2 shows the region of interest on Δ , embedded in 3-space.





FIGURE 4.2




FIGURE 4.3

Let $\mathbf{D} = d \cup d'$ or $\delta \cup \delta'$. Then \mathbf{D} is a disc, and its boundary $\partial\mathbf{D}$ is a curve in F . Since F is incompressible, the curve $\partial\mathbf{D}$ also bounds a second disc, say \mathbf{D}' , in F . Now $\mathbf{D} \cup \mathbf{D}'$ is a 2-sphere in $\mathbf{M} = S^3 - L$, and since \mathbf{M} is irreducible, this 2-sphere must bound a

3-ball \mathbf{B} in \mathbf{M} . Then we can use \mathbf{B} to isotope d' to d . Similarly, we can isotope δ' to δ . After the change we tip the discs d and δ a little bit, to make the new digon transverse to the fibers of \mathbf{H}' . The effect of the isotopy is to eliminate the given singularity of the foliation. See Figure 4.3. This reduces c_5 , keeping c_1, \dots, c_4 fixed. Since \mathbf{C} is by hypothesis minimal, this cannot happen. Notice that these changes can't increase $|\mathbf{H}' \cdot (\mathbf{F} \cup \text{Int} \Delta')|$, even if there is another digon Δ' in the picture, because distinct digons which intersect always do so in a very controlled fashion (cf. the bottom picture in Figure 3.11), because near a point which is in $\Delta \cap \Delta'$ the digons Δ and Δ' will always be attached to opposite sides of \mathbf{F} . \square

A region T on a digon Δ is said to be *degenerately foliated* if it is foliated by parallel arcs of decreasing lengths, with a single degenerate leaf as in Figure 4.1a. Our first observation is that degenerately foliated regions sometimes tell us that there is a fiber \mathbf{H}'_ϕ which is tangent to \mathbf{F} at a ϕ -value near the value at the instant of degeneracy:

LEMMA 4.2. (1) *Let T be a degenerately foliated region on Δ . Suppose that each nondegenerate leaf in T has name 24 (respectively 46, 62, 13, 35, 51). Then within a small ϕ -neighborhood of the instant of degeneracy a braiding of type b_1^{-1} (resp. b_2^{-1} , b_3^{-1} , b_1 , b_2 , b_3) occurs in \mathbf{F} .*

(2) *If a leaf α in the foliation of a digon Δ has type [24] then the side of $\Delta - \alpha$ which does not contain the crossing dots is degenerately foliated.*

Proof. (1) See Figure 4.4 (next page), where we have depicted T as a “tongue” which is orthogonal to some non-singular \mathbf{H}'_ϕ , its boundary being glued to \mathbf{F} . Typical leaves α , α' , α'' are shown. Let $\{\alpha(\phi); \phi \in [\phi_0, \phi_1]\}$ be the leaves in the foliation of T , with $\alpha(\phi_0)$ the given leaf of name 24 and $\alpha(\phi_1)$ the degenerate leaf. (See also Figures 3.6, which can now be reinterpreted as a series of pictures on fibers of \mathbf{H}' , with the dotted arc which is labeled β there now reinterpreted as our $\alpha(\phi_0)$.) Each $\alpha(\phi)$ has name 24 and each begins and ends in \mathbf{F}^+ . The length of $\alpha(\phi)$ decreases monotonically to 0 as $\phi \rightarrow \phi_1$. This means that near ϕ_1 there must be an \mathbf{F} -singularity between arcs of $\mathbf{F}^+ \cap \mathbf{H}'_\phi$. By our conventions it is a singularity of type b_1^{-1} .

(2) Let $\alpha = \alpha(\phi)$ be the given leaf of name 24. Then α divides Δ , and since both endpoints of α are in the same component of




FIGURE 4.4

$\partial\Delta - \partial\Delta \cap L$, the two crossing dots are on the same side of Δ split along α . Let T be the closure of the region of $\Delta \setminus \alpha$ which does not contain the crossing dots. Then T must meet H_ϕ transversally, as in Figure 4.4, also the 2 (respectively 4) endpoint of α lies on the arc $F_{2\phi}$ (resp. $F_{4\phi}$) of $F^+ \cap H'_\phi$ which runs out to p'_1 (resp. p'_2). Now, $\partial T \subset \partial\Delta$ is in F , so F necessarily forms a bridge over T , joining $F_{2\phi}$ to $F_{4\phi}$. But then, possibly after a small change in H' , the region T will be degenerately foliated. \square

LEMMA 4.3. *If α has type [22] then the side of $\Delta - \alpha$ which does not contain the crossing dots is degenerately foliated.*

Proof. The assertion follows easily with the help of ideas used in the proof of Lemma 3.5. Refer to Figure 3.10, and consider the leaf β which is shown there. Suppose that both sides of Δ split along β contain singularities of the foliation. Without loss of generality we may assume that β is innermost on H'_ϕ , i.e. the interior of d_H has empty intersection with Δ . Since the portion of ∂d_Δ which is in F lies entirely in F^+ , we conclude that d_Δ cannot contain the crossing dots. As in the proof of Lemma 3.5 we may construct a new digon by discarding d_Δ and replacing it by d_H . If d_Δ contained a singularity of the foliation then c_5 would not be minimal. Thus d_Δ must be degenerately foliated. \square

LEMMA 4.4. *Surgeries between two leaves which have sign $++$ (or two leaves which have sign $--$) do not occur.*



FIGURE 4.5

Proof. Let α_1 and α_2 be the leaves to be surgered. From Table 4.1 we see that each α_j has type [22] or [24]. Let T_j , $j = 1, 2$, be the region of $\Delta - \alpha_j$ which does not contain the crossing dots. Then T_j is degenerately foliated, by Lemmas 4.2 and 4.3. Now $\partial T_j = \alpha_j \cup \beta_j$ where β_j is in $\partial\Delta^+$ for $j = 1$ and 2 . There are two possibilities (depicted in Figure 4.5):

- (i) β_1 contains β_2 (or β_2 contains β_1)
- (ii) $\beta_1 \cap \beta_2$ is empty.

In case (i) the singularity is in T_1 , contradicting the fact that T_1 is degenerately foliated, so this can't happen. In case (ii), after the surgery we will have two new leaves of type [22] or [24], and one of them will bound a region T_3 which contains both T_1 and T_2 and does not contain the crossing dots, and is not degenerately foliated, contradicting Lemma 4.3. \square

LEMMA 4.5. *Δ -surgeries between leaves of sign ++ (or --) and leaves of sign +- do not occur.*

Proof. Suppose that the surgery in question is $pq + rs \rightarrow qr + sp$, where pq has sign ++ and rs has sign +-. Our first assertion is that the surgery either begins or ends with a leaf in the equivalence class [24]. For, suppose it does not. Then $[pq] \neq [24]$, but $[pq]$ has sign ++ or --, so from Table 4.1 we must have $[pq] = [22]$. Assume that $pq = 22$. Since rs has sign type +-, and since the names 25, 12, 23 are ruled out by Lemma 4.1, we see from Table 4.1 that $rs = 14, 36, 34, 56, 45$ or 61 . The six surgeries may therefore be described as $22 + r4 \rightarrow 2r + 42$ and $22 + r6 \rightarrow 2r + 26$, where $r = 1, 3, 5$. Up to equivalence the latter is $44 + r2 \rightarrow 4r + 24$, where r is still 1, 3 or 5. Thus, up to equivalence we may assume that our surgery produces a leaf of name 24.

By Lemma 4.2 the leaf with name 24 which is produced by the surgery bounds a degenerately foliated region on Δ , and near the

instant of degeneracy there must be a braiding of type b_1^{-1} . Now, there is no obstruction to changing the fibration a little bit so that the braiding occurs before the surgery. The effect of the braiding of type b_1^{-1} is to change the names of the leaves which are involved in the surgeries to $44 + r4 \rightarrow 4r + 44$ or $22 + r2 \rightarrow 2r + 22$. However, that violates Lemma 4.1.

It remains to consider the case when the surgery both begins and ends with a leaf in the class [24]. Up to equivalence, we may assume that it begins with a leaf of name 24. Consulting Table 4.1, we see that the surgery must be between our leaf and one of type [14], [12] or [23], and on running through all the possibilities we see there are exactly 2 which produce a leaf of type [24], namely $24 + 61 \rightarrow 46 + 21$ and $24 + 56 \rightarrow 45 + 26$, and these two are equivalent up to a change in the sense of the fibration, so we may assume our surgery is $24 + 56 \rightarrow 45 + 26$. The argument now proceeds as before. By Lemma 4.2 the leaf with name 46 which is produced by the surgery bounds a degenerately foliated region, so the surgery must be followed by a braiding of type b_2^{-1} . By a small change in fibration, we may assume that the braiding occurs before the surgery, so the sequence may be assumed to be b_2^{-1} , $26 + 16 \rightarrow 66 + 12$. Reversing the order of fibration, this becomes $66 + 12 \rightarrow 16 + 26$, b_2 . Now, the leaf 26 which is produced by the surgery has type [24], so it bounds a degenerately foliated region and so must be followed by a braiding, of type b_3^{-1} , that is the sequence must be $66 + 12 \rightarrow 16 + 26$, b_3^{-1} , b_2 . Changing the fibration, we may reverse the order of the surgery and the first braiding, obtaining the sequence b_3^{-1} , $66 + 16 \rightarrow 61 + 66$, b_2^{-1} . However, this violates Lemma 4.1. \square

LEMMA 4.6. *Let Δ_1 and Δ_2 be the components of Δ split along a leaf α of type [12]. Then (possibly after a small modification in \mathbf{H}') we may assume that either Δ_1 or Δ_2 is degenerately foliated.*

Proof. Figure 4.6a shows a leaf α of type [12] as it appears on the fiber \mathbf{H}'_ϕ . Let x, y be the endpoints of α and let β be the arc of $\mathbf{F} \cap \mathbf{H}'_\phi$ which runs between x and y . Then $\alpha \cup \beta$ is the boundary of a disc D on \mathbf{H}'_ϕ which is pierced once by the link \mathbf{L} at the point z .

Figure 4.6b shows the same leaf α on the digon Δ . It divides Δ into subdiscs Δ_1 and Δ_2 , with $\partial\Delta_j = \alpha \cup \delta_j$, $j = 1, 2$, where δ_j lies in \mathbf{F} . The link \mathbf{L} meets Δ twice, at the crossing dots $w_j \subset \delta_j$,




FIGURE 4.6

$j = 1, 2$. The sub-arcs δ_{1+} and δ_{2+} are in $\partial\Delta^+$, while δ_{1-} and δ_{2-} are in $\partial\Delta^-$.

Figure 4.6c puts the pictures in 4.6a and b together to give view in 3-space. The disc Δ_1 , together with a triangular region in F , is a cone C_1 on $\alpha \cup \beta$ with vertex w_1 . The union of C_1 and D is then a 2-sphere S_1 which is pierced twice by L , at w_1 and at z . There is also another cone C_2 with vertex w_2 on the other side of H'_ϕ , with $C_1 \cap C_2 = \alpha \cup \beta$, determining a second 2-sphere S_2 which is pierced twice by L at z and w_2 . Let $j = 1$ or 2 . Then S_j divides 3-space into two 3-balls, B_j and E_j and one of them meets L in a trivial arc, because by hypothesis L is prime. Suppose the notation to have been chosen so that $B_1 \cup B_2$ is the 3-ball which contains D . We claim that $L \cap B_j$ must be unknotted for either $j = 1$ or $j = 2$. For, if $L \cap B_j$ were knotted for both $j = 1$ and 2 , then $L \cap E_j$ would be unknotted

for $j = 1$ and 2 , and since $\mathbf{L} = (\mathbf{L} \cap E_1) \cup (\mathbf{L} \cap E_2)$ that's impossible.

We assume the notation to have been chosen so that $\mathbf{L}_1 = \mathbf{L} \cap B_1$ is unknotted. This means that the triangular region of \mathbf{F} with vertices z, y, w_1 is a disc. Let \mathbf{N} be a neighborhood of \mathbf{L}_1 in 3-space which includes Δ_1 . Then we can modify \mathbf{H}' , if necessary, by pushing the fibers up a little bit through \mathbf{N} , to insure that Δ_1 is degenerately foliated. \square

LEMMA 4.7 (*cf. Figure 4.1b*). *Every Δ -surgery between two leaves in the foliation of a digon Δ may be assumed to be type $pp + qq \leftrightarrow pq + qp$, where one of p, q is even and the other is odd.*

Proof. Since every leaf in the foliation has sign $++$, $--$, or $-+$ Lemmas 4.4 and 4.5 imply that the only surgeries which are possible are ones which exchange a leaf of sign $++$ and a leaf of sign $--$ for two leaves of sign $-+$. Up to equivalence we may assume the leaf of sign $++$ has name 22 or 24. Since there are six leaves of sign $--$ in Table 4.1, there are at most twelve surgeries which exchange a leaf of sign $++$ and a leaf of sign $--$ for two of sign $-+$. Our task is to show that the nine surgeries 22 with 13, 35, 51 and 24 with 11, 33, 55, 13, 35, 51 do not occur.

The first observation is that the first three surgeries in our group of nine are identical with the second three up to equivalence, so we forget the first three. Also, we can eliminate the surgeries 24 with 13 and 35 because the leaves in question are not disjoint and therefore cannot both appear on a single same fiber. Thus we are left with four possibilities:

$$\begin{aligned} [24 + 11 \rightarrow 41 + 12], & \quad [24 + 33 \rightarrow 43 + 32], \\ [24 + 55 \rightarrow 45 + 52], & \quad [24 + 51 \rightarrow 45 + 12]. \end{aligned}$$

We consider the first, second and last, choosing a new representative of the equivalence class so that the surgeries in question begin with $12 + 41$, $12 + 61$, $12 + 45$ respectively. (In the case of the second surgery this requires us to reverse the sense of the fibration.) By Lemma 4.6 the leaf 12 bounds a degenerately foliated region T on Δ which contains one of the crossing dots. Since the degenerate leaf in that region does not coincide with the dot (by our choices, in §4, of the attaching curves for the digons) the region T necessarily also contains leaves of type 11 or 22. By a small change in the fibration the surgery can be changed to one between 11 or 22 and 41, 61 and

45. However, since 11 has sign $--$ and 22 has sign $++$, while 41, 61 and 45 all have sign $+ -$ this contradicts Lemma 4.5.

The surgery $24 + 55 \rightarrow 45 + 52$ remains. The leaf 24 divides Δ into two regions, one of which contains the singularity of the given foliation. The other one must be degenerately foliated by Lemma 4.2. Moreover, there is a braiding of type b_1^{-1} at the instant of degeneracy. By a small change in the fibration, we may assume that the braiding occurs *before* the surgery. This braiding changes the leaf of type 24 to a leaf of type 44 or 22 and leaves the leaf of type 55 unaltered. The new surgery is $44 + 55 \rightarrow 45 + 54$ or $22 + 55 \rightarrow 25 + 52$. Each is equivalent to one of those given in the statement of Lemma 4.7. \square

We may now simplify our notation. The symbol $[pq + rs \leftrightarrow qr + sp]$ which we originally needed to describe the digon singularities has now been replaced by the simpler symbol $[pp + qq \leftrightarrow pq + qp]$. We now simplify it still further to $[p - q]$, with the symbols $[p - q]^\rightarrow$ and $[p - q]^\leftarrow$ used to distinguish the cases $[pp + qq \rightarrow pq + qp]$ and $[pp + qq \leftarrow pq + qp]$. The symbol $[p - q]$ means either $[p - 1]^\rightarrow$ or $[p - q]^\leftarrow$. Similarly, we have representatives $(p - q)$, $(p - 1)^\rightarrow$ and $(p - q)^\leftarrow$ of $[p - q]$, $[p - q]^\rightarrow$ and $[p - q]^\leftarrow$. Using this notation, the results of this section are summarized below.

PROPOSITION 4.8. (1) *A leaf of type [12], [22] or [24] necessarily splits off a degenerately foliated region on a digon, as in Figure 4.1a. In the cases of types [22] and [24] the region is the one which does not contain the two crossing dots.*

(2) *If a digon is not standardly foliated (Figure 3.4) then its singularities are one of the following three types:*

type [1 - 2], with representatives (1 - 2), (3 - 4) and (5 - 6);

type [1 - 4], with representatives (1 - 4), (3 - 6) and (5 - 2).

type [1 - 6], with representatives (1 - 6), (3 - 2) and (5 - 4).

(3) *If a singularity of type [1 - 4] or [1 - 6] occurs in a digon Δ , then a neighborhood of the singular leaves in Δ is foliated as in Figure 4.1b.*

(4) *If a singularity of type [1 - 2] occurs in a digon Δ , then it is the only singularity in Δ , and all of Δ is foliated as in Figure 4.1c.*

Proof. The only part which is not simply a restatement of earlier results is assertion 4, and even that assertion is essentially obvious. For, assume that the given singularity is $11 + 22 \rightarrow 12 + 21$. The

leaves of type 11 and 22 each divide Δ into two pieces, one of which contains the two crossing dots. By Lemma 4.3, the side which does not contain the dots is degenerately foliated. Two leaves of name 12 are produced by the surgery. Each divides Δ into two regions, one of which contains the singularity. By Lemma 4.6 the other must be degenerately foliated. But then, all of Δ is foliated, with exactly one singularity. \square

We are now able to use information about the foliation of the digons to deduce a strong property of the braid word $W(\mathbf{L})$.

COROLLARY 4.9. *If the non-standard foliation of a digon $\Delta_i(k)$ contains an Ω -singularity in the equivalence class $[1 - 2]^\rightarrow$ then there is exactly one digon joining disc \mathbf{D}_i to \mathbf{D}_{i+1} , and one of the braid generators $(a_i)^{\pm 1}$ occurs exactly once in $W(\mathbf{L})$. Moreover, the non-standard foliation of the entire 2-sphere $\partial\mathbf{B}_i$ contains exactly one digon-singularity, namely the given singularity of type $[1 - 2]^\rightarrow$.*

Proof. By hypothesis, there is a digon $\Delta_i = \Delta_i(1)$ whose foliation contains a singularity of type $11 + 22 \rightarrow 12 + 21$. By Proposition 4.8 we conclude that Δ_i is foliated with exactly one singularity, as in Figure 4.1c.

We examine a split-open fiber of \mathbf{H}' just before, during and after the singularity at $\phi = \phi_0$. The first thing that will happen is the appearance of a pair of leaves, of names 11 and 22. The surgery produces a pair of parallel leaves of name 12. After the surgery these leaves will “migrate” across \mathbf{L} , changing their names to 11 and 22. The latter will then shrink and disappear. See Figure 4.7.

Let α and α' denote the parallel leaves of name 12, as in Figure 4.8a. Notice that α and α' and two subarcs β and β' of $\mathbf{H}'_{\phi_0} \cap \mathbf{F}$ form a simple closed curve $x = \alpha \cup \beta \cup \alpha' \cup \beta'$ type. We claim that this simple closed curve is in the boundary of one of the 3-balls \mathbf{B}_i . For, if we split \mathbf{H}'_{ϕ_0} open along x we obtain three components, and since the one which is closest to the axis \mathbf{A}' must be in \mathbf{T} the middle component is in some \mathbf{B}_i . Thus x is in $\partial\mathbf{B}_i$.

Figure 4.8b shows x as it would appear to an observer inside \mathbf{B}_i who is looking out at the boundary. Notice that x bounds a disc d on $\partial\mathbf{B}_i$ which contains at least one undercrossing arc $\mathbf{L}_i = \mathbf{L}_i(1)$, and in fact it contains exactly one because the entire disc d is foliated during the ϕ -interval studied in the sequence of pictures in Figure 4.7. Since there is a digon singularity between α and α' , the two half-digons

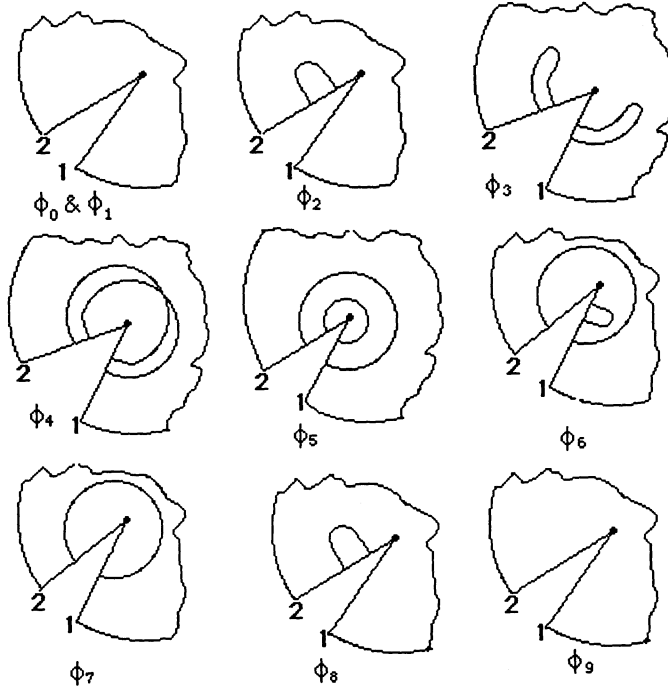


FIGURE 4.7




FIGURE 4.8

which are illustrated in Figure 4.8b must be part of the same digon. But then, this digon must be joined to itself by the undercrossing arc L_i , so our digon is the only digon which joins the discs D_i and D_{i+1} . Since the number of digons between D_i and D_{i+1} is equal to the number of bands between D_i and D_{i+1} this shows that a_i or its inverse occurs exactly once in $W(L)$. \square

5. Moving A' into T . We have dealt with the minimization of every entry in C except c_4 . In this section we study the minimization of c_4 and consequences. The principal result is:

PROPOSITION 5.1. *The non-standard axis A' is inside the T component in the decomposition of S^3 split open along $F \cup \Omega$.*

The proof will occupy most of the section. After it is completed we will establish a significant corollary.

Proof of Proposition 5.1. We begin by describing a set of partial “coordinates” for the non-standard and standard axes. By Corollary 2.3 we know that A and A' each pierce F three times, and so each is divided into 3 segments by its intersections with F . The segments on A (respectively A') have a natural cyclic order on A (resp. A'), determined by the orientations on A and A' . By construction we know that A is disjoint from the union Ω of the digons. By Proposition 3.4 we know that A' is too. Thus each segment of A and each segment of A' is in one of the four components B_1, B_2, B_3, T of the torus-3-ball decomposition of S^3 . We assign labels to the segments of A and A' to obtain a *cycle of labels*, a cyclic word of length three in the symbols B_1, B_2, B_3, T . The cycle of labels for A is, of course, $B_1B_2B_3$. Clearly A' is isotopic to A if its cycle of labels is also $B_1B_2B_3$. Our goal is to show that with our definition of complexity the cycle of labels for A' is TTT .

We can make the information in the cycle of labels more precise by placing restrictions on the points at which A' pierces F . Recall that we gave a decomposition of F into subsets #1, #2, #3 in §3. See Figure 3.3. There are three components of set #3, one in each disc. The manner in which set #2 meets a band can be seen from Figure 3.1(b). There are two components, one of which connects to D_i and the other to D_{i+1} . Recall that the digons which join the discs D_1 and D_{i+1} (respectively D_{i-1}) are attached to D_i along arcs parallel to ∂D_i . We now make the convention that when there are digons attached to both D_{i-1} and D_{i+1} the attaching arcs for the digons going to D_{i+1} are closer to ∂D_i than those going to D_{i-1} . Thus set #2 separates sets #1 and #3 on the disc parts of F . Set #1 will be discussed in more detail shortly.

Clearly F is the disjoint union of #1, #2, and #3. Therefore we have a second cyclic word, now in the symbols 1, 2, 3, which describes where A' pierces F . The possible *cycles of pierce points* are: 111, 222, 333, 122, 133, 211, 233, 311, 322, 123, 132.

The cycle of pierce points and the cycle of labels are, of course, related. If A' pierces F in set #1 its label will not change; if it pierces F in set #2 its label will change from T to B_i or B_i to T ; if it

pierces \mathbf{F} in set #3 its label will change from \mathbf{B}_i to \mathbf{B}_{i+1} . This shows immediately that the cycles 222, 133, 211, 233, 311, 123, 321 do not occur because they do not correspond to cycles of labels. Also, the cycle of labels is $\mathbf{B}_1\mathbf{B}_2\mathbf{B}_3$ if and only if the cycle of pierce-points is 333. This reduces things to the pierce-point sequences 111, 122 and 322.

The subset #1 of \mathbf{F} , which contains \mathbf{L} , merits special attention. As noted earlier, it splits into two subsets. The first contains the r undercrossing arcs $\mathbf{L}_i(k)$ and the adjacent triangular regions $T_i(k)$, $i = 1, 2, 3$, $k = 1, \dots, r_i$. It is subset #1''. The second is everything else. It also has r components, each of which is made up from a triangular region opposite some $\mathbf{L}_i(k)$, together with long narrow strips on either side of the triangular region, which run parallel to \mathbf{L} , out to the disc parts of \mathbf{F} and along the discs from one band to the next. It is subset #1'.

LEMMA 5.2. *Distinct points of $\mathbf{A}' \cap \mathbf{F}$ are in distinct components of subset #1' or 1''.*

Proof. Suppose that p'_i and p'_{i+1} are in the same component \mathbf{X} of subset 1' or 1''. Each component \mathbf{X} is a topological disc. This is clear if \mathbf{X} is set 1''. As for set 1', the only way this could fail to be the case is if there is a pair of bands $b_i(k)$ and $b_{i+1}(k)$ which are not separated by any band $b_{i\pm 1}(q)$, and the bands have opposite senses; however in that case \mathbf{W} would not be freely reduced. Thus \mathbf{X} is a disc. Its boundary is a union of two arcs. One of them (call it \mathbf{L}_0) is a subarc of \mathbf{L} and the other is in the boundary of one of the digons. See Figure 5.1 (next page). The fact that \mathbf{F} is a Bennequin surface relative to \mathbf{H}' tells us there is a neighborhood of p'_i and of p'_{i+1} which is radially foliated, also the foliation of \mathbf{X} is everywhere transverse to \mathbf{L} . This implies, immediately, that there is a singular leaf in the foliation which joins p'_i to p'_{i+1} . Even more, there must be a transverse singular leaf, say β , with both of its endpoints on \mathbf{L}_0 . The endpoints of β divide \mathbf{L}_0 into three components. Let \mathbf{L}_1 be the middle one. The leaf β separates \mathbf{X} into two subdiscs. Let \mathbf{X}_1 be the one with $\partial\mathbf{X}_1 = \beta \cup \mathbf{L}_1$. Then \mathbf{X}_1 must be radially foliated, as in Figure 5.1, and if we split \mathbf{F} along β we will obtain a new surface \mathbf{F}' , with $\partial\mathbf{F}' = \mathbf{L} - \mathbf{L}_1 + \beta$ a new braid representative of the same link type 1. The new representative is a 2-braid. That is impossible. \square

With our refinement of set #1, the pierce-point cycles which can




FIGURE 5.1




FIGURE 5.2

occur are $1'1'1'$, $1''1''1''$, $1'22$, $1''22$ and 322 . Notice that the axis A' necessarily passes from T to T if it punctures F in set $\#1'$, from B_i to B_i if it punctures F in $1''$. Thus our five pierce-point cycles correspond to the cycles of labels and pierce-points which are listed in Figure 5.2. Note that we have fixed notation, so that (for example) in case (ii) the pierce-points p'_1 and p'_2 are in subset 2 of F and p'_3 is in subset 3, so that α_1 is in B_i and α_2 and α_3 are inside T.

Case (i) of Figure 5.2, with pierce-point cycle $1'1'1'$, is the only one in which the cycle of labels is TTT. Thus the truth of Proposition 5.1 is equivalent to the assertion that cases (ii), (iii), (iv) and (v) do not occur. This, in turn, is equivalent to the assertion that the entry c_4 in the complexity function C is zero, because c_4 is the number of points in $A' \cap F$ which are *not* in set $1'$. We will prove that cases (ii), (iii), (iv), (v) cannot occur.

In what follows, the reader may find it helpful to refer back to Figures 3.3 and 3.7 as needed. The former illustrates the decomposition




FIGURE 5.3

of F into subsets $1'$, $1''$, 2 and 3 and the latter illustrates the induced decomposition of ∂B_i .

Case (ii). Assume that case (ii) occurs. We begin by examining the geometry near the pierce-point p'_3 , which (unlike p'_1 and p'_2) is in subset #1' of F . See Figure 5.3 (and set $i = 3$). The facts that:

- (1) L is transverse to every fiber of H' ,
- (2) the foliation of F is radial near each pierce-point,
- (3) points in subset $1'$ are arbitrarily close to L , and
- (4) the point p'_3 is the only point of $A' \cap F$ in subset $1'$ of F

imply that (possibly after a small modification in H') we may choose a fiber H'_{ϕ_0} which meets F in an arc β which joins p'_3 directly to L in subset $1'$, without crossing any digons. Therefore, when S^3 is split open along $F \cup \Omega$ the arc β will go over to an arc $\beta^- \cup \beta^+$ in ∂T which joins 5 to 6 without crossing any digons, as in the right sketch in Figure 5.3. The arc β is drawn in thickly, as it occurs on H'_{ϕ_0} , in Figure 5.4(a) (next page). Call it $\mu_{56}(\phi_0)$ in what follows, to stress the fact that this arc is the component of $H'_{\phi_0} \cap (F \cup \Omega)$ which joins 5 to 6 .

We now turn our attention to other components of $H'_{\phi_0} \cap (F \cap \Omega)$. Stepping inside B_i , we look out at ∂B_i . See Figure 3.7. The subarc α_1 of A' is inside B_i , with its 2 endpoint in the F^+ part and its 3 endpoint in the F^- part of ∂B_i , and since $\alpha_1 \subset \partial H'_{\phi_0}$ there must be a component $\mu_{23}(\phi_0)$ of $H'_{\phi_0} \cap (F \cup \Omega)$ which joins 2 to 3 . We have sketched in the part of it which we know with a thick line.

Using Figure 5.2 and the right picture in Figure 3.7 we may describe $\mu_{23}(\phi_0)$ in another way. The arc $\mu_{23}(\phi)$ begins at the point 2 , which




FIGURE 5.5




FIGURE 5.6

L. Thus, going to Figure 5.4(c) we conclude that the only possibility is that $\mu_{23}(\phi_0)$ begins at 2 in F^+ , proceeds in set 2 to one of the adjacent digons, crosses that digon to subset 2 of F^- using a single digon crossing-arc of name 23, and then proceeds in subset 2 of F^- to 3. From Figure 5.4(c), we notice that the subarc α_1 of A' and $\mu_{23}(\phi)$ bound a disc δ in H'_ϕ . But then, we may push α_1 across δ into T . The fact that $\mu_{23}(\phi_0)$ only crosses one digon shows that this move decreases the complexity. Thus case (ii) does not occur.

Several of the techniques which we used to eliminate case (ii) will be used again when we treat the other cases, so we record them for future use. Call a fiber of H' *generic* if it does not contain a point of tangency with F or with Ω and if in addition it does not contain any of the crossing dots.

LEMMA 5.3. *Assume that the pierce-point p'_i is in subset #1' (or 1''). Then (after a modification of H') we may choose a generic fiber H'_ϕ with the property: $H'_\phi \cap F$ includes an arc β which is entirely in set #1' (or #1'') and joins p'_i to L .*




FIGURE 5.7

LEMMA 5.4. Suppose that p'_i and p'_{i+1} are in subsets #2 and/or #3 of $\mathbf{F} \cap \partial \mathbf{B}_i$. Let \mathbf{H}'_{ϕ_0} be any generic fiber. Then (after a modification of \mathbf{H}') we may assume that an arc in $\mathbf{H}'_{\phi_0} \cap \partial \mathbf{B}_i$ which joins p_i to p_{i+1} does not cross \mathbf{L} .

LEMMA 5.5. Let \mathbf{H}'_{ϕ_0} be any generic fiber. Then we may assume that $\mathbf{H}'_{\phi_0} \cap \Omega$ does not include any digon crossing-arcs of type [12], [11] or [24].

In addition, using the picture proof in Figure 5.7, we obtain:

LEMMA 5.6. Let \mathbf{H}'_{ϕ_0} be any generic fiber. Then we may assume that $\mathbf{H}'_{\phi_0} \cap \partial \mathbf{B}_i$ (or $\mathbf{H}'_{\phi_0} \cap \partial \mathbf{T}$) does not contain any simple closed curves.

We continue the proof of Proposition 5.1.

Case (iii). Assume that case (iii) of Figure 5.2 occurs. Our first task is to prove that if \mathbf{H}'_{ϕ} is a non-singular fiber of the non-standard fibration, then (as in case (ii)), there are no arcs of type [23] in $\mathbf{H}'_{\phi} \cap \Omega$. Suppose the contrary. Among all such arcs, choose one which is outermost, i.e. closest to A' , as in Figure 5.4(c). Without loss of generality we may assume that its name is 23. Passing to the picture on \mathbf{H}'_{ϕ} , as in Figure 5.8(b), we see that α_1 cobounds with $\beta^+ \cup \omega \cup \gamma^-$ a disc δ in \mathbf{H}'_{ϕ} . We may push α_1 across δ into \mathbf{T} , and the complexity will be reduced *if* we can show that this does not introduce other, unexpected, intersections of α_1 with other digons. To prove that, notice that the fact that ω was chosen to be innermost shows that $\text{int}(\beta^+) \subset \mathbf{F}^+$ and $\text{int}(\gamma^-) \subset \mathbf{F}^-$ are disjoint from all digons. But then, since p'_1 is in subset 1'' of \mathbf{F} , we see that the arcs $\beta^- \subset \mathbf{F}^-$ and $\gamma^+ \subset \mathbf{F}^+$ also do not meet any digons (see Figure 5.8(a) again).




FIGURE 5.8

Thus pushing α_1 across δ into \mathbf{T} , reduces complexity. Therefore there are no digon-crossing arcs of name 23.

The fact that p'_1, p'_2 and p'_3 are in distinct components of subset $1''$ of \mathbf{F} shows that we may apply Lemma 5.3 three times to choose a fiber \mathbf{H}'_{ϕ_0} such that $\mathbf{H}'_{\phi_0} \cap \partial\mathbf{B}_i$ is a union of arcs $\{\mu_{2i-1, 2i}(\phi_0), i = 1, 2, 3\}$ which join $2\mathbf{i} - \mathbf{1}$ to $2\mathbf{i}$, crossing \mathbf{L} once and missing all digons. Pushing \mathbf{H}'_{ϕ_0} forward in the fibration, we follow the evolving arcs. There must be a singularity in the foliation of $\partial\mathbf{B}_i$. By Proposition 4.8 the surgery is limited to the equivalence classes $[2 - 3], [2 - 5]$ or $[2 - 1]$. Since digon crossing arcs of type $[23]$ cannot occur, we conclude that the first surgery which occurs must be $[2 - 5]$ or $[2 - 1]$. Suppose it is $[2 - 5]$. Reversing the order of the fibration if necessary, we may assume it is $22 + 55 \rightarrow 25 + 52$. See Figure 5.9 (next page), with $2j = 2, 2i - 1 = 5$. On the left is a region on $\partial\mathbf{B}_i$ which contains the singularity, and on the right is the same region foliated without a singularity. There are no changes anywhere else in the foliation. The change in \mathbf{H}' can be realized by a local modification which does not introduce new singularities. Thus the complexity was not minimal. We are reduced to the case where the first surgery is type $[2 - 1]$. But then, by Corollary 4.9 the entire 3-ball \mathbf{B}_i is foliated with exactly one singularity. However, this contradicts the fact that p'_1, p'_2 and p'_3 are




FIGURE 5.9

in distinct components of subset $1''$. Thus case (iii) does not occur.

Case (iv). Assume that case (iv) occurs. By Lemma 5.4 applied 3 times, we may choose a fiber H'_{ϕ_0} of H' which meets ∂B_i in a simple arc $\mu_{23}(\phi_0)$ which joins 2 to 3 without crossing L , meets ∂B_{i+1} in a simple arc $\mu_{45}(\phi_0)$ which joins 4 to 5 without crossing L , and meets ∂T in a simple arc $\mu_{61}(\phi_0)$ in ∂T which joins 6 to 1 without crossing L . By Lemma 5.5 we may assume that $\mu_{23}(\phi_0)$ and $\mu_{45}(\phi_0)$ do not contain digon-crossing arcs of type [12], [11] or [24], so that the only possibilities for digon-crossing arcs are names 23, 45, 61, 14, 36, 52. Using these, the only way we can construct the required arcs $\mu_{23}(\phi_0)$ and $\mu_{45}(\phi_0)$ are by using digon-crossing arcs of name 45 and 61. Thus $H'_{\phi_0} \cap (F \cap \Omega)$ is as illustrated in Figure 5.10. The endpoint of either β_1 or of β_2 is closest to p'_2 , say β_1 . Then we may isotope α_1 through the subdisc δ_1 of H'_{ϕ_0} , pushing it into T . After that we may isotope α_2 through δ_2 to push A' into T . This reduces the complexity. We conclude that case (iv) does not occur.

Case (v). Assume that case (v) occurs. Using Lemma 5.5, select a fiber H'_{ϕ_0} such that $H'_{\phi_0} \cap \partial B_i$ includes an arc which joins 5 to 6, crossing L once but not intersecting any digons. Since the points 1, 4, 5, 6 are in ∂B_i it follows that there is a simple arc $\mu_{14}(\phi_0)$ in $H'_{\phi_0} \cap \partial B_i$ which joins the points $1 \subset F^-$ and $4 \subset F^+$ without crossing L . We may then apply Lemma 5.4 to conclude that when




FIGURE 5.10

$\mu_{14}(\phi_0)$ crosses a digon, the digon crossing-arc will have a name $pq \in \{13, 14, 24 \text{ or } 23\}$. However, the names 13 and 24 cannot occur by Lemma 5.5, and the name 23 cannot occur by Lemma 5.6, so pq can only be 14. Thus, we may assume that we have a disc fiber H'_{ϕ_0} for which $H'_{\phi_0} \cap (F \cap \Omega)$ is as illustrated in Figure 5.11(a) (next page). That is, $\mu_{14}(\phi_0)$ begins at $1 \subset F^-$, proceeds in F^- until it meets a digon Δ , crosses it in a single arc of name 14, and then proceeds to 4 in F^+ . The view we would have of $H'_{\phi_0} \cap B_i$ when we are inside B_i is shown in Figure 5.11(b). The point **5** (resp. **6**) will be in the F^- (resp. F^+) part of some $1''$ region on ∂B_i . The points **1** (resp. **4**) are in the F^- (resp. F^+) parts of subset 2 on ∂B_i . We suggest that the reader go back to Figure 3.7 (the decomposition of ∂B_i into subset $1'', 2, 3$) as an aid in understanding the various subcases which we now discuss.

Subcase v.1. The $1''$ region which contains **5** and **6** is adjacent to the digon Δ which $\mu_{14}(\phi_0)$ intersects. See Figure 5.11(c). Then we may push a portion of A' , say α_2 , into T as follows. First, we push the arc which has endpoints **5** and **6** into Δ to form a digon-intersection arc of name 55, as illustrated in Figure 5.11(c) and the sequence of pictures in Figure 5.11(d). Then we may form the digon singularity $14 + 55 \rightarrow 45 + 15$. The 45 arc which is created by this surgery is parallel in a disc fiber to α_2 , so we may isotope α_2 through the disc δ which is illustrated in the final snapshot in Figure 5.11(d). However, this contradicts minimum complexity.

Subcase v.2. The $1''$ region which contains **5** and **6** is not adjacent to Δ . Then we deform $\mu_{14}(\phi_0)$ as in Figure 5.11(e), so that a surgery $11 + 44 \leftrightarrow 14 + 41$ occurs. Now there are again two possibilities (easily understood by looking back at Figure 3.7): Either the deformed arc $\mu_{14}(\phi)$ intersects subset 3, or it does not. If it does, then digon arcs

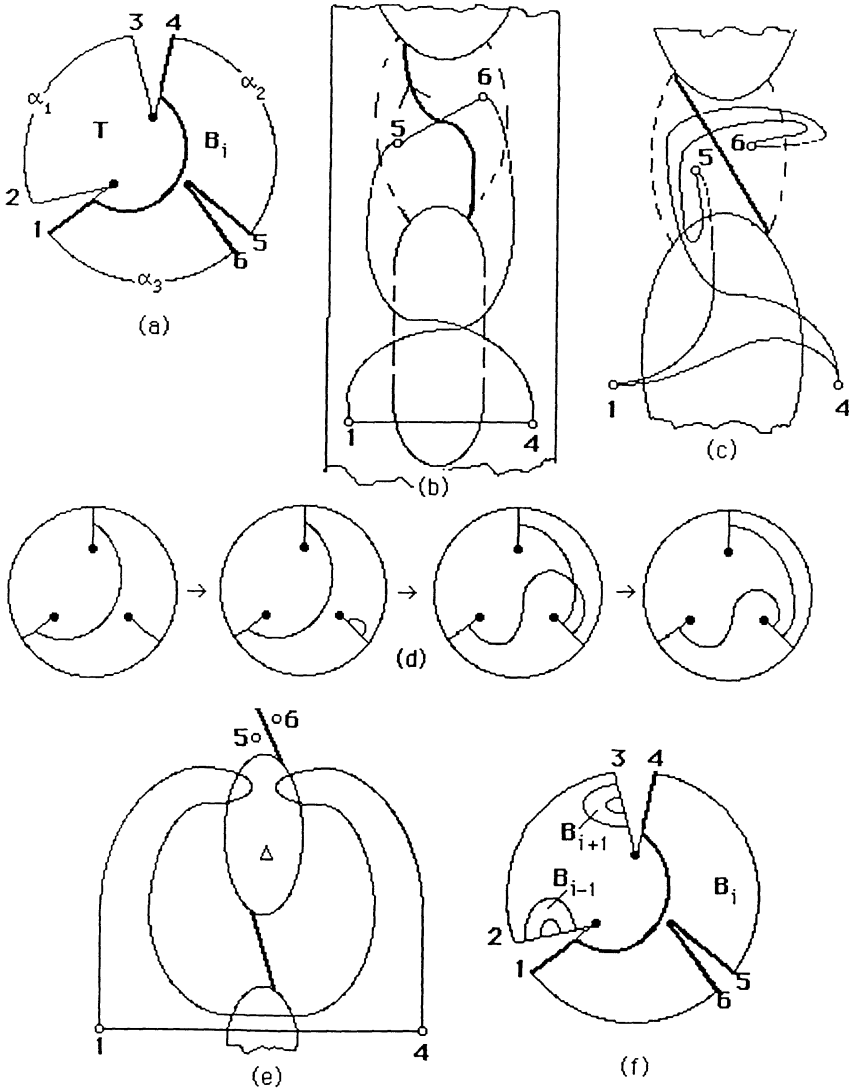


FIGURE 5.11

of name 22 or 33 will be introduced in pairs (two 22's or two 33's), as in Figure 5.11(f), because $\mu_{14}(\phi)$ will meet a type 3 region on both sides. Since the surgeries $pp + pp \rightarrow pp + pp$ cannot occur, the only possibility for foliating the intermediate digons is with a surgery $22 + 33 \rightarrow 23 + 32$. However, notice that the regions between the two 22's in H'_ϕ has label B_{i-1} and the region between the two 33's has label B_{i+1} , so this is impossible. We conclude that when $\mu_{14}(\phi)$ is deformed, it cannot intersect subset 3.

But then, after the $11 + 44 \leftrightarrow 14 + 41$ surgery, we will be back in

the situation of subcase v.1, with the same contradiction to minimum complexity as before. Thus case (v) does not occur.

We are reduced to case (i) of Figure 5.2. The cycle of labels is \mathbf{TTT} and the cycle of pierce-points is $1'1'1'$. The integer c_4 in the complexity 5-tuple is 0. The proof of Proposition 5.1 is complete. \square

A simple closed curve (scc) $\mathbf{c} = \mathbf{c}(\phi_0)$ in the foliation of $\partial\mathbf{T}$ is *essential* if there exists a parallel scc $\mathbf{c}(\phi)$ in the foliation which bounds a disc $\delta(\phi)$ which has non-empty intersection with \mathbf{L} . Otherwise, \mathbf{c} is inessential. A scc $\mathbf{c} = \mathbf{c}(\phi_0)$ in the foliation of $\partial\mathbf{B}_i$ is *special* if \mathbf{c} is disjoint from \mathbf{L} and is a union $\alpha_1 \cup \beta_1 \cup \alpha_2 \cup \beta_2$ of four arcs: an arc β_1 in \mathbf{F}^- , an arc β_2 in \mathbf{F}^+ , and two arcs α_1 and α_2 , each of which crosses a single digon in $\partial\mathbf{B}_i$ once. A special scc \mathbf{c} is *essential* if there exists a parallel simple closed curve $\mathbf{c}(\phi)$ in the foliation such that each component of $\partial\mathbf{B}_i$ split open along $\mathbf{c}(\phi)$ has non-empty intersection with \mathbf{L} . Otherwise, \mathbf{c} is inessential.

LEMMA 5.7. *We may assume that every scc (resp. special scc) \mathbf{c} on $\partial\mathbf{T}$ (resp. $\partial\mathbf{B}_i$) is essential.*

Proof. Let \mathbf{c} be an inessential scc in the foliation of $\partial\mathbf{T}$. Then \mathbf{c} is a union of arcs in $\mathbf{F}^+ \cup \mathbf{F}^- \cup \Omega$, and can be viewed both as a scc in $\partial\mathbf{T}$ and as a scc in a fiber of \mathbf{H}' . Looking at \mathbf{c} as it occurs in fibers of \mathbf{H}' it is easy to see that by replacing \mathbf{c} by parallel copies if necessary, we may assume that a digon-crossing arc α in \mathbf{c} has one of its endpoints in \mathbf{F}^+ and the other in \mathbf{F}^- . Thus α is type [12] or [14] or [16]. However, if α were type [12] then by Proposition 4.8, part 1, we could push \mathbf{c} through the foliation to eliminate α without encountering any singularities. But then there is a \mathbf{c}' which is parallel to \mathbf{c} and which meets \mathbf{L} non-trivially, so \mathbf{c} was essential. Thus α has type [14] or [16]. In fact, since \mathbf{c} cannot cross \mathbf{L} , the only possibility is that up to symmetries \mathbf{c} contains exactly two parallel digon crossing arcs, both of type 14 or both of type 16. Thus an inessential scc \mathbf{c} in $\partial\mathbf{T}$ is a special scc.

From now on we allow \mathbf{c} to be in either $\partial\mathbf{B}_i$ or $\partial\mathbf{T}$. Pushing \mathbf{c} forward or backward in the foliation, we encounter the first homoclinic point, which will be on a singular leaf \mathbf{c}' parallel to \mathbf{c} . See Figure 5.12 (next page). Now \mathbf{c}' bounds two discs, one in a fiber of \mathbf{H}' and the other in $\partial\mathbf{T}$ or $\partial\mathbf{B}_i$. These discs fit together to form an S^2 which bounds a B^3 which is disjoint from \mathbf{L} (because \mathbf{L} is non-split). We can then modify the fibration by pushing fibers of \mathbf{H}' across the B^3




FIGURE 5.12

to eliminate the scc. Thus without loss of generality we may assume that every simple closed curve in the foliation of $\partial\mathbf{T}$ and every special simple closed curve in the foliation of $\partial\mathbf{B}_i$ is essential. \square

COROLLARY 5.8. *We may assume that there is a fiber \mathbf{H}'_{ϕ_0} with the property: $\mathbf{H}'_{\phi_0} \cap \partial\mathbf{T}$ is a union of three arcs μ_{12} , μ_{34} , and μ_{56} , where $\mu_{2i-1, 2i}$ joins the pierce points $2i-1$ and $2i$, crossing L once, without intersecting any digons.*

Proof. This follows directly from Lemmas 5.3 and 5.7. For, by Lemma 5.3 we know we may find a fiber of \mathbf{H}' which satisfies the property we need for at least one i , say $\mathbf{H}'_{\phi_0} \cap \partial\mathbf{T}$ contains the arc μ_{12} . But then if μ_{34} and μ_{56} did not have the stated properties there would be inessential simple closed curves in the foliation of $\partial\mathbf{T}$, a possibility which is ruled out by Lemma 5.7. \square

6. The Classification Theorem. All of the pieces are in place, and we are ready to prove the theorem which is the central result of this manuscript. Our proof will occupy the rest of this section.

THE CLASSIFICATION THEOREM (Version 2). *Let \mathcal{L} and \mathcal{M} be oriented link types of braid index ≤ 3 in oriented 3-space. Let \mathbf{L} , \mathbf{M} be closed 3-braids which represent \mathcal{L} and \mathcal{M} , and have the same braid axis \mathbf{A} . Let $\mathbf{W} = W(a_1, a_2, a_3)$ and $\mathbf{X} = X(a_1, a_2, a_3)$ be words which define the conjugacy classes of \mathbf{L} and \mathbf{M} in B_3 . Then the link types \mathcal{L} and \mathcal{M} , coincide if and only if \mathbf{W} is conjugate to \mathbf{X} , with*

the following exceptions:

(i) \mathbf{W} is conjugate to a_1a_2 and \mathbf{X} to $a_1a_2^{-1}$ or $(a_1a_2)^{-1}$ and \mathbf{L} represents the unknot, which has braid index 1.

(ii) \mathbf{W} is conjugate to $a_1^p a_2$ and \mathbf{X} to $a_1^p a_2^{-1}$, where $|p| \neq 1$, and \mathbf{L} represents a type $(2, p)$ torus link, which has braid index 2.

(iii) \mathbf{W} is conjugate to $(a_i)^p (a_n)^q (a_m)^r$ or $(a_i)^p (a_n)^q (a_m)^r (a_n)^{\pm 1}$, where $\{i, n, m\} = \{1, 2, 3\}$ and \mathbf{X} is conjugate to \mathbf{W} read backwards. The 3-braid \mathbf{L} represents a prime, non-split, invertible link of braid index 3 which has at most two conjugacy classes of 3-braid representatives. These classes are related by braid-preserving flypes.

Proof of the Classification Theorem. W. Magnus and A. Pelluso proved in [M-P] that the unknot has three conjugacy classes of 3-braid representatives. The three are defined by the admissible words $(a_1a_2)^{\pm 1}$ and $a_1a_2^{-1}$. This is case (i) of the classification theorem.

A complete list of the links of braid index 2 includes the type $(p, 2)$ torus links, $|p| \neq 1$, and the 2-component unlink. It was shown by Murasugi in [Mu] that each of these has two conjugacy classes of 3-braid representatives, defined by admissible words $a_1^p a_2^{\pm 1}$, $p \in \mathbb{Z}$. This covers case (ii) of the classification theorem. Thus we are reduced to links of braid index 3.

Morton proved in [Mo, 1] that composite links of braid index 3 admit unique conjugacy classes of 3-braid representatives. See [B-M, 1] for a proof that split links of braid index 3 admit unique conjugacy classes of 3-braid representatives. Thus we are reduced to the case where \mathcal{L} is prime, non-split, and has braid index 3. These are the hypotheses we used in §§2–5 of this manuscript. We ask when such an \mathcal{L} can be represented by more than one conjugacy class in B_3 ?

It will be convenient to adopt the point of view which has been used everywhere in this paper up to now, i.e. we begin with a single representative \mathbf{L} of the link type \mathcal{L} and two axes \mathbf{A} and \mathbf{A}' for 3-braid representatives. As was observed at the end of §2, there is then a homeomorphism $h: S^3 \rightarrow S^3$ such that $h(\mathbf{A}') = \mathbf{A}$; also h maps fibers of \mathbf{H} to fibers of \mathbf{H}' and maps \mathbf{L} to a new representative $\mathbf{M} = h(\mathbf{L})$. If $\mathbf{W}' = W'(b_1, b_2, b_3)$ describes \mathbf{L} relative to the axis \mathbf{A}' , we obtain a defining word \mathbf{X} in the elementary braids a_1, a_2, a_2 which represents \mathbf{M} by setting $\mathbf{X} = W'(a_1, a_2, a_3)$.

We review what we learned in §§2–5. By Corollary 2.3 we may assume that both \mathbf{W} and \mathbf{W}' are admissible, i.e. they are shortest words, also among all shortest words they have shortest syllable length,

and also they use all three generators. By Proposition 5.1 we know that the non-standard axis A' is in the interior of T . By Proposition 3.2 we know that there is a natural link diagram for L on ∂T . We begin to investigate the non-standard foliation of ∂T . We are interested in its Ω -singularities and F -singularities. By Proposition 4.8 the Ω -singularities are restricted to types $[1 - 2]$, $[1 - 4]$ and $[1 - 6]$. By Corollary 4.9 there cannot be more than three singularities of type $[1 - 2]$, since there are only three distinct letters a_i .

Suppose that the non-standard foliation contains three singularities of type $[1 - 2]$. Then Corollary 4.9 implies that W has letter length three and uses all three generators. Up to conjugation and possibly replacing L and M by their mirror images we may assume the first letter is a_1 . So $W = a_1(a_n)^\beta(a_m)^\gamma$, where $\beta, \gamma = \pm 1$ and $\{n, m\} = \{2, 3\}$. Of the 8 possible words of this form, we see from relations (2.1) that the only one which cannot be replaced by a word of syllable length 2 is $W = a_1a_2a_3$. Now W' is also admissible, so it also contains all three letters. Since the genus of F is an invariant of the link type \mathcal{L} , we know that W' also has letter length 3. By [J], the exponent sum of W is a link type invariant for links of braid index 3, so every letter in W' is positive. But then, by the argument just given for W the only possibility is $W' = b_1b_2b_3$, which implies that $X = a_1a_2a_3 = W$. Thus \mathcal{L} is represented by a unique conjugacy class in B_3 .

Now suppose that there are *two* singularities of type $[1 - 2]$ in the nonstandard fibration. Then two of the braid generators a_i and a_k each occur once in W , with exponent ± 1 . Up to conjugation we may assume that $(a_1)^{\pm 1}$ and $(a_k)^{\pm 1}$ occur once, where $k = 2$ or 3 . Replacing W by its inverse if necessary we may assume that a_1 and $(a_k)^{\pm 1}$ occur once. The syllables a_1 and $(a_k)^{\pm 1}$ in W could either be separated by powers of a_r or not.

If a_1 and $(a_k)^{\pm 1}$ are separated by powers of a_r then the cyclic word W is $(a_k)^{\pm 1}(a_r)^p a_1(a_r)^q$. But then, using relations (2.4), we see that if $(k, r) = (2, 3)$ we may replace W by $(a_2)^{\pm 1}(a_3)^p(a_2)^q a_1$, whereas if $(k, r) = (3, 2)$ we may replace W by $(a_3)^{\pm 1} a_1(a_3)^p(a_2)^q$. In the revised forms for W the only braid generator which occurs exactly once is a_1 . Therefore after the change in admissible word the non-standard fibration will have at most one singularity of type $[1 - 2]$.

If a_1 and $(a_k)^{\pm 1}$ are not separated by powers of a_r then $W = a_1(a_k)^\beta(a_k)^p$, where $\beta = \pm 1$ and where $\{k, r\} = \{2, 3\}$. We now argue as we did in the case where there are three singularities of type




FIGURE 6.1

[1 – 2], to conclude that $\mathbf{W} = a_1 a_2 a_3^p$, with $|p| \geq 2$. Interchange the roles of the standard and non-standard fibrations. The new non-standard fibration has $n = 3, 2, 1$ or 0 singularities of type [1 – 2]. Now if $n = 3$ we are done by the argument used earlier for the standard fibration. If $n = 2$, then \mathbf{W}' or $(\mathbf{W}')^{-1}$ must also have the form $a_1 a_2 (a_3)^q$, where $|q| \geq 2$. If $\mathbf{W} = a_1 a_2 a_3^p$, $\mathbf{W}' = a_1 a_2 a_3^q$, then the fact that \mathbf{W} and \mathbf{W}' have the same exponent sum implies that $p = q$ and $\mathbf{W} = \mathbf{W}'$. If $\mathbf{W} = a_1 a_2 a_3^p$, $(\mathbf{W}')^{-1} = a_1 a_2 a_3^q$, then the exponent sum argument gives $p + 2 = -q - 2$, which implies that $p = -q - 4$. On the other hand, the fact that \mathbf{F} is a surface of maximal Euler characteristic for both \mathbf{L} and \mathbf{L}' , together with the fact that the Euler characteristic depends upon the number of bands (but not upon their signs) shows that $|p| = |q|$ or $p = \pm q$. The case $p = -q$ cannot occur because $p = -q - 4$. So $p = q = -2$. Thus \mathbf{W} and \mathbf{W}' both represent the Figure 8 knot and $\mathcal{L} = \mathcal{L}'$. Therefore from now on we may assume that there is at most one occurrence of [1 – 2].

The \mathbf{F} singularities are types $(b_1)^{\pm 1}$, $(b_2)^{\pm 1}$, $(b_3)^{\pm 1}$. They may be further subdivided, as follows:

(i) It may happen that there is solid cylinder $D^2 \times 1$ in $S^3 - \mathbf{A}$ such that each fiber $D^2 \times \{t\}$ is included in both a fiber of \mathbf{H} and a fiber of \mathbf{H}' ; also the link \mathbf{L} meets each fiber $D^2 \times \{t\}$ of this cylinder twice, transversally. In this situation, if an \mathbf{F} -singularity occurs inside the cylinder, it will be an \mathbf{F} singularity in both fibrations. That is, a band in the standard disc-band decomposition is foliated in the same way by the standard and non-standard fibrations. Call this a *standard* \mathbf{F} -singularity. A *block* is a maximal sequence of standard \mathbf{F} -singularities, as depicted in Figure 6.1. We use bold-faced type $(\mathbf{b}_i)^p$ to indicate that the syllable $(b_i)^p$ in \mathbf{W}' comes from a block.

(ii) The second possibility is that the foliation near an \mathbf{F} -singularity is not isotopic to the standard foliation. We call this a *non-standard* or NSF-singularity.

Thus the non-standard fibration of $S^3 - \mathbf{L}$ is described by a sequence of events, each of which is either a block of standard \mathbf{F} -singularities, an NSF-singularity or an Ω -singularity. We call the NSF and Ω -singularities *events*. The *cycle of events* is the cyclic list of NSF-singularities (described by a single letter $(b_j)^{\pm 1}$, $j = 1, 2, 3$) and Ω -singularities (described by a type symbol $(p - q)^{\rightarrow}$ or $(p - q)^{\leftarrow}$). The *augmented cycle of events* is obtained from the cycle of events by adding the blocks (each being described by a symbol $(\mathbf{b}_k)^m$), where $k = 1, 2, 3$ and where $m \in \mathbb{Z} - \{0\}$.

By Corollary 5.8 we may choose an *initial fiber* \mathbf{H}'_{ϕ_0} which is disjoint from all of the digons and meets $\partial\mathbf{T}$ in three arcs, an arc μ_{12} which joins the pierce-points $\mathbf{1}$ and $\mathbf{2}$, an arc μ_{34} which joins $\mathbf{3}$ and $\mathbf{4}$ and an arc μ_{56} which joins $\mathbf{5}$ and $\mathbf{6}$. In each case the arc beings in \mathbf{F}^- and crosses \mathbf{L} once and passes into \mathbf{F}^+ , meeting no digons. In what follows we will always assume that our cycles begin with such an initial fiber.

We now prove that the cycle of events contains at least four distinct events. The *initial forward* (resp. *backward*) sequence is the cycle of events, starting with an initial fiber and pushing it forward (resp. backward). Notice that the initial forward (resp. backward) sequence contains at least two events, for the following reason: The foliation of \mathbf{F} is radial in a neighborhood of each of the points where \mathbf{A}' pierces \mathbf{F} . Therefore we know that each of the three arcs μ_{12} , μ_{34} , μ_{56} must be modified by a singularity which “changes its connectivity” at some time during the complete augmented cycle. For example, after a singularity of type $(2 - 3)^{\rightarrow}$, the arcs $\mu_{12}(\phi)$ and $\mu_{34}(\phi)$, which join $\mathbf{1}$ to $\mathbf{2}$ and $\mathbf{3}$ to $\mathbf{4}$ will be replaced by arcs $\mu_{14}(\phi)$ and $\mu_{23}(\phi)$ which join $\mathbf{1}$ to $\mathbf{4}$ and $\mathbf{2}$ to $\mathbf{3}$. Since blocks do not change connectivity, whereas NSF and Ω -singularities do, this means that there must be at least two NSF and/or Ω -singularities in the initial forward sequence. But then, the same must be true for the initial backward sequence, because the forward and backward sequences cannot meet unless the connectivities agree. Therefore the cycle of events has length at least four.

The proof of the Classification Theorem will be completed in three main steps: the first main step is to prove (Lemma 6.1) that the conditions of the theorem are satisfied if the augmented cycle of events

contains only blocks and Ω -singularities of type $[1 - 4]$. The second main step is to prove (Lemma 6.4) that the conditions of the theorem are also satisfied if the augmented cycle of events contains no singularities of type $[1 - 4]$. The third main step (Lemma 6.5) is to prove that the augmented cycle of events either contains only blocks and Ω -singularities of type $[1 - 4]$, or else it does not contain any Ω -singularities of type $[1 - 4]$. Thus every case is covered by Lemmas 6.1 and 6.4.

The basic technique which we use is explained in some detail in the proof of Lemma 6.1, below. That lemma will be followed by two others, Lemmas 6.2 and 6.3, which apply the same general methods to many different cases. After we have all the cases in hand it will be easy to prove the main Lemmas 6.4 and 6.5.

LEMMA 6.1. *Assume that the cycle of events contains only Ω -singularities in the equivalence class $[1-4]$. Then \mathbf{W} and \mathbf{X} both have syllable length three, and \mathbf{X} is \mathbf{W} read backwards.*

Proof of Lemma 6.1. We push \mathbf{H}'_{ϕ_0} forward in the fibration to a fiber \mathbf{H}'_{ϕ_1} which is just after the first event. By our hypotheses that event is a singularity in the equivalence class $[1 - 4]$. Since there are no arcs of type $[14]$ in \mathbf{H}'_{ϕ_0} we know that singularity cannot be type $[1 - 4]^-$, so it must be $[1 - 4]^+$. Up to symmetry we may assume it is $(1 - 4)^+$.

The evolving situation in fibers of \mathbf{H}' , as we approach and pass the $(1 - 4)^+$ singularity, is depicted in the three snapshots at the top of Figure 6.2 (next page). The corresponding picture on $\partial\mathbf{T}$ near this first singularity is then as illustrated in the bottom picture in Figure 6.2. To see this, the first thing to notice is that since the singularity is preceded by the appearance of a pair of digon-crossing arcs of name 11 and 44, the leaves μ_{12} and μ_{34} which are closest to the pierce-points 1 and 4 in \mathbf{H}'_{ϕ_0} must be deformable in $\partial\mathbf{T} - \mathbf{L} \cup \Omega$ to arcs which are on opposite sides of the same digon. Since the pierce-points 1 and 2, and also 3 and 4, are on opposite sides of \mathbf{L} this determines the location of the four pierce-points (up to an isotopy in which they are moved without crossing \mathbf{L} or penetrating any digon). The second thing to notice is that the $(1 - 4)^+$ singularity creates a pair of parallel digon-crossing arcs α_1 and α_2 of type 14 in, say, \mathbf{H}'_{ϕ_1} . These arcs α_1 and α_2 , together with an arc β_1 in $\mathbf{F}^+ \cap \mathbf{H}'_{\phi_1}$ and an arc β_2 in $\mathbf{F}^- \cap \mathbf{H}'_{\phi_1}$, form a simple closed curve $\mathbf{c} = \mathbf{c}(\phi_1) = \alpha_1 \cup \beta_1 \cup \alpha_2 \cup \beta_2$.




FIGURE 6.2

This curve bounds a disc d in $H'_{\phi_{-1}}$, which (by the discussion in the proof of Lemma 5.7) is necessarily in one of the 3-balls, say B_i . The subarcs α_1 and α_2 are also in ∂T , but β_1 and β_2 are not.

Figure 6.3 shows the curve c as it might look on ∂B_i . In our example there is a chain of two digons, joined up by undercrossing arcs (cf. Figures 3.7 and 3.9). Our curve c meets one of the digons in two arcs (α_1 and α_2) and bounds a disc δ , which also intersects the digon. The curve c is a *special* simple closed curve in ∂B_i so by Lemma 5.7 it must be essential. This means that the disc δ must be a subdisc of a larger disc δ' on ∂B_i , which contains at least one undercrossing arc and which is foliated by non-singular simple closed curves parallel to c . Also $\partial\delta'$ is a special simple closed curve c' which




FIGURE 6.3

is parallel to c and which intersects distinct digons. Thus δ' contains one or more of the undercrossing arcs. This means that the singularity must be followed by a block $P = \mathbf{b}_i^p$, $p \neq 0$. The integer p is the number of undercrossing arcs between α_1 and α'_1 , in δ' , on $\partial \mathbf{B}_i$, when ϕ_2 is chosen so that p is maximal. The digon-crossing arcs α_1 and α'_1 both have name 23, and they are separated on ∂T by the block. Thus we have shown: if the cycle of singularities begins with an initial fiber, which is followed by the singularity $(1 - 4)^{\rightarrow}$, then the second event must be a non-empty and maximal block \mathbf{b}_i^p .

The argument we have just given is sufficiently important so that we interrupt our proof to state a generalized version, for future use:

SUBLEMMA 6.1.1. *A digon singularity of type $[1 - 4]^{\rightarrow}$ or $(2 - 3)^{\rightarrow}$ creates a special simple closed curve in the boundary of some \mathbf{B}_i . The digon singularity must be followed by a non-empty block $(\mathbf{B}_i)^p$, for if not this curve is not essential.*

Proof of Sublemma 6.1.1. We proved Sublemma 6.1.1 for the case of $(1 - 4)^{\rightarrow}$. The proof for $(2 - 3)^{\rightarrow}$ is identical, with one small modification: the pierce points 1, 2 and 3, 4 on ∂T must be changed from their positions in Figure 6.2 to their positions in Figure 6.4 (next page).

We return to the proof of Lemma 6.1. What can happen after the block P ? Let \mathbf{H}'_{ϕ_3} be a fiber of \mathbf{H}' just after the singularity which follows the block. By the hypotheses of the lemma, we are only allowed to have singularities of type $[1 - 4]$ or blocks. Since p is maximal, and since there is only one arc in $\mathbf{H}'_{\phi_2} \cap \partial T$ which crosses L twice, no further blocks can occur. Singularities of type $(2 - 5)^{\rightarrow}$ or $(3 - 6)^{\rightarrow}$ are impossible, because the two digon-crossing arcs of name 14 are an obstruction in \mathbf{H}'_{ϕ_2} . All singularities of type $[1 - 4]^{\leftarrow}$ except for




FIGURE 6.4

$(1-4)^{\leftarrow}$ are impossible, because the initial arcs which are required for such a singularity don't exist. So we are reduced to $(1-4)^{\rightarrow}$ or $(1-4)^{\leftarrow}$.

We now claim that the initial forward sequence cannot begin with $(1-4)^{\rightarrow}$, $(1-4)^{\leftarrow}$. For, suppose that it did. After the second $(1-4)^{\rightarrow}$ the picture in H'_{ϕ_3} will be as illustrated in Figure 6.5. There are now two pairs of parallel digon crossing arcs of name 14. Each pair, together with subarcs of F^+ and F^- , splits off a disc in the fiber of H' , and this disc will be in one of the regions B_1 , B_2 or B_3 . We label it accordingly. After the second singularity there are two such regions, with labels B_1 and B_k , and a region between them with label T . In principle, we could have $i = k$ or $i \neq k$. There are three special simple closed curves in the final snapshot: c in ∂B_i , c' in ∂B_k and c'' in ∂T . All must be essential. The four digon-crossing arcs of name 14 are all in ∂T .

We now observe that if $i = k$ there is no way for the curve c'' to be essential, because the second singularity would necessarily occur in a digon which is directly below the block P of Figure 6.2 and by Sublemma 6.1.1 that digon would necessarily be followed by a block Q which joins the same two strands, so the two blocks could be amalgamated and the digon-singularity eliminated. Thus $i \neq k$.

If $i \neq k$ the local picture on ∂T after the second singularity must be as in Figure 6.5. The curve c'' is essential. However W is seen to include a subword $(a_1)^p(a_2)^{-1}(a_1)^{-1}(a_2)^q$, where $p \neq 0$, $q \neq 0$. By relations (2.2) such a word is equivalent to $(a_1)^{p+q}(a_2)^{-1}(a_1)^{-1}$,




FIGURE 6.5

which has shorter syllable length. So, W is not admissible and this case does not occur.

We are reduced to the case where the initial forward sequence begins with $(1 - 4)^\rightarrow$, followed by $(1 - 4)^\leftarrow$. Recall that the $(1 - 4)^\rightarrow$ and subsequent block resulted in the foliation of a disc δ' , with $\partial\delta' = c'$, on ∂B_i . After the $(1 - 4)^\leftarrow$ singularity we will have foliated a second disc, also with boundary c' , in ∂B_i . Thus we will have foliated an S^2 in ∂B_i . But then we must have foliated all of B_i . This means that the braid generator b_i occurs as a single block b_i^p in W' . Thus b_i occurs as a single syllable b_i^p in W' , and a_i occurs as a single syllable a_i^p in W .

When the sequence $(1 - 4)^\rightarrow$, block, $(1 - 4)^\leftarrow$ is completed we will have a new initial fiber. The next singularity could be either of the two singularities which are symmetrically equivalent to $(1 - 4)^\rightarrow$, i.e. $(3 - 6)^\rightarrow$ or $(5 - 2)^\rightarrow$. The argument we gave for $(1 - 4)^\rightarrow$ applies equally well to these cases. Thus there is exactly one syllable a_1^p , one syllable a_2^q and one syllable a_3^r in W , also W' contains three syllables b_1^p, b_2^q, b_3^r . Thus X contains three syllables a_1^p, a_2^q, a_3^r . But then, the only possible difference between W and X is the order in which the three syllables occur. By hypothesis L has two distinct axes, so X and W cannot be identical cyclic words. We conclude that the cyclic word X is the cyclic word W read backwards. \square

In the proof of Lemma 6.1 we showed that certain initial sequences could not occur. The next lemma generalizes the phenomena which ruled them out.

LEMMA 6.2. *The cycle of events does not begin with any of the following, or with any other sequences equivalent to these under symmetries:*

$$(6.2.1) \quad (1 - 4)^{\rightarrow}, (1 - 4)^{\leftarrow},$$

$$(6.2.2) \quad (2 - 3)^{\rightarrow}, (2 - 3)^{\leftarrow},$$

$$(6.2.3) \quad (b_1)^{-1}, (b_1)^{-1},$$

$$(6.2.4) \quad b_1, (b_1)^{-1} \text{ or } (b_1)^{-1}, b_1.$$

Proof of Lemma 6.2.

Proof of 6.2.1 and 6.2.2. Assertion 6.2.1 was proved in the course of the proof of Lemma 6.2.1. The proof of 6.2.2 is identical, using the modification described in the proof of Sublemma 6.1.1.

Proof of 6.2.3. To prove the assertion we need to understand the foliation of $\partial\mathbf{T}$ in the region near the singular leaves for an NSF-singularity. See Figure 6.6, which depicts the case of a singularity of type b_1^{-1} . (All other cases can be obtained from this one by symmetries.) It will be helpful to begin by studying that singularity on \mathbf{F} , rather than on $\partial\mathbf{T}$. The singularity affects the leaves of the foliation of \mathbf{F} which emerge from the pierce-points p'_1 and p'_2 . See Figure 6.6(a). Recall that by Lemma 5.2 the points p'_1 and p'_2 are in distinct components of subset $1'$ of \mathbf{F} . By our assumptions about \mathbf{H}'_ϕ at $\phi = \phi_0$ there will be a leaf $\pi(\mu_{12}(\phi_0))$ in the foliation of \mathbf{F} which joins p'_1 to \mathbf{L} in set $1'$ and another arc $\pi(\mu_{34}(\phi_0))$ which joins p'_2 to \mathbf{L} in subset $1'$. As ϕ increases these arcs will experience the \mathbf{F} singularity under investigation. Since $\pi(\mu_{12}(\phi_0))$ and $\pi(\mu_{34}(\phi_0))$ are in *distinct* components of subset $1'$, it follows that the evolving arcs must cross the digon attaching curves, which are illustrated in Figure 6.6(a) as dotted arcs, because they are on \mathbf{F}^- .

Figure 6.6(b) shows the identical situation on $\partial\mathbf{T}$. It can be understood by noticing that subset $\#1'$ of \mathbf{F} opens up to two regions on $\partial\mathbf{T}$ which meet along \mathbf{L} . There is a subtle and important point. As the positive sides of $\mu_{12}(\phi)$ and $\mu_{34}(\phi)$ come together for the future singularity, we will see them penetrate into nearby digons, and since the digon attaching curve which is in the region of interest is on \mathbf{F}^- , the penetration will be by digon crossing arcs on the \mathbf{F}^- part of $\partial\mathbf{T}$. Figure 6.6(c) shows the same situation on $\partial\mathbf{T}$, but redrawn to emphasize that in order for this singularity to be possible, there must be a joining arc between μ_{12} and μ_{34} on the \mathbf{F}^+ side of the braid strand. Figure 6.6(d) shows the same situation on a sequence of fibers of \mathbf{H}' . As the




FIGURE 6.6

arcs of $H'_\phi \cap F$ come together for the future type $(b_1)^{-1}$ -singularity we see a pair of digon crossing arcs of name 11 and 33 appear. The endpoints of each are separated by the endpoint of the joining arc for the singularity. Therefore, after the singularity there will be a pair of parallel digon-crossing arcs of name 13. If, instead, the singularity had been type b_1 , the corresponding digon-crossing arcs would have name 24.

It is now easy to see why the sequence $(b_1)^{-1}, (b_1)^{-1}$ cannot occur. The first $(b_1)^{-1}$ creates a pair of digon-crossing arcs of name 13 and these arcs are obstructions in the fibers of H' to a joining arc for the second $(b_1)^{-1}$.

Proof of 6.2.4. \mathbf{W}' is freely reduced, so b_i and $(b_i)^{-1}$ cannot be adjacent in \mathbf{W}' . \square

For future use, we note that events $b_j^{\pm 1}$, $[1 - 6]^{\rightarrow}$ and $[1 - 4]^{\rightarrow}$ each result in the creating of a pair of parallel digon crossing arcs in fibers of \mathbf{H}' . We call these arcs *bridges*. The *name* of the bridge is the name of its digon-crossing arcs. The *bridge region* is the region on fibers of \mathbf{H}' cut off by the bridge together with subarcs of \mathbf{F}^{\pm} . The *boundary* of the bridge is the simple closed curve \mathbf{c} which bounds the bridge region. The *label* of the bridge is \mathbf{B}_i if the bridge region is contained in the 3-ball \mathbf{B}_i .

LEMMA 6.3. *We may assume that the initial forward sequence in the cycle of events does not begin with any of the following sequences, or with other sequences equivalent to these under symmetries:*

$$(6.3.1) \quad (1 - 4)^{\rightarrow}, (4 - 5)^{\rightarrow},$$

$$(6.3.2) \quad (b_1)^{-1}, (2 - 3)^{\rightarrow},$$

$$(6.3.3) \quad (b_1)^{-1}, b_2, (2 - 3)^{\rightarrow},$$

$$(6.3.4) \quad (b_1)^{-1}, b_2, (4 - 5)^{\rightarrow},$$

$$(6.3.5) \quad (b_1)^{-1}, (b_2)^{-1}, (4 - 5)^{\rightarrow},$$

$$(6.3.6) \quad (2 - 3)^{\rightarrow}, (b_1)^{-1},$$

$$(6.3.7) \quad (b_2)^{-1}, (2 - 3)^{\rightarrow}, (1 - 6)^{\rightarrow},$$

$$(6.3.8) \quad (1 - 4)^{\rightarrow}, (2 - 3)^{\rightarrow},$$

$$(6.3.9) \quad (2 - 3)^{\rightarrow}, (4 - 5)^{\rightarrow}.$$

Also, the following sequences are equivalent after a change in \mathbf{H}' :

$$(6.3.10) \quad (2 - 3)^{\rightarrow}, (b_2)^{-1} \text{ and } (b_2)^{-1}, (2 - 3)^{\rightarrow}.$$

Proof of Lemma 6.3. The arguments are all variations on ones used earlier. The reader will not lose the thread of the argument by omitting the proofs.

Proof of 6.3.1. Suppose that $(1 - 4)^{\rightarrow}, (4 - 5)^{\rightarrow}$ occurs. Let \mathbf{H}'_{ϕ_2} be a fiber on \mathbf{H}' after the $(4 - 5)^{\rightarrow}$ event. By Sublemma 6.1.1 the $(4 - 5)^{\rightarrow}$ event must be followed by a block \mathbf{b}_1^p which will move one of the type 45 arcs onto a different digon in $\partial\mathbf{B}_i$. However, from Figure 6.7 we see that there are only three leaves in the foliation of $\partial\mathbf{T}$ at $\phi = \phi_2$, and the two which contain the 45 digon crossing arcs are the leaves which join **4** to **5** and **6** to **1**. The former does not cross \mathbf{L} at all and the latter crosses \mathbf{L} once. Since a leaf in the foliation can't sweep through a block unless it crosses \mathbf{L} twice, we conclude





FIGURE 6.7

that $p = 0$, contradicting Sublemma 6.1.1. Therefore this case does not occur.

Proof of 6.3.2, 6.3.3, 6.3.4, 6.3.5. The arguments are all essentially the same as the one used to prove 6.3.1. The final event is a singularity of type $(q - r)^{\rightarrow}$, where $qr = 23, 23, 45, 45$ in the four cases. By Sublemma 6.1.1 that bridge must be followed by a non-empty block. Part of the picture on a fiber of H' after the $(q - r)^{\rightarrow}$ singularity is like the right picture in Figure 6.7. Neither of the type or bridge arcs is on a leaf of the foliation of ∂T which crosses L twice, so no such block is possible.

Proof of 6.3.6. The right picture in Figure 6.7, with $p, q, r, s = 1, 2, 3, 4$, may be identified with the picture on a fiber of H' after the two events $(2 - 3)^{\rightarrow}, (b_1)^{-1}$. The singularities create two bridges, the first with name 23 and label B_k and the second with name 13 and label B_i . Figure 6.8 (next page) shows the foliation of ∂T in the two cases $i = k$ and $i \neq k$. The techniques for constructing it are just like those we used in the proof of Lemma 6.1. Both begin with the local picture in Figure 6.4, after the $(2 - 3)^{\rightarrow}$ singularity. In order for a $(b_1)^{-1}$ singularity to occur next, there must be a joining arc in F^+ between the “2” side and the “4” side of the arc which joins pierce-points 1 and 3 in Figure 6.4. If $i = k$ the local picture must be as illustrated on the left in Figure 6.8. But then, the block P is not maximal. If P is extended to include the crossing just below it the complexity will be reduced. So this case cannot occur.



Singular leaves in the foliation of ∂T after $(2-3) \rightarrow, b_1^{-1}$

FIGURE 6.8

If $i \neq k$ the local picture on ∂T must be the one on the right in Figure 6.8. It is essentially identical to Figure 6.5 and is ruled out for the same reasons as were used to rule out the sequence $(1-4) \rightarrow, (1-4) \rightarrow$ in the proof of Lemma 6.1.

Proof of 6.3.7. The events b_2^{-1} and $(2-3) \rightarrow$ result in the creation of bridges in fibers of H' : first a bridge of name 35 with label B_i , and then a bridge of name 23 and label B_k . There are two cases: $i \neq k$ and $i = k$. We will discuss the case $i \neq k$ first, and then we will show that the argument which rules out the case $i \neq k$ also rules out the case $i = k$.

The local picture after the initial b_2^{-1} singularity can be deduced from Figure 6.5. After the singularity, there will be a leaf joining the pierce-points 3 and 6, and the leaves μ_{34} and μ_{12} , as in Figure 6.9. In order for the $(2-3) \rightarrow$ singularity to be possible, the “3” side of μ_{34} and the “2” side of μ_{12} must be accessible to opposite sides of the same digon, and this dictates that the six pierce-points are located as shown, relative to the evolving link projection. By Sublemma 6.1.1 there must be a non-empty block $P = b_1^p$ after the type $(2-3) \rightarrow$ singularity, and a leaf which we have labeled ζ which joins 1 to 6 and runs below the block P .

Our task now is to extend this picture to include the third key event, i.e. $(1-6) \rightarrow$. Referring to Figures 3.5 and 3.6, we recall that we will need to find joining arcs α in ∂T and β in fibers of H' , between




FIGURE 6.9

the “1” side and the “6” side of the leaf ζ . The arcs α and β must cobound a disc Δ in S^3 , and the singularity will be realized by an isotopy along Δ . Since the block P is maximal, the only way that α can exist is if P is followed (in the projection of L onto ∂T) by a crossing a_2^{-1} , as illustrated in Figure 6.9. The $(1-6)^{\rightarrow}$ singularity produces digon crossing arcs of name 16 in a digon which lies between the 3rd and 1st braid strands, Sublemma 6.1.1 says there must be a non-empty block $Q = b_3^q$ after the $(1-6)^{\rightarrow}$ singularity.

Passing to the standard word W , we see from Figure 6.9 that W contains the subword $\dots(a_1)^p(a_2)^{-1}(a_3)^q\dots$. By relation (2.3) that subword is equivalent to $a_1^{p+q-1}a_2^{-1}$. Thus W does not have shortest syllable length.

The identical argument applies if $i = k$. The only change is in the picture above the block P , which is now as in Figure 6.10. The block




FIGURE 6.10

P , the crossing we called a_2^{-1} and the block Q exist as before. Again W does not have shortest syllable length.

Proof of 6.3.8. If the sequence $(1 - 4)^{\rightarrow}, (2 - 3)^{\rightarrow}$ occurs, the local picture on ∂T begins with the configuration which was illustrated in Figure 6.2. We leave it to the reader to extend it to the full picture, using the following considerations: In order for the $(2 - 3)^{\rightarrow}$ event to be possible, the “2” side and the “3” side of the arc which joins 2 to 3 in Figure 6.4 must be on opposite sides of the same digon. The only way this can happen is if the block P shown there is followed by a crossing of type $(a_2)^{-1}$ in the projection of L onto ∂T . Then after the $(2 - 3)^{\rightarrow}$ event there must be a second non-empty block $(b_3)^q$ between braid strands 2 and 3. So W contains the subword $(a_1)^p(a_2)^{-1}(a_3)^q$. This word is equivalent to $(a_1)^{p+q}(a_2)^{-1}$. Therefore W does not have minimum syllable length and this case does not occur.

Proof of 6.3.9. Suppose that the sequence $(2 - 3)^{\rightarrow}$ followed by $(4 - 5)^{\rightarrow}$ occurs. After the two singularities there will be a bridge of name 23 and label B_i and another of name 45 and label B_k . The partial projections on ∂T in the two cases are shown in Figure 6.13 (see p. 98). Sublemma 6.1.1 implies that the blocks P and Q in both pictures are both non-empty. This shows immediately that the case $i = k$ is impossible because the block P is not maximal.

We pass to the case $i \neq k$. Recall that the leaves of the foliation

in the initial fiber, i.e. the arcs μ_{12} , μ_{34} , μ_{56} , cross \mathbf{L} once, without crossing any digons. In Figure 6.13, right picture, we have shown, in addition to the blocks P and Q , a (possibly empty) block $R = (a_3)^r$ located above μ_{56} and below P . To show that the block R must be non-empty, we notice that the three arcs $\mu_{2i-1, 2i}$ cut directly across \mathbf{L} shows that \mathbf{A}' loops directly under \mathbf{L} between the pierce-points $2i - 1$ and $2i$. But then, if $r = 0$, \mathbf{A}' would cut directly under all three strands and so be isotopic to \mathbf{A} , which is impossible. Thus $r \neq 0$.

But then, \mathbf{W} contains the subword $(a_1)^p(a_3)^r(a_2)^q$, where $p \neq 0$, $r \neq 0$, $q \neq 0$. This implies that $r \neq -1$, for if $r = -1$ we could use the relations $(a_1)^p(a_3)^{-1} = (a_3)^{-1}(a_2)^p$ to amalgamate syllables, contradicting the assumption that \mathbf{W} has shortest syllable length.

We now consider the initial backward sequence. It cannot contain an Ω -singularity of type $[1-2]$ because if it did then (by Corollary 5.8) one of the 3-balls would be completely foliated after this singularity, however we already have three non-empty clocks P , Q and R , one in each 3-ball. Thus the initial backward sequence satisfies all of the same restrictions as did the initial forward sequence. In particular, this means that it begins with a singularity of type $[2-3]^\rightarrow$ or b_i^{-1} . However, it cannot begin with an NSF singularity, because from Figure 6.13 we see that there are obstructions to the required joining arcs on $\partial\mathbf{T}$. It also cannot be an Ω -singularity of type $(2-3)^\rightarrow$, for if so the forward sequence would contain a singularity of type $(2-3)^\leftarrow$ followed by another of type $(2-3)^\rightarrow$ without any intervening block, and this is ruled out by Sublemma 6.1.1. The same argument rules out a singularity of type $(4-5)^\rightarrow$, because in the forward sequence we can reverse the order and do the $(4-5)^\rightarrow$ before the $(2-3)^\rightarrow$. It also cannot be an Ω -singularity of type $(1-6)^\rightarrow$ because there is no way for arcs of name 11 and 66 to be on opposite sides of the same digon. Thus there is no way to complete the projection to a cycle. Thus the sequence $(2-3)^\rightarrow$, $(4-5)^\rightarrow$ cannot occur.

Proof of 6.3.10. The sequence $(b_2)^{-1}$, $(2-3)^\rightarrow$ was considered earlier, in 6.3.7. It was shown that the case $i = k$ does not occur. We now claim that the case $(2-3)^\rightarrow$, $(b_2)^{-1}$ also does not occur if $i = k$. For, the projection of \mathbf{L} onto $\partial\mathbf{T}$ is unchanged. The region which is foliated is unchanged. One need only reverse the order of the two singularities. But then, the same reason as was used earlier to rule out $(b_2)^{-1}$, $(2-3)^\rightarrow$ shows that the case $(2-3)^\rightarrow$, $(b_2)^{-1}$ also does not occur if $i = k$.

For the case $(2-3)^{\rightarrow}$, $(b_2)^{-1}$ with $i \neq k$, we refer to Figure 6.10, focusing on the foliation shown there on ∂T after $(b_2)^{-1}$, $(2-3)^{\rightarrow}$ and before the subsequent $(1-6)^{\rightarrow}$. There is no obstruction on ∂T to reversing the order of the singularities and doing the $(2-3)^{\rightarrow}$ singularity first. There is also no obstruction in fibers of H' . Therefore the two cases are equivalent, after a change in fibration. \square

LEMMA 6.4. *If the singularities in the non-standard foliation do not include any singularities of type $[1-4]$, then W either has syllable length three, or else it has syllable length four and one of the braid generators occurs exactly once.*

Proof of Lemma 6.4. We investigate the initial forward sequence. We have already shown that we may assume that there is at most one singularity of type $[1-2]$. We may therefore assume that if one occurs, it is in the initial backward sequence. Therefore the initial forward sequence only involves Ω -singularities of type $[1-6]$ and NSF singularities, augmented by blocks. We study the possibilities for the first two events.

Our first assertion is that the first two events create bridges which are in distinct sectors of fibers of H' . This is so because all possibilities for two events in the same sector are ruled out by Lemma 6.2 and by Lemma 6.3, parts 6.3.6 and 6.3.2.

Our next assertion is that the first two events cannot both be Ω -singularities. For if they are, then up to symmetry we may assume they are $(2-3)^{\rightarrow}$, $(4-5)^{\rightarrow}$, however that possibility is ruled out by 6.3.9.

We next claim that the first two events cannot both be NSF-singularities. If they are, then up to symmetry we may assume they are $(b_1)^{-1}$, $(b_2)^{\delta}$, where $\delta = \pm 1$. What can happen next? The possibilities for the third key event are: another NSF singularity or an Ω -singularity of type $[1-6]^{\rightarrow}$. The only NSF singularities for which there are no obstructions to a joining arc in the fibers of H' are $(b_2)^{-\delta}$ and b_1 . However if $(b_2)^{-\delta}$ occurred then W' would not be reduced. If b_1 occurred W' would contain the subword $(b_1)^{-1}(b_2)^{\delta}b_1$, which could be replaced by the shorter subword $(b_3)^{\delta}$. So, the third event must be an Ω -singularity of type $[1-6]$.

The sequence $(b_1)^{-1}$, $(b_2)^{\pm 1}$, $(4-5)^{\rightarrow}$ is ruled out by 6.3.4 and 6.3.5. The sequence $(b_1)^{-1}$, b_2 , $(2-3)^{\rightarrow}$ is ruled out by 6.3.3. The sequence $(b_1)^{-1}$, $(b_2)^{-1}$, $(2-3)^{\rightarrow}$ is ruled out with the help of

6.3.10, for we can exchange the order of the $(b_2)^{-1}$ and the $(2-3)^{\rightarrow}$, which reduces things to a sequence in which the first two events are not both, NSF singularities. The same reason rules out $(b_1)^{-1}$, b_2 , $(1-6)^{\rightarrow}$ because b_2 , $(1-6)^{\rightarrow}$ is equivalent under symmetries to $(b_2)^{-1}$, $(2-3)^{\rightarrow}$. The sequence $(b_1)^{-1}$, $(b_2)^{-1}$, $(1-6)^{\rightarrow}$ is ruled out because if it occurs, then after the $(1-6)^{\rightarrow}$ there must be a block \mathbf{b}_3^p . But then, \mathbf{W}' will contain the subword $(b_1)^{-1}(b_2)^{-1}(b_3)^p$, which is equivalent to $(b_1)^{p-1}(b_2)^{-1}$, so \mathbf{W}' does not have shortest syllable length.

Thus we have shown at least one of the first two events must be an Ω -singularity, and that the two events occur in distinct sectors on fibers of \mathbf{H}' . Up to symmetry we may assume the Ω -singularity is $(2-3)^{\rightarrow}$. The distinct possibilities are then:

- $(2-3)^{\rightarrow}$, $(b_1)^{-1}$,
- $(b_2)^{-1}$, $(2-3)^{\rightarrow}$,
- $(2-3)^{\rightarrow}$, b_2 ,
- b_2 , $(2-3)^{\rightarrow}$.

We claim that the sequences $(b_2)^{-1}$ followed by $(2-3)^{\rightarrow}$ and $(2-3)^{\rightarrow}$ followed by $(b_2)^{-1}$ do not occur. By 6.3.10 the two cases are equivalent, so we may pass back and forth between them as convenient. What can the third event be? Assuming first that the order of the first two events is $(b_2)^{-1}$, $(2-3)^{\rightarrow}$, we see (by 6.2.2) that the third event cannot be an Ω -singularity of type $(2-3)^{\rightarrow}$. By 6.3.7 it cannot be $(1-6)^{\rightarrow}$ and by 6.3.9 it cannot be $(4-5)^{\rightarrow}$. To see that it cannot be b_2 , reverse the order of the first two events to $(2-3)^{\rightarrow}$, $(b_2)^{-1}$ and use the fact that \mathbf{W}' is freely reduced. There are obstructions in fibers of \mathbf{H}' to the third event being $b_1^{\pm 1}$ or $(b_2)^{-1}$. The only remaining possibility is b_3^q . Now use the fact that by Sublemma 6.1.1 the $(2-3)^{\rightarrow}$ event is necessarily followed by a nonempty block \mathbf{b}_1^p , so the augmented cycle is $(2-3)^{\rightarrow}$, $(\mathbf{b}_1)^p$, $(b_2)^{-1}$. But then \mathbf{W}' contains a subword of the word $\dots(\mathbf{b}_1)^p(b_2)^{-1}(b_3)^q\dots$, which is equivalent to the subword $\dots(b_2)^{-1}(b_3)^{p+q}\dots$. Since the latter has fewer syllables, this can't occur.

We are reduced to exactly two cases, up to symmetry, i.e. $(2-3)^{\rightarrow}$, b_2 and b_2 , $(2-3)^{\rightarrow}$. Each of these results in the creation of two bridges. Let the labels on the bridges be \mathbf{B}_i and \mathbf{B}_k . There are two cases to consider: $i = k$ and $i \neq k$. When we take this fact into account the two remaining cases expand to four cases, namely:

Case 1. $(2-3)^{\rightarrow}$, b_2 with $i = k$,




FIGURE 6.11

Case 2. $b_2, (2-3)^{\rightarrow}$ with $i = k$,

Case 3. $(2-3)^{\rightarrow}, b_2$ with $i \neq k$,

Case 4. $b_2, (2-3)^{\rightarrow}$ with $i \neq k$.

To handle Case 1 we reason as follows. Since the first singularity is type $(2-3)^{\rightarrow}$ we know that the initial local picture is as in Figure 6.4. In Figure 6.11 we have illustrated the same situation, but omitted the digons and the leaves of the foliation in order to emphasize new aspects of the geometry. In particular, we have shown how the non-standard axis must look in the part of the picture which is of interest. The non-standard axis is inside T , however it pierces ∂T at the point-pairs 1-2, and 3-4. Each of these points actually represents a single point on F but two points on ∂T . The non-standard axis A' is actually disconnected when we cut open F in our torus-3-ball decomposition. We have shown it instead in Figure 6.11 as looping under the braid (as it must) between 1 and 2 and also between 3 and 4. In between the points 3 and 4 it is inside T . But then, we may always push the block P over the axis, without changing the fact that A' is an axis for a closed braid representation of L or changing the complexity, to change things to the case where the first singularity is b_2 instead of $(2-3)^{\rightarrow}$. Thus Case 1 is reduced to Cases 2 and 4.

We next consider Case 2. Figure 6.12 shows the partial foliation on ∂T after the sequence $b_2, (2-3)^{\rightarrow}$, with the two bridges having the same label. The portion of the braid which is labeled $P(a_1, a_2)$ might or might not occur. By Sublemma 6.1.1 we know that the $(2-3)^{\rightarrow}$ singularity must be followed by a nonempty block which moves one of the 23 bridge arcs. A picture on fibers of H' shows that the leaf which joins the pierce points 1 and 6 must be the arc which sweeps through the block, because it crosses L twice and contains a 23 bridge




FIGURE 6.12

arc. However, we see from Figure 6.12 that no such block can occur. Thus case 2 cannot occur.

We now claim that in Case 3 there is a non-standard axis, the word W satisfies the conditions of Lemma 6.4. (Remark: we could reduce Case 3 to Case 4 in the same way as we reduced Case 1 to Cases 2 and 4, but it will be easier if we do not.) The first step in the proof is the construction of the projection onto ∂T of that part of the sequence which corresponds to the initial forward sequence. The procedure for so-doing should be familiar by now. See the left picture in Figure 6.14. We have eliminated the digons to avoid cluttering up the picture. The portion of the figure which we are looking at now includes the pierce-points $1, \dots, 6$, the non-empty block $P = (a_1)^p$ which necessarily follows the initial $(2-3) \rightarrow$ event and the subsequent block $Q = (a_2)^q$, which might be empty. The location of the pierce-points 5 and 6 is determined by two considerations: first, the necessity for a joining arc between μ_{14} and μ_{56} in F^- , and second the fact that the bridge created by the b_2 singularity is required to have a different label from the bridge created by the initial singularity of type $(2-3) \rightarrow$. We




FIGURE 6.13




FIGURE 6.14

have also sketched in part of the axis A' in Figure 6.14. We know it looks as sketched because the cyclic order of the six points on A' must be $1, 2, \dots, 6$ and because the axis A' can't do anything more complicated than to simply loop under the i th braid strand between the pierce-points $2i-1$ and $2i$, because we began with an initial fiber.

Now notice that if $q = 0$ the non-standard axis will loop directly around the link, indicating that the braid has a trivial loop and the braid index is not minimal. Thus $q \neq 0$. Looking at the right picture in Figure 6.14, and mentally translating the axis A' from the left picture to the right, we see that if $q = 1$ then the non-standard axis is isotopic to the standard axis. Thus $q \neq 1$ either. Therefore, after


the initial forward sequence we know there is a subword $(a_1)^p(a_2)^q$ of \mathbf{W} with $p \neq 0$, $q \neq 0, 1$.

We now study the possibilities for the initial backward sequence (IBS), using this initial forward sequence (IFS). The IBS begins with the same initial fiber, with the pierce-points arranged as in the left picture in Figure 6.14. What can the first *backward* event be? The first thing to notice is that there are leaves μ_{12} , μ_{34} and μ_{56} in the foliation which join the point pairs **1-2**, **3-4** and **5-6**. Since we know that after the IFS there will be leaves in the foliation which join **2** to **3**, **1** to **6** and **4** to **5**, it follows that the IBS contains singularities which change the connectivities of the leaves, before the two sequences overlap.

Suppose that the IBS begins with an Ω -singularity. If it is type $(2-3)^\rightarrow$, then it would necessarily be followed in the backward direction by a non-empty block P' . However, in that case the two blocks P and P' could be amalgamated, reducing the complexity, so that can't happen. If it begins with $(4-5)^\rightarrow$ then the **5** side of the arc μ_{56} and the **4** side of the arc μ_{34} must be on opposite sides of the same digon. The only way this can occur is if $q = 1$, as indicated in the right picture in Figure 6.14. However, as observed earlier, that is impossible. Similarly, if the IBS begins with $(1-6)^\rightarrow$ the **6** side of μ_{56} and the **1** side of μ_{12} must be on opposite sides of the same digon, and again this means $q = 1$. So, the IBS cannot begin with an Ω -singularity.

Now suppose that the IBS begins with an NSF-singularity. There are six types to consider (types b_1^ε , $i = 1, 2, 3$ and $\varepsilon = \pm 1$). If it is type b_1^{-1} then, referring to Figure 6.6(c) we know that we need to have a joining arc in the backward direction, in \mathbf{F}^+ , between the **2** side of μ_{12} and the **4** side of μ_{34} . Call that a "joining arc of type **24**". However, from the partial projection of Figure 6.14 that is clearly impossible. Since joining arcs of type **24**, **46** and **61** in \mathbf{F}^+ and **13**, **35**, and **51** in \mathbf{F}^- are all impossible in the situation of Figure 6.14, we conclude that the IBS can't begin with an NSF-singularity either.

Thus, the only possibility is that the first event in the IBS is a type $(1-2)^\rightarrow$ singularity, which implies that \mathbf{W} now has the subword $(a_3)^{-1}(a_1)^p(a_2)^q$ of \mathbf{W} with $p \neq 0$, $q \neq 0$, as in Figure 6.15 (next page). This results in the foliation of the entire 3-ball \mathbf{B}_3 . After the $(1-2)^\rightarrow$ singularity it becomes possible to do a $(1-6)^\rightarrow$ singularity and so to sweep out block P . But then, we will have foliated all of \mathbf{B}_2 and well as \mathbf{B}_3 . So there can only be one more event, namely a



completed cycle, with non-standard axis A' (opposite sides of rectangle to be identified)
 $(2-3)^{\rightarrow}, b_2$

FIGURE 6.15

$(2 - 3)^{\rightarrow}$, followed by a block $(a_1)^r$, which results in the foliation of all of B_1 too. Thus the cyclic word W is $(a_1)^r(a_3)^{-1}(a_1)^p(a_2)^q$, as claimed. Straightening out A' in Figure 6.15, and referring back to Figure 1.3 at the beginning of this paper, we see that $W' = (a_1)^p(a_3)^{-1}(a_1)^r(a_2)^q$, as claimed.

Case 4 remains, and we now show that it reduces to Case 3. To see this, we first construct the partial projection associated to the initial forward sequence. See the right picture in Figure 6.14, where again we have omitted digons. The non-empty block $R = (a_1)^r$ is the one which is associated to the $(2 - 3)^{\rightarrow}$ singularity. In principle, the block P could be empty. However, now notice that after we sketch in part of the axis we can see that block P must be non-empty because if it were empty the braid would be reducible, since the axis loops once about the braid between the points 4 and 6. With this partial projection in hand, go back to Figure 6.15, and notice that we can push the block P over the axis A' over the subarc 1234 of A' . This changes part of the projection for the cycle in Figure 6.15 to the partial projection of Figure 6.16. Since there was only one way to complete the sequence $(2-3)^{\rightarrow}, b_2$ to a cycle, the two cases must be identical. This completes the proof of Lemma 6.4. \square

LEMMA 6.5. *If the non-standard foliation of one of the digons contains a singularity of type [1 - 4] then every singularity is type*




FIGURE 6.16

[1 – 4]. In particular, there are no NSF-singularities and there are no Ω -singularities of type [1 – 6] or [1 – 2].

Proof of Lemma 6.5. We are given a non-standard axis and a fibration \mathbf{H}' which includes a singularity of type [1 – 4]. By symmetry, we may suppose that the singularity is type $(1 - 4)^{\rightarrow}$. Irrespective of whether we begin with an initial fiber or not, there will be a 14 bridge in fibers of \mathbf{H}' immediately after the singularity. Therefore the possibilities for the next event are quite limited: up to the usual symmetries, they are $(1 - 4)^{\rightarrow}$, $(1 - 4)^{\leftarrow}$, $(2 - 3)^{\rightarrow}$, $(4 - 5)^{\rightarrow}$, $(b_1)^{-1}$ and $(b_2)^{-1}$. In view of 6.2.1, 6.3.8, 6.3.1 (which are easily seen to hold even if we drop the assumption that we started with an initial fiber) we are reduced to $(1 - 4)^{\leftarrow}$, $(b_1)^{-1}$ and $(b_2)^{-1}$. We want to prove that $(1 - 4)^{\rightarrow}$ followed by $(b_1)^{-1}$ and $(b_2)^{-1}$ cannot occur.

Consider $(1 - 4)^{\rightarrow}$, $(b_2)^{-1}$ first. It will be convenient to replace this sequence by its symmetric equivalent $(2 - 5)^{\rightarrow}$, b_2 . Now, let ϕ_1 and ϕ_2 be polar angles just after the $(2 - 5)^{\rightarrow}$ and b_2 singularities. Then $\mathbf{H}'_{\phi_1} \cap \mathbf{F}$ is a union of three arcs, which are divided by the three points of $\mathbf{H}'_{\phi_1} \cap \mathbf{L}$ into 6 arcs μ_1, \dots, μ_6 , where μ_i runs from the pierce-point \mathbf{i} to one of the three points of $\mathbf{H}'_{\phi_1} \cap \mathbf{L}$. The 25 bridge further divides μ_5 into 3 subarcs. Let ξ be the one which is closest to $\mathbf{5}$ and let ζ be the union of the other two. Then the joining arc for the b_2 meets μ_5 along ξ , whereas the 25 bridge arcs meet μ_5 along ζ . The pictures in Figure 3.12 of §3 then show that after the




FIGURE 6.17

b_2 singularity the names of the 25 bridge arcs will have changed to 23. Thus $\mathbf{H}'_{\phi_2} \cap \mathbf{F} \cup \Omega$ includes bridges of name 24 and 23. There is another sequence of singularities which also leads to bridges of name 24 and 23, i.e. $b_2, (2-3)^{\rightarrow}$. We investigated it in the proof of Lemma 6.4. It is not difficult to see that region which is foliated by the pair $b_2, (2-3)^{\rightarrow}$ can also be foliated by the pair $(2-5)^{\rightarrow}, b_2$. That is, the sequences $(2-5)^{\rightarrow}, b_2$ and $b_2, (2-3)^{\rightarrow}$ are equivalent after a change in fibration. Since the latter does not use any singularities of type $[1-4]$, we can make this change in fibration and so eliminate the pair $(2-5)^{\rightarrow}, b_2$ (and its symmetric equivalent $(1-4)^{\rightarrow}, (b_2)^{-1}$) from consideration.

The case $(1-4)^{\rightarrow}, (b_1)^{-1}$ remains. Let ϕ_1, ϕ_2 and ϕ_3 be ϕ -values after the $(1-4)^{\rightarrow}$ singularity, the subsequent block, and at the moment of the $(b_1)^{-1}$ singularity. The local picture on $\partial\mathbf{T}$ after the $(1-4)^{\rightarrow}$ singularity and the block was depicted earlier in Figure 6.2. The arc α'_1 is subdivided into 5 segments, which we will refer to as i, ii, iii, iv and v, by its intersections with \mathbf{L} and with the digon-attaching curves, where segment #i is closest to the pierce-point 2. In order for the $(b_1)^{-1}$ singularity to be possible there must be a joining arc in \mathbf{F}^+ between segments #i and #iii. Pass to Figure 6.17, which shows the local picture on $\partial\mathbf{T}$ near the second singularity, omitting unnecessary

details from Figure 6.2. We have cut \mathbf{T} open between the first and second strands. We see that \mathbf{W} contains the subword $b_1^p(b_2)^{-1}b_3$, which is equivalent to $(b_2)^{-1}(b_3)^{p+1}$. We conclude that \mathbf{W} does not have shortest syllable length and is therefore not admissible.

Thus the only possibility is that each singularity of type $[1 - 4]$ is followed by another of type $[1 - 4]$. This completes the proof of Lemma 6.5, and so also of the Classification Theorem. \square

7. The algorithm and applications. Our task in this section is to translate the Classification Theorem of §6 into an algorithm to decide whether two links which are defined by closed 3-braids have the same oriented link type in oriented 3-space. This problem is reduced, via the Classification Theorem, to two other problems:

Problem A. Find unique representatives for conjugacy classes in B_3 .

Problem B. Identify the unique representatives of the conjugacy classes which correspond to links which are represented by more than one conjugacy class of 3-braids.

There are several known solutions to Problem A and our algorithm will use the solution which was given by Schreier in 1924 [Sc]. As noted earlier, a new and different solution, due to P. J. Xu, is given in [X]. In the latter conjugacy classes are represented by shortest words in the elementary braids a_1 , a_2 , a_3 which were used in the body of this manuscript; the unique representatives of conjugacy classes are shortest words in these generators, and therefore determine a surface of maximum Euler characteristic with the link as its boundary. Our reason for choosing Schreier's solution over Xu's at this time are that it is a little bit easier to use than Xu's. Moreover, Xu has not yet succeeded in finding normal forms which are expressed in terms of shortest syllables as well as shortest words, so we do not have the ability at this time to translate her work into an optimum algorithm which orders braids by the complexity function which was the basis of the work in this paper.

7.1. *Schreier's solution to the conjugacy problem [Sc].*

Step 0. We assume that we are given a 3-braid $\mathbf{W} = W(\mathbf{L})$ which is defined as a product of powers of the elementary braids σ_1 and σ_2 and their inverses, where $a_1 = \sigma_1$, $a_2 = \sigma_2$, $a_3 = \sigma_2\sigma_1\sigma_2^{-1}$. Our goal is to find a unique representative for the conjugacy class of \mathbf{W} in B_3 .

Step 1. Introduce new generators x and y , where:

$$(7.1) \quad x = (\sigma_1 \sigma_2 \sigma_1)^{-1} \quad \text{and} \quad y = \sigma_1 \sigma_2,$$

so that:

$$(7.1)^* \quad \begin{aligned} \sigma_1 &= y^2 x, \\ \sigma_2 &= x y^2, \\ \sigma_1^{-1} &= x y, \\ \sigma_2^{-1} &= y x. \end{aligned}$$

Use (7.1)* to rewrite \mathbf{W} as a cyclically reduced word \mathbf{W}_0 in powers of x and y .

Step 2. Notice that if we use x and y as generators then B_3 is a free product with amalgamations of the cyclic groups generated by x and y , the amalgamation being defined by $x^{-2} = y^3$. Set

$$\mathbf{C} = x^{-2} = y^3 = (\sigma_1 \sigma_2 \sigma_1)^2 = (\sigma_1 \sigma_2)^3.$$

Then \mathbf{C} generates the center of B_3 . We may therefore replace the cyclic word \mathbf{W}_0 by:

$$(7.2) \quad \mathbf{W}_1 = \mathbf{C}^k x y^{a_1} x y^{a_2} x y^{a_3} \cdots x y^{a_r},$$

where k is an arbitrary integer and where each $a_i = 1$ or 2 .

Step 3. Group together the terms in (7.2) in the form:

$$(7.3) \quad \mathbf{W}_2 = \mathbf{C}^k \mathbf{T},$$

where \mathbf{T} is one of the cyclic words:

$$(7.4) \quad \mathbf{T} = (xy)^{p_1} (xy^2)^{q_1} (xy)^{p_2} (xy^2)^{q_2} \cdots (xy)^{p_r} (xy^2)^{q_r},$$

$$r \geq 1, p_1, q_1, \dots, q_r \geq 1,$$

$$\begin{aligned} &(xy)^p, \quad p \geq 1, \\ &(xy^2)^q, \quad q \geq 1, \\ &y, \quad \text{or} \\ &y^2, \quad \text{or} \\ &x, \quad \text{or} \\ &1. \end{aligned}$$

The word $\mathbf{C}^k \mathbf{T}$, where $k \in \mathbb{Z}$ is Schreier's unique representative of the conjugacy class of \mathbf{W} . Using the elementary braid generators σ_1

and σ_2 , the list of the distinct conjugacy classes is:

$$(7.4)^* \quad \begin{aligned} & \mathbf{C}^k \sigma_1^{-p_1} \sigma_2^{q_1} \cdots \sigma_1^{-p_r} \sigma_2^{q_r}, & (k, r, p_i, q_i \in \mathbf{Z}; r, p_i, q_i \geq 1), \\ & \mathbf{C}^k \sigma_1^p, & (k, p \in \mathbf{Z}), \\ & \mathbf{C}^k \sigma_1 \sigma_2, & (k \in \mathbf{Z}), \\ & \mathbf{C}^k \sigma_1 \sigma_2 \sigma_1, & (k \in \mathbf{Z}), \\ & \mathbf{C}^k \sigma_1 \sigma_2 \sigma_1 \sigma_2, & (k \in \mathbf{Z}). \end{aligned}$$

We shall refer to the first case in (7.4)* as the *generic case*.

7.2. *The algorithm.* Let \mathcal{L} be a link which is represented by a closed 3-braid \mathbf{L} . Let $\mathbf{W} = W(\mathbf{L})$ be a cyclic word in the standard elementary braids σ_1, σ_2 which defines \mathbf{L} . Let $\mathbf{C}^k \mathbf{T}$ be Schreier's unique representative of the conjugacy class of \mathbf{W} . Then the conjugacy class of \mathbf{W} is the only conjugacy class in B_3 which represents \mathcal{L} , with the following exceptions:

Links of braid index 1. The unknot, with three conjugacy classes of 3-braid representatives:

$$\sigma_1 \sigma_2, \sigma_1^{-1} \sigma_2 \quad \text{and} \quad \mathbf{C}^{-1}(\sigma_1 \sigma_2)^2.$$

Links of braid index 2. (a) The 2-component unlink, with two conjugacy classes of 3-braid representatives:

$$\sigma_1 \quad \text{and} \quad \sigma_1^{-1}.$$

(b) A type $(2, p)$ torus link, $p \neq 0, \pm 1$, with two conjugacy classes of 3-braid representatives:

$$\begin{aligned} & \sigma_1^{-1} \sigma_2^p \quad \text{and} \quad \mathbf{C} \sigma_1^{p-2} & \text{if } p \geq 4, \\ & & (\sigma_1 \sigma_2)^2 & \text{if } p = 3, \\ & & \sigma_1 \sigma_2 \sigma_1 & \text{if } p = 2, \\ & \sigma_1^p \sigma_2 \quad \text{and} \quad \mathbf{C}^{-1} \sigma_1^{p+4} \sigma_2 & \text{if } p \leq -5, \\ & & \mathbf{C}^{-1} \sigma_1 & \text{if } p = -4, \\ & & \mathbf{C}^{-1} \sigma_1 \sigma_2 & \text{if } p = -3, \\ & & \mathbf{C}^{-1} \sigma_1 \sigma_2 \sigma_1 & \text{if } p = -2. \end{aligned}$$

Links of braid index 3 which "admit flypes". \mathcal{L} or its mirror image has two representatives, namely:

$$\mathbf{C}^k \sigma_1^{-1} \sigma_2^u \sigma_1^{-v} \sigma_2^w \quad \text{and} \quad \mathbf{C}^k \sigma_1^{-1} \sigma_2^w \sigma_1^{-v} \sigma_2^u,$$

where $k = 0$ or 1 , $v \geq 2$ and u and w are distinct positive integers,

or

$$\mathbf{C}^s \sigma_1^{-1} \sigma_2^u \sigma_1^{-1} \sigma_2^v \sigma_1^{-1} \sigma_2^w \quad \text{and} \quad \mathbf{C}^s \sigma_1^{-1} \sigma_2^u \sigma_1^{-1} \sigma_2^w \sigma_1^{-1} \sigma_2^v,$$

where $s = 1$ or 2 and u, v, w are pairwise distinct positive integers.

Proof. Our starting point is the classification theorem (version 2).

Braid index 1. The 3 cases are Schreier's representatives of the conjugacy classes of $\sigma_1 \sigma_2$, $\sigma_1^{-1} \sigma_2$ and $\sigma_1^{-1} \sigma_2^{-1}$ respectively.

Braid index 2. The cases which are listed are Schreier's representatives of the conjugacy classes of $\sigma_1^{-1} \sigma_2^p$ and $\sigma_1 \sigma_2^p$.

Braid index 3. After making the substitutions $a_1 = \sigma_1$, $a_2 = \sigma_2$, $a_3 = \sigma_2 \sigma_1 \sigma_2^{-1}$ one obtains a uniform description of \mathbf{W} and \mathbf{W}^{\leftarrow} as $(\sigma_1)^\delta (\sigma_2)^a (\sigma_1)^b (\sigma_2)^c$ and $(\sigma_1)^\delta (\sigma_2)^c (\sigma_1)^b (\sigma_2)^a$, where $\delta = \pm 1$. The integers δ, a, b, c could each be > 0 or < 0 , so our task is to put these into Schreier's normal form. We only need to consider the cases where \mathbf{W} and \mathbf{W}^{\leftarrow} are in distinct conjugacy classes. The first thing to notice is that we only need to look at cases where the Schreier representative is generic with $r \geq 2$, because in all other cases the Schreier representative and its reverse coincide. Up to replacing the cyclic word \mathbf{W} with \mathbf{W}^{-1} , \mathbf{W}^{\leftarrow} or $(\mathbf{W}^{\leftarrow})^{-1}$, we may therefore assume that the signs of δ, a, b, c are $-+++, -+--, -+-+, +++++, +++-, ++-+$. Let $\alpha, \beta, \gamma = |a|, |b|, |c|$ respectively.

If the signs are $-+++$, then $\mathbf{W} = (xy)(xy^2)^\alpha (y^2x)^\beta (xy^2)^\gamma$, where α, β, γ are all positive. Computing, we find that \mathbf{W} has Schreier representative $\mathbf{C}(xy)(xy^2)^{\alpha-1}(xy)(xy^2)^{\beta-2}(xy)(xy^2)^{\gamma-1}$. In view of the fact that we are only interested in the generic case we may assume that the integers $\alpha-1, \beta-2, \gamma-1$ are all positive, so the unique representative of \mathbf{W} is $\mathbf{C}\sigma_1^{-1}\sigma_2^u\sigma_1^{-1}\sigma_2^v\sigma_2^{-1}\sigma_2^w$, where u, v and w are positive, whereas \mathbf{W}^{\leftarrow} has Schreier representative $\mathbf{C}\sigma_1^{-1}\sigma_2^u\sigma_1^{-1}\sigma_2^w\sigma_1^{-1}\sigma_2^v$. The two words are cyclically distinct if and only if u, v, w are pairwise distinct.

The other five cases are similar. We compute the Schreier representative of \mathbf{W} to be:

$$\begin{array}{ll} (xy)(xy^2)^u(xy)^v(xy^2)^w & \text{if the signs are } -++- \\ (xy)(xy^2)^u(xy)^v(xy^2)^w & \text{if the signs are } -+-+ \\ \mathbf{C}^2(xy)(xy^2)^u(xy)(xy^2)^v(xy)(xy^2)^w & \text{if the signs are } +++++ \\ \mathbf{C}(xy)(xy^2)^u(xy)^v(xy^2)^w & \text{if the signs are } +++- \\ \mathbf{C}(xy)(xy^2)^u(xy)^v(xy^2)^w & \text{if the signs are } ++-+. \end{array}$$

The assertion follows. \square

7.3. *Applications.* In what follows \mathcal{L} denotes a link of braid index 3, \mathbf{L} is a 3-braid representative and $\mathbf{W} = W(\mathbf{L})$ has Schreier normal form.

COROLLARY 7.1. *A. The link type of \mathbf{L} is invertible if and only if the conjugacy class of \mathbf{W} in (7.4)* satisfies one of the following:*

- (a) \mathbf{W} has any of the 3-braid non-generic normal forms, or
- (b) \mathbf{W} is generic and the arrays $(p_1, q_1, p_2, \dots, p_r, q_r)$ and $(p_r, q_r, p_{r-1}, q_{r-1}, \dots, p_1, q_1)$ differ by an even cyclic permutation, or
- (c) \mathbf{W} is one of the special 3-braids which “admits a flype”.

B. The link of type \mathbf{L} is amphicheiral if and only if the conjugacy class of \mathbf{W} in (7.4) satisfies one of the following:*

- (a) $\mathbf{W} = 1$, or
- (b) \mathbf{W} is generic, the arrays $(p_1, q_1, p_2, \dots, p_r, q_r)$ and $(q_1, p_2, q_2, \dots, p_r, q_r, p_1)$ differ by an even cyclic permutation, and $k = 0$.

Proof of Corollary 7.1. A. Reversal of orientation on \mathbf{L} (but not 3-space) is achieved by reversing orientation on \mathbf{W} , i.e. reading \mathbf{W} backwards. Let \mathbf{W}^{\leftarrow} be the word obtained from \mathbf{W} by this change. In the non-generic cases $\mathbf{W} = \mathbf{W}^{\leftarrow}$ because \mathbf{C} is in the center of B_3 , and also conjugation by the Garside braid $\sigma_1\sigma_2\sigma_1$ interchanges σ_1 and σ_2 . In the generic case we use the same facts to arrive at possibility (b). As for (c), the special braids which admit flypes are invertible for a different reason: their two distinct conjugacy classes are represented by words \mathbf{W} and \mathbf{W}^{\leftarrow} .

B. Reversal of the orientation of 3-space (but not of \mathbf{L}) corresponds to changing the sign of each crossing in the knot diagram. Let \mathbf{W}^* be the word obtained from \mathbf{W} by this change. In the non-generic case the only way we can have $\mathbf{W} = \mathbf{W}^*$ is if $\mathbf{W} = 1$. This proves (a). In the generic case, if $\mathbf{W} = \mathbf{C}^k \sigma_1^{-p_1} \sigma_2^{q_1} \sigma_1^{-p_2} \sigma_2^{q_2} \dots \sigma_1^{-p_r} \sigma_2^{q_r}$ then $\mathbf{W}^* = \mathbf{C}^{-k} \sigma_1^{p_1} \sigma_2^{-q_1} \sigma_1^{p_2} \sigma_2^{-q_2} \dots \sigma_1^{p_r} \sigma_2^{-q_r}$. Conjugating \mathbf{W}^* by the Garside braid we see that its normal form is $\mathbf{C}^{-k} \sigma_1^{-q_1} \sigma_2^{p_2} \sigma_1^{-q_2} \sigma_2^{p_3} \dots \sigma_1^{-q_r} \sigma_2^{p_1}$. Assertion (b) follows. \square

REMARK. Since a typical word in Schreier’s normal form does not have the symmetries of Corollary 7.1 we conclude that most links of braid index 3 are neither invertible nor amphicheiral. On the other hand, since the standard knot tables were organized by crossing number, the examples which formed the core body of data for early workers in knot theory had many symmetries.

COROLLARY 7.2. *\mathcal{L} is a composite knot or link if and only if the Schreier representative of the conjugacy class of \mathbf{W} or \mathbf{W}^{-1} has one of the following normal forms*

$$\begin{aligned} \sigma_1^{-u} \sigma_2^v, & \quad \text{where } u \geq v \geq 2. \\ C^{-1} \sigma_1^{-u} \sigma_2 \sigma_1^{-v} \sigma_2, & \quad \text{where } u \geq v \geq 0. \end{aligned}$$

Proof. By [Mo, 1] or [B-M, IV] a composite link which is a closed 3-braid is conjugate to $(\sigma_1)^a (\sigma_2)^b$, where $|a| \geq 2$, $|b| \geq 2$. The first case occurs when a and b have opposite signs and the second when they have the same sign. \square

Our next corollary is an analogue of the Tait conjecture which holds for links of braid index 3.

COROLLARY 7.3. *Let \mathcal{L} and \mathcal{L}' be links of braid index 3, and let \mathbf{D} and \mathbf{D}' be diagrams which represent \mathcal{L} and \mathcal{L}' and which have 3 Seifert circles. Then $\mathcal{L} = \mathcal{L}'$ only if the algebraic crossing numbers of \mathbf{D} and \mathbf{D}' agree.*

Proof. By a theorem of Yamada [Y] every diagram with 3 Seifert circles may be changed to a 3-braid diagram without changing the algebraic crossing number. Since algebraic crossing number is a conjugacy class invariant, and since the two conjugacy classes in the exceptional case (iii) of §7.2 have the same exponent sum, the assertion follows. \square

REMARK. A different proof of Corollary 7.3 is given in [J] for $n = 3$ and 4. We hope to be able to prove in a later paper in this series that Corollary 7.3 is true for every braid index.

COROLLARY 7.4. *There is an algorithm to compute the genus of a link of braid index 3, if one begins with a closed 3-braid representative.*

Proof. Let \mathbf{F} be a Bennequin surface, with $\partial\mathbf{F} = \mathbf{L}$ a 3-braid representative of the given link \mathcal{L} . Let μ be the number of components in \mathbf{L} and let r be the number of bands in the disc-band decomposition of \mathbf{F} . Then the genus g of \mathbf{F} , which is also the genus of \mathcal{L} , is determined by r and μ and the number of discs in the disc-band decomposition, which is 3. The number r is also the letter length of any shortest word \mathbf{W} in the generators a_1, a_2, a_3 of B_3 . Using

results in [X], this letter length may be computed. This yields an algorithm to compute the genus. We refer the reader to [X] for the algorithm. \square

The Classification Theorem gives an affirmative answer to an interesting question of Murasugi (see the remark at the bottom of page 75 of [Mu]):

COROLLARY 7.5. *Let \mathcal{L} be represented by the conjugacy class of $\mathbf{C}^k\mathbf{T}$ and let \mathcal{K} be represented by the conjugacy class of \mathbf{T} . Then the integer k and the link type of \mathcal{K} are invariants (even more, they are complete invariants) of the link type of \mathcal{L} , except in the special case where \mathcal{L} is one of the links covered by case (iii) of the algorithm of §7.2. In the latter case \mathcal{L} is invertible, but \mathcal{K} need not be, so one might have two distinct link types $\mathcal{K}^{\rightarrow}$ and \mathcal{K}^{\leftarrow} which correspond to the single link type $\mathcal{L}^{\rightarrow}$.*

REMARK. The assumption that \mathcal{L} has braid index 3 is essential here, as may be seen by inspecting the exceptional cases (i) and (ii) in §7.2 above.

COROLLARY 7.6. *The stabilization index of a link of braid index 3 is one.*

Proof. By the Classification Theorem, a link of braid index 3 is either represented by a unique conjugacy class (in which case the stabilization index is zero) or else it is represented by two conjugacy classes which are related to one-another by flypes. Therefore it is only necessary to show that two classes which are related by flypes become conjugate after a single stabilization. For a proof, see page 100, equations (2–27) of [Bi, 1]. \square

Our next application concerns the Jones polynomial [J] and its various generalizations, i.e. link type invariants which are constructed from traces on finite-dimensional matrix representations of the braid groups. A central question about such invariants is whether a finite collection of them can be a complete type invariant. For links which have braid index ≤ 3 . A complete answer is beyond the scope of this paper, however we have a small contribution.

In the manuscript [Bi, 2] the first author initiated a study of the question of when distinct links of braid index 3 have the same 2-variable Jones polynomial. She encountered a fundamental difficulty

in that study, as she lacked a method for distinguishing link types, other than ad hoc techniques which applied to one example at a time. Now that a method is available, it becomes possible to extend the results of [Bi, 2]:

COROLLARY 7.7. *For each integer N there exist infinitely many sets, each of which contains exactly N links of braid index 3 which share the same 2-variable Jones polynomial.*

Proof. Let $\zeta: B_3 \rightarrow \mathrm{SL}(2, Z[t, t^{-1}])$ be the reduced Burau representation of B_3 . Choose $\alpha, \beta \in B_3$. Then the associated closed braids L_α and L_β will have distinct link types if the conjugacy classes of α and β are distinct and are not one of the exceptional cases listed in the Classification Theorem. By [Bi, 2] the links L_α and L_β will have the same 2-variable Jones polynomial if and only if

- (1) α and β have the same exponent sum (as words in σ_1 and σ_2), and
- (2) $\zeta(\alpha)$ and $\zeta(\beta)$ have the same trace.

Let $\rho: \zeta(B_3) \rightarrow \mathrm{SL}(2, \mathbb{Z})$ be the mapping defined by setting $t = -1$. Then ρ is surjective (see [Bi, 2]) and the kernel of ρ is the infinite cyclic group generated by $\zeta((\sigma_1\sigma_2\sigma_1)^4)$. Thus we may replace (2) by the condition:

- (3) $\rho\zeta(\alpha)$ and $\rho\zeta(\beta)$ have the same trace.

Choose any integer $N > 1$. In example 8.2 of [H] Horowitz has constructed just the sort of examples we seek.

8. Hints of additional pathology. It was shown in the introduction to this paper that the exceptional case (iii) of the Classification Theorem generalizes to produce examples of links of braid index n , for every n , which have at least two conjugacy classes of n -braid representatives, related by generalized flypes (see Figure 1.3). While flypes are “natural” for $n = 3$ in the sense that the links in question are invertible links and the two classes represent the same link with two orientations, this need not be the case for $n > 4$ because of the extra twists introduced into the weighted strands by generalized flypes.

We now assert that there is additional pathology, unrelated to flypes, which appears in a rudimentary form in B_3 . A composite link $\mathcal{K}_{p,q}$ of braid index 3 is a connected sum of a type $(2, p)$ and a type $(2, q)$ torus link. By [Mo, 1] such links admit a *unique* conjugacy class of 3-braid representatives. This means that if $\mathbf{K}_{p,q}$ is a representative




FIGURE 8.1

of $\mathcal{H}_{p,q}$ and if A_r, A_s are disjoint simple closed curves in $S^3 - \mathbf{K}_{p,q}$ which are axes for 3-braid structures on $\mathcal{H}_{p,q}$, then the links $\mathbf{K}_{3,5} \cup A_r$ and $\mathbf{K}_{3,5} \cup A_s$ are equivalent, i.e. there is a homeomorphism $h_{r,s}: (S^3, \mathbf{K}_{p,q}, A_r) \rightarrow (S^3, \mathbf{K}_{p,q}, A_s)$.

Figure 8.1 shows a representative $\mathbf{K}_{3,5}$ of $\mathcal{H}_{3,5}$. Two simple closed curves, A_0 and A_1 , are illustrated in the complement of $\mathbf{K}_{3,5}$. A few moments reflection should suffice to convince the reader that A_0 and A_1 are both axes for 3-braid structures. Also, A_0 and A_1 are *not isotopic* in $S^3 - \mathbf{K}_{3,5}$. By iterating the “winding” process which produced A_1 from A_0 one can obtain infinitely many pairwise disjoint 3-braid axes $\dots, A_{-2}A_{-1}, A_0, A_1, A_2, \dots$, which are pairwise non-isotopic in $S^3 - \mathbf{K}_{3,5}$. On the other hand, by our earlier observations, the links $\mathbf{K}_{3,5} \cup A_r$ and $\mathbf{K}_{3,5} \cup A_s$ are equivalent for each pair r, s under a homeomorphism $h_{r,s}$ of S^3 .

We first encountered these examples in the following way. Bennequin surfaces were an essential tool in the work in this paper. However Bennequin’s original proof that they exist was unclear to us at a crucial point. We were thus led to study his theorem in detail, and to fill in the gaps in his proof, and when we did so our work (see [B-M, II]) revealed the existence of an infinite family $(F_s; s \in N)$ of “pathological” Markov surfaces which Bennequin had missed. They are *not* Bennequin surfaces, yet their boundaries are (composite) links

of braid index 3. In fact, if we choose a Bennequin surface F for the axis A_0 of Figure 8.1, the surface F_s is $h_{0,s}(F)$.

The generalization of the pathology just described to composite links (and also to split links) of braid index n is the principal focus of the work in [B-M, IV].

REFERENCES

- [A1] J. W. Alexander, *A lemma on systems of knotted curves*, Proc. Nat. Acad. Sci. U.S.A., **9** (1923), 93–95.
- [Ar] E. Artin, *Theorie der Zöpfe*, Abh. Math. Sem. Univ. Hamburg, **4** (1925), 47–72.
- [Be] Daniel Bennequin, *Entrelacements et equations de Pfaff*, Astérisque, **107-108** (1983), 87–161.
- [Bi,1] Joan S. Birman, *Braids, Links and Mapping Class Groups*, Ann. of Math. Studies #82, Princeton University Press, 1974.
- [Bi,2] ———, *On the Jones polynomial of closed 3-braids*, Invent. Math., **81** (1985), 287–294.
- [B-M,I] Joan S. Birman and William W. Menasco, *Studying links via closed braids I: A finiteness theorem*, Pacific J. Math., **154** (1992), 17–36.
- [B-M,II] ———, *Studying links via closed braids II: On a theorem of Bennequin*, Topology Appl., **40** (1991), 71–82.
- [B-M,IV] ———, *Studying links via closed braids IV: Composite and split links*, Invent. Math., **102** (1990), 115–139.
- [B-M,V] ———, *Studying links via closed braids V: The unlink*, Trans. Amer. Math. Soc., **329** (1992), 585–606.
- [B-M,VI] ———, *Studying links via closed braids VI: A non-finiteness theorem*, Pacific J. Math., **156** (1992), 265–285.
- [F] Elizabeth Finkelstein, *Computer program available on request*, Dept. of Mathematics, Columbia University, New York, NY 10027.
- [G] F. Garside, *The braid groups and other groups*, Quart. J. Math. Oxford, **20** No. 78 (1969), 235–254.
- [H] Robert D. Horowitz, *Characters of free groups represented in the 2-dimensional special linear group*, Comm. Pure Appl. Math., **25** (1972), 635–659.
- [J] V. F. R. Jones, *Annals of Math.*, **126** (1987), 335–388.
- [Ma] A. A. Markov, *Über die freie äquivalenz der geschlossenen Zöpfe*, Recueil de la Soc. Math. de Moscou, **43** (1936), 73–78.
- [Mo,1] Hugh R. Morton, *Closed braids which are not prime knots*, Math. Proc. Cambridge Philos. Soc., **93** (1979), 421–426.
- [Mo,2] ———, *Threading knot diagrams*, Math. Proc. Cambridge Philos. Soc., **99** (1986), 247–260.
- [Mo,3] ———, *Infinitely many fibered knots with the same Alexander polynomial*, Topology, **17** (1978), 101–104.
- [Mu] K. Murasugi, *On closed 3-braids*, Mem. Amer. Math. Soc., **151** (1976).
- [M-P] W. Magnus, and A. Pelusco, *On a theorem of V. I. Arnold*, Comm. Pure Appl. Math., **22** (1968), 683–692.
- [M-T] K. Murasugi and R. S. D. Thomas, *Isotopic closed non-conjugate braids*, Proc. Amer. Math. Soc., **33** (1972), 137–139.

- [Ro] D. Rolfsen, *Knots and Links*, Publish or Perish, Inc., 1976.
- [Ru] L. Rudolph, *Braided surfaces and Seifert ribbons for closed braids*, *Comment. Math. Helv.*, **58** (1983), 1–37.
- [Sc] O. Schreier, *Über die Gruppen $A^a B^b = 1$* , *Abh. Math. Sem. Univ. Hamburg*, **3** (1924), 167–169.
- [T] Takahashi, *Explicit formulas for the Jones polynomials of closed 3-braids*, preprint.
- [X] P. J. Xu, *The genus of closed 3-braids*, *J. of Knot Theory and its Ramifications*, **1**, No. 3 (1992), 303–326.
- [V-B] J. Van-Buskirk, *Prime positive knots with inconjugate minimal-string braid representatives*, preprint, Univ. of Oregon.
- [Y] S. Yamada, *The minimum number of Seifert circles equals the braid index of a link*, *Invent. Math.*, **89** (1987), 347–356.

Received November 15, 1990 and in revised form July 15, 1992. First author partially supported by NSF Grants #DMS-88-055672 and #DMS-91-06584. Second author partially supported by NSF Grant #DMS-90-02673.

COLUMBIA UNIVERSITY
NEW YORK, NY 10027

AND

SUNY AT BUFFALO
BUFFALO, NY 14222

