Invent. math. 102, 115–139 (1990)

Inventiones mathematicae

Studying links via closed braids IV: composite links and split links

Joan S. Birman[1,*] **and William W. Menasco**[2]

[1] Department of Mathematics, Columbia University, New York, NY 10027, USA
[2] Department of Mathematics, SUNY at Buffalo, Buffalo, NY 14222, USA

Oblatum 29-XI-1988 & 25-I-1990

Summary. The main result concerns changing an arbitrary closed braid representative of a split or composite link to one which is obviously recognizable as being split or composite. Exchange moves are introduced; they change the conjugacy class of a closed braid without changing its link type or its braid index. A closed braid representative of a composite (respectively split) link is composite (split) if there is a 2-sphere which realizes the connected sum decomposition (splitting) and meets the braid axis in 2 points. It is proved that exchange moves are the only obstruction to representing composite or split links by composite or split closed braids. A special version of these theorems holds for 3 and 4 braids, answering a question of H. Morton. As an immediate Corollary, it follows that braid index is additive (resp. additive minus 1) under disjoint union (resp. connected sum).

1. Introduction

This paper is part of a series of papers ([B-M, I, II, III, IV and V] in which the authors study representations of oriented links in oriented S^3 by closed braids. The long-term goal is to produce computable link type invariants by making use of one of the known solutions to the conjugacy problem in the braid group (e.g. [G]). Any such effort must deal with the various mechanisms which produce more than one conjugacy class of braids in a braid group B_n which represent the same link type. Two such mechanisms are treated in this paper.

A link $\boldsymbol{K}$ is *split* if there is a 2-sphere $\mathbf{X}$ in S^3-$\boldsymbol{K}$ which does not bound a 3-ball. A closed braid representative $\mathbf{K}$ of a split link $\boldsymbol{K}$ is a *split closed braid* if the splitting 2-sphere $\mathbf{X}$ can be chosen to meet the braid axis $\mathbf{A}$ in exactly two points. The closed 4-braid α_1 in Fig. 1 is an example of a split closed 4-braid which represents the disjoint union of two trefoil knots.

Since every link can be represented by a closed braid, it is obvious that every split link can be represented by a split closed braid. For, if $\boldsymbol{K}$ is the disjoint union of

* Partially supported by NSF Grant #DMS-88-05672

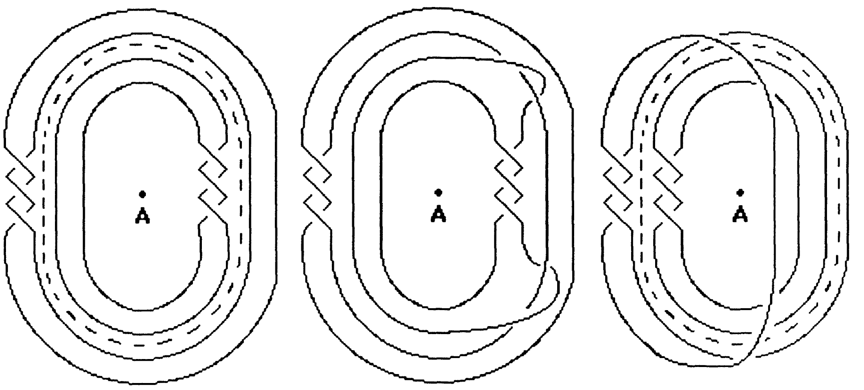

Fig. 1. Closed 4-braid representatives of the disjoint union of two trefoils

links $\boldsymbol{K}_1$ and $\boldsymbol{K}_2$, choose an n_1-braid representative $\mathbf{K}_1$ of $\boldsymbol{K}_1$ and an n_2-braid representative $\mathbf{K}_2$ of $\boldsymbol{K}_2$ and use them to construct the obvious $(n_1 + n_2)$-braid representative $\boldsymbol{K}$ by banding together their respective axes. There is, however, a more subtle question, i.e. are there n-braid representatives of split links which are not isotopic in the complement of the axis to split closed n-braids?

The answer is "yes", and an example is given by the closed 4-braid α_2 in Fig. 1, which also represents the disjoint union $\boldsymbol{K}$ of two trefoil knots. If α_2 were isotopic in the complement of the axis to a split 4-braid representative of $\boldsymbol{K}$, it would necessarily be conjugate to α_1 because there is a unique closed 2-braid which represents the trefoil, and so a unique split closed 4-braid which represents the disjoint union of two trefoils. Let the symbol "i" (resp. "i^{-1}") denote an elementary braid in which the i^{th} strand crosses once over (under) the $(i+1)^{\text{st}}$ strand. Then α_1 and α_2 may be described by the braid words:

$$\alpha_1 = 1^3 3^3 ,$$

$$\alpha_2 = 1^3 2^{-2} 3^3 2^2 .$$

Now, there is a homomorphism $\phi\colon B_4 \to B_3$ defined by $\phi(1) = \phi(3) = 1$, $\phi(2) = 2$. If α_1 and α_2 were conjugate in B_4 then $\phi(\alpha_1)$ would be conjugate to $\phi(\alpha_2)$ in B_3. Using the solution to the conjugacy problem in B_3 which is given in [Mu] we obtain unique representatives 1^6 for the conjugacy class of $\phi(\alpha_1)$ and $\Delta^2 1^2 2^{-2} 1^2 2^{-2}$ for the class of $\phi(\alpha_2)$. Thus α_1 and α_2 are not conjugate in B_4. By a theorem of Morton [Mo, 3] this shows that the corresponding closed braids are not isotopic in the complement of $\mathbf{A}$. Infinitely many such examples can be obtained by iterating the winding process which is shown in Fig. 1.

Remark. The example α_3 in Fig. 1 is also a split closed 4-braid representative of $\boldsymbol{K}$. We give it to show that the intersection of the splitting 2-sphere with the plane of projection need not be disjoint from the projected image of $\mathbf{K}$.

Similar phenomena arise when we consider composite links. A link $\boldsymbol{K}$ is *composite* if there is a 2-sphere $\mathbf{Y}$ in S^3 which meets the link in 2 points and decomposes it into sublinks $K_1 \# K_2$, neither of which is an unknotted arc. A closed braid representative **K** of K is *a composite closed braid* if the 2-sphere **Y** can be chosen to meet the braid axis **A** in 2 points. An example is given in Fig. 2. One may ask whether every closed n-braid representative of a composite link is conjugate to a composite closed n-braid? That question has received some attention in the literature. Morton showed that every closed 3-braid which represents a composite link is conjugate to an obviously composite 3-braid. In the same paper he also gave an example of a closed 5-braid which represents a composite link, but is not conjugate to a composite 5-braid. He was unable to decide whether 4-braids resemble 3 or 5-braids in this regard.

We will show in this paper how an arbitrary braid representative of a split (or composite) link can be changed to a split (or composite) braid representative. To state our contributions, we need the concept of an *exchange move*. In Fig. 3 the labels n_j on the strands are weights, indicating n_j parallel strands. We allow any type of braiding on the $n_1 + n_2$ (resp. $n_1 + n_3$) strands in the boxes which are labeled X(resp. Y). Assume that the braid axis **A** is the z axis, and that the arc which is labeled n_3 lies in the $x - y$ plane. Up to isotopy of S^3, an exchange is defined to be an isotopy of **K** which moves the arc which is labeled n_2 from a position which is a little bit above (or below) the $x - y$ plane to a position which is a little bit below (or above) the $x - y$ plane, keeping the rest of **K** invariant. An exchange move takes n-braids to n-braids, but need not preserve conjugacy class because the isotopy of **K** in S^3 is in general not realizable in the complement of the axis **A**. The first two results in this paper show that exchange moves are the only obstruction to representing split (or composite) links by split (or composite) closed braids:

The split braid theorem. *Let K be a split link, and let* **K** *be an arbitrary closed n-braid representative of K. Then there exists a split n-braid* $\mathbf{K}^\bullet$ *which represents* **K** *and a*

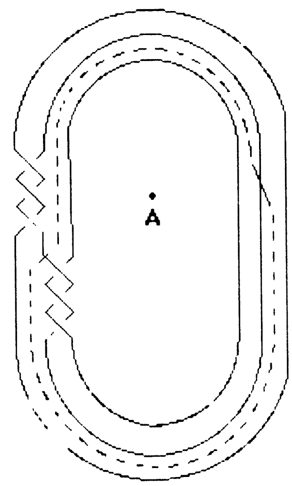

Fig. 2. Composite closed braid representing the connected sum of two trefoils

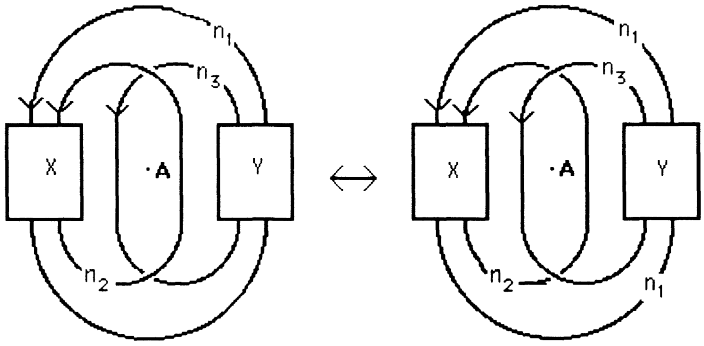

Fig. 3

finite sequence of closed n-braids:

$$\mathbf{K} = \mathbf{K}_0 \to \mathbf{K}_1 \to \mathbf{K}_2 \to \ldots \to \mathbf{K}_m = \mathbf{K}^{\bullet}$$

such that each $\mathbf{K}_{i+1}$ *is obtained from* $\mathbf{K}_i$ *by either isotopy in the complement of the axis or an exchange.*

The composite braid theorem. *Let* K *be a composite link, and let* $\mathbf{K}$ *be an arbitrary closed n-braid representative of* K*. Then there exists a composite n-braid* $\mathbf{K}^{\bullet}$ *which represents* K *and a finite sequence of closed n-braids*:

$$\mathbf{K} = \mathbf{K}_0 \to \mathbf{K}_1 \to \mathbf{K}_2 \to \ldots\ldots \to \mathbf{K}_m = \mathbf{K}^{\bullet}$$

such that each $\mathbf{K}_{i+1}$ *is obtained from* $\mathbf{K}_i$ *by either isotopy in the complement of the axis or an exchange.*

The theorem of Markov ([Ma] or [Bi] or [Be] or [Mo, 2]) shows that, since the 4-braids α_1 and α_2 of Fig. 1 (or Morton's two 5-braids) represent the same oriented link type, they are related by a sequence of moves which involve first increasing and then decreasing the string index. This process is known as *stabilization*. Our theorems replace the stabilization process with a more direct process which changes one conjugacy class to another, preserving string index. It is thus a step in the proof of a version of "Markov's theorem without stabilization". A subsequent paper in this series will prove such a theorem in full generality.

The *braid index* of a link $\mathbf{K}$ is the smallest integer n such that $\mathbf{K}$ can be represented by an n-braid. As an immediate corollary of the Split and Composite Braid Theorems we obtain:

The braid index theorem. *Braid index is additive* (*respectively additive minus* 1) *under disjoint union* (*respectively connected sum*).

Our next result relates to special versions of the Split and Composite Braid Theorems which hold when $n = 3$ and 4.

The 3 and 4-braid theorem. (*i*) *Every* 3*-braid representative of a split link is conjugate to a split* 3*-braid representative. However, there are examples of* 4*-braid representatives of split links which are not conjugate to split* 4*-braids.*

(*ii*) *Every* 3-*braid representative of a composite link is conjugate to a composite* 3-*braid representative.*

(*iii*) *Let* **K** *be a* 4-*braid representative of a composite link* ***K***. *If* ***K*** *cannot be represented by a closed braid of braid index* < 4, *then* **K** *is conjugate to a composite* 4-*braid. However, if* ***K*** *can be represented by a closed braid of braid index* < 4, *then* **K** *need not be conjugate to a composite* 4-*braid.*

2. Proof of the split braid theorem

Our work begins with a split link ***K***, and an arbitrary oriented closed braid representative **K** of ***K***. The braid axis is **A**. We are also given a 2-sphere **X** which realizes the splitting of **K**. We orient **K**, and then orient **A** so that **K** is oriented in the positive sense about **A**. We assign (arbitrarily) an orientation to **X**, so that at each point of **X** there is a well-defined outward-drawn normal to **X**. The 2-sphere **X** will in general be pierced by the braid axis **A** many times. Our goal is to modify **K** until there is a 2-sphere **X**′ which realizes the splitting and is pierced twice by **A**.

The principle tool in our work will be the study of the singular foliation of **X** which is induced by the fibration **H** of the open solid cylinder $R^3 - \mathbf{A}$ by half-planes $\{\mathbf{H}_t;\, t \in [0, 2\pi]\}$. The leaves of this foliation are the components of $\mathbf{X} \cap \mathbf{H}_t$, $t \in [0, 2\pi]$. The first step in the proof is to put **X** into a nice position relative to the fibration, in order to partially standardize the foliation.

By standard general position arguments we may assume:

(*i*) The intersections of **A** with **X** are finite in number and transverse.
(*ii*) There is a solid torus neighborhood $N(\mathbf{A})$ of **A** in S^3-**K** such that each component of $\mathbf{X} \cap N(\mathbf{A})$ is a disc.
(*iii*) The foliation in each component of $\mathbf{X} \cap N(\mathbf{A})$ is the standard radial foliation.
(*iv*) All but finitely many $\mathbf{H}_t$'s meet **X** transversally, and those which do not (the *singular fibers*) are each tangent to **X** at exactly one point in the interior of $\mathbf{H}_t$.
(*v*) The tangencies in (*iv*) are local maxima or minima or saddle points.

A *singular leaf* in the foliation of **X** will be one which contains a point of tangency. All other leaves are *non-singular*. Note that it follows from (*iv*) and (*v*) that:

(*a*) Each non-singular leaf is either an arc with both endpoints on $\mathbf{A} = \partial \mathbf{H}_t$ or a simple closed curve.
(*b*) A singular fiber $\mathbf{H}_\theta$ contains exactly one singular point p_θ.
(*c*) Each singular point p_θ is either a center or a saddle.

We now introduce a measure of the complexity of the pair (**X**, **H**). Let $|\mathbf{X} \cap \mathbf{A}|$ be the number of points in $\mathbf{X} \cap \mathbf{A}$. Let $|\mathbf{H} \cdot \mathbf{D}|$ be the number of singular points in the foliation of **X**. The *complexity* $c(\mathbf{X}, \mathbf{H})$ is the pair $(|\mathbf{X} \cap \mathbf{A}|, |\mathbf{H} \cdot \mathbf{D}|)$. We assign the standard lexicographic ordering to this complexity function. We will say that (**X**, **H**) is *equivalent* to (**X**′, **H**) if there is an isotopy taking **X** to **X**′ which takes $(\mathbf{X} \cap \mathbf{H}_t, \mathbf{X} \cap \partial \mathbf{H}_t)$ to $(\mathbf{X}' \cap \mathbf{H}_t, \mathbf{X}' \cap \partial \mathbf{H}_t)$, for all $t \in [0, 2\pi]$. Notice that by our definition of equivalence, every representative of an equivalence class has the same complexity.

We investigate simple closed curves (SCC's) in $\mathbf{X} \cap \mathbf{H}$. We say that a pair (**X**, **H**) *has SCC's* if there exists a non-singular fiber $\mathbf{H}_t$ such that a component of $\mathbf{X} \cap \mathbf{H}_t$ is a SCC.

Lemma 1. *Assume that* $(\mathbf{X}, \mathbf{H})$ *satisfies* (i)–(v) *and has SCC's. Then there exists a splitting 2-sphere* $\mathbf{X}'$ *such that* $(\mathbf{X}^{\bullet}, \mathbf{H})$ *also satisfies* (i)–(v), *and in addition has no SCC's, moreover,* $c(\mathbf{X}^{\bullet}, \mathbf{H}) < c(\mathbf{X}, \mathbf{H})$.

Proof of Lemma 1. Our proof, sketched below, is very similar to the argument given in Lemma 2 of [B-M, 1]. The main difference is that our surface $\mathbf{X}$ is closed whereas the surface $\mathbf{D}$ in the proof of Lemma 2 of [B-M, 1] had non-empty boundary $\partial\mathbf{D} = \mathbf{K}$.

If there is a SCC $\alpha(t)$ in $\mathbf{X} \cap \mathbf{H}_t$ for some non-singular $\mathbf{H}_t$, we proceed as in the proof of Lemma 2 of [B-M, 1], following $\alpha(t)$ as it evolves in the fibration, until we arrive at a simple closed curve $\alpha(\theta)$ which contains a singularity of the foliation. Now $\alpha(\theta)$ lies on a singular fiber $\mathbf{H}_\theta$, and bounds a disc Δ in $\mathbf{H}_\theta$. Notice that Δ cannot be punctured by $\mathbf{K}$. For, the algebraic and geometric intersection numbers of $\mathbf{K}$ with Δ coincide, because Δ is a subdisc of $\mathbf{H}_\theta$ and $\mathbf{K}$ is a closed braid with axis $\partial\mathbf{H}_\theta$. There is also a second disc Δ', in $\mathbf{X}$, with $\partial\Delta = \partial\Delta'$. Since $\mathbf{K}$ does not intersect $\mathbf{X}$, it also does not intersect Δ'. Now, $\Delta \cup \Delta'$ forms a 2-sphere $\mathbf{S}$ in S^3, and the algebraic intersection number of $\mathbf{S}$ with $\mathbf{K}$ is zero because $\mathbf{S}$ is a 2-sphere. Since the geometric intersection number of Δ with $\mathbf{K}$ is equal to the algebraic intersection number of $\mathbf{K}$ with $\mathbf{S}$, it follows that $\mathbf{K}$ does not pierce Δ.

If the interior of Δ has empty intersection with $\mathbf{X}$, we surger $\mathbf{X}$ along Δ. The surgered surface $\mathbf{X}^{\sim}$ will be a pair of two 2-spheres, $\mathbf{X}^{\sim}(1)$ and $\mathbf{X}^{\sim}(2)$. At least one of these will be a new splitting 2-sphere, say $\mathbf{X}^{\sim}(1)$, and as in the proof of Lemma 2 of [B-M, 1] the new splitting 2-sphere will have complexity which is no more than that of the original $\mathbf{X}$. If both $\mathbf{X}^{\sim}(1)$ and $\mathbf{X}^{\sim}(2)$ are splitting 2-spheres, it won't matter which one we retain. If the interior of Δ meets $\mathbf{X}$, we find an innermost subdisc δ of Δ whose boundary is a component of $\mathbf{X} \cap \mathbf{H}_\theta$ and surger $\mathbf{X}$ along δ. Ultimately, we will arrive at a splitting 2-sphere $\mathbf{X}'$ which has the property that the induced foliation has no simple closed curves, and $c(\mathbf{X}^{\bullet}, \mathbf{H}) < c(\mathbf{X}, \mathbf{H})$. □

By Lemma 1, each component of intersection of a non-singular $\mathbf{H}_t$ with our 2-sphere $\mathbf{X}$ is an arc with both of its endpoints on $\mathbf{A} = \partial\mathbf{H}_t$. Such an arc divides $\mathbf{H}_t$ into two components, so it bounds two discs Δ and Δ' on $\mathbf{H}_t$. If either of these, say Δ, is *not* pierced by the link $\mathbf{K}$, we can push our 2-sphere $\mathbf{X}$ inward along a 3-space neighborhood of Δ to remove two points of intersection of $\mathbf{A}$ with $\mathbf{X}$. See Fig. 4. Of course, this could add simple closed curves to the foliation of $\mathbf{X}$, but the definition of our complexity function was chosen so that even if it does, the complexity will be reduced, because we will have reduced $|\mathbf{A} \cap \mathbf{X}|$. Moreover, even if SCC's are added, they can be removed with the aid of Lemma 1.

This motivates us to call a leaf β in $\mathbf{H}_t \cap \mathbf{X}$ *essential* if $\mathbf{K}$ pierces both components of $\mathbf{H}_t$ split along β. See Fig. 4 again. In view of the above remarks and Lemma 1, we can replace assumption (v) by the stronger assumptions:

$(v)^*$ The tangencies in (iv) are saddle points.
(vi) If $\mathbf{H}_t$ is non-singular, then each component of $\mathbf{X} \cap \mathbf{H}_t$ is essential.

We return to the main thread of the argument. The surface $\mathbf{X}$ is closed, so the braid axis $\mathbf{A}$ must pierce $\mathbf{X}$ an even number of times, say 2μ. If $2\mu = 2$, our closed braid is a split closed braid, and we are done, so assume $2\mu > 2$. Thus each component of intersection of $\mathbf{X}$ with a non-singular $\mathbf{H}_t$ is an arc which joins two of the 2μ points where the braid axis $\mathbf{A}$ pierces the 2-sphere $\mathbf{X}$. Recall that we assumed (see (iii) above) that the foliation of $\mathbf{X}$ is radial near each point of $\mathbf{A} \cap \mathbf{X}$. From this and the

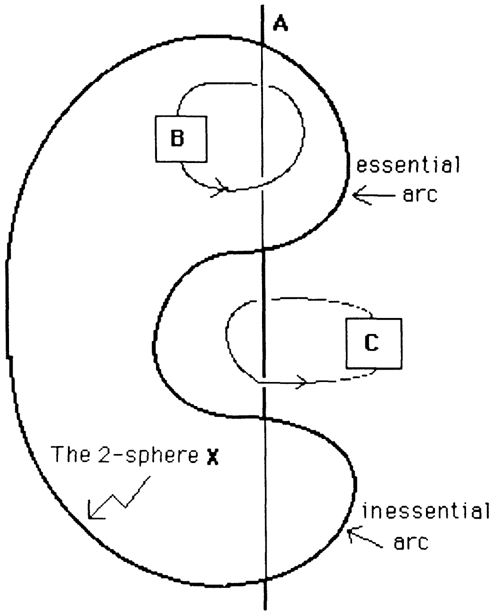

Fig. 4

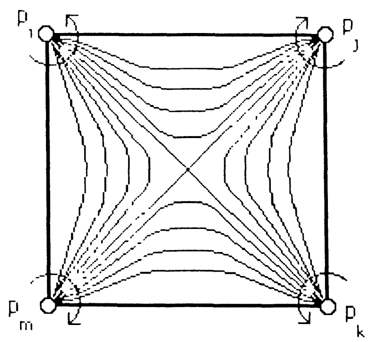

Fig. 5

fact that $2\mu > 2$ it follows that there must be singularities in the foliation. Each singularity is the result of a saddle-point type tangency of a singular fiber $\mathbf{H}_\theta$ with $\mathbf{X}$. Figure 5 shows how the foliation looks in a neighborhood of a singular point. There will be four singular leaves which go out from the singular point like the spokes on a wheel, and end at points where the axis $\mathbf{A}$ pierces $\mathbf{X}$.

Let $p_1, \ldots, p_{2\mu}$ be the points of $\mathbf{A} \cap \mathbf{X}$. Let $\theta_1, \ldots, \theta_q$ be the t-values at which singularities occur. The singular leaves are then the leaves in $\{\mathbf{H}_\theta \cap \mathbf{X};$

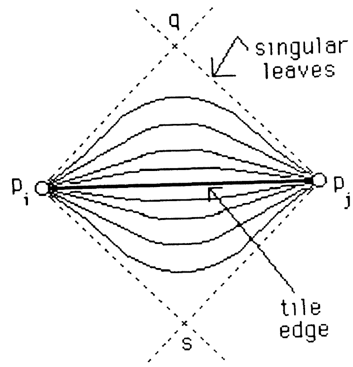

Fig. 6

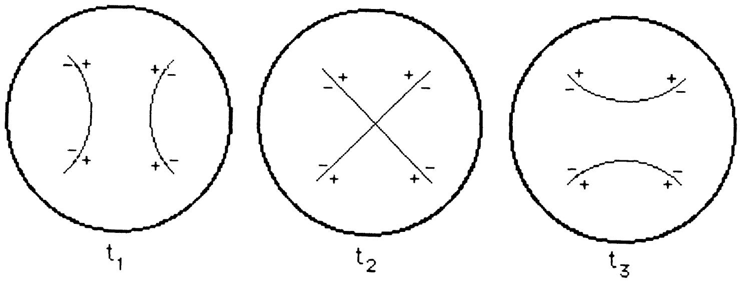

Fig. 7

$\theta = \theta_1, \ldots, \theta_q\}$ which contain the singular points. Their complement in **X** can only be a union of regions $R_1, \ldots, R_k$, as in Fig. 6, each of which is foliated without singularities. Choose a leaf e_j from each R_j. The union of all of these leaves, $\{e_1, \ldots, e_k\}$ gives a cell decomposition of **X**. The 2μ points $p_1, \ldots, p_{2\mu}$ where **A** pierces **X** are the 0-cells, the non-singular arcs $\{e_1, \ldots, e_k\}$ are the 1-cells, and the 2-cells each contain one singularity of the foliation. Every 2-cell has four vertices and 4 edges (Fig. 5). We call our 2-cells *tiles*, and the foliated cell decomposition a *tiling* of **X**.

The *sign* of a tile is defined as follows: each tile contains exactly one singularity of the foliation, which occurs at a point of tangency between **X** and a fiber $\mathbf{H}_\theta$ of **H**. At the point of tangency the normals to **X** and $\mathbf{H}_\theta$ coincide. The sign of a tile is *positive* or *negative*, according as the orientations on the normals to **X** and $\mathbf{H}_\theta$ agree or disagree. Notice that if $(\theta - \varepsilon, \theta + \varepsilon)$ is a t-interval about θ which does not include any other singularities, then we can distinguish the two cases in the following way: Choose any $t \in (\theta - \varepsilon, \theta + \varepsilon)$ and label the sides of the arcs in $\mathbf{H}_t \cap \mathbf{X}$ "$+$" or "$-$" according as they correspond to the positive or negative sides of the oriented surface **X**. Then as $t \to \theta$ at a positive (resp. negative) singularity the $+$ (resp. $-$) sides of two arcs of $\mathbf{H}_t \cap \mathbf{X}$ will appear to coalesce, as in the left-to-right (resp. right-to-left) sequence in Fig. 7.

The *valence* of a vertex in the tiling is the number of tile edges which meet at that vertex. The next few lemmas concern the existence of vertices of low valence in the tiling. We will ultimately use these to recognize when our link admits a complexity-reducing exchange move.

Lemma 2. *The tiling of* **X** *always contains either a vertex of valence* 2 *or* 3.

Proof of Lemma 2. Let V, E and F denote the number of 0, 1 and 2-cells in the tiling. Then $V-E+F=2$, the Euler characteristic of **X**. Now, each 2-cell in the tiling has four edges in its boundary, and each edge is an edge of exactly two 2-cells, so that $2F = E$, and so $2V - E = 4$. Now let V_i be the number of vertices of valence i in the tiling. Since there are no vertices of valence 1, we have $V = V_2 + V_3 + V_4 + \dots$. Since each edge has 2 vertices in its boundary, we also have that $2E = 2V_2 + 3V_3 + 4V_4 + \dots$. Thus

$$2V_2 + V_3 = 8 + V_5 + 2V_6 + 3V_7 + \dots .$$

The terms on both sides of this equation are non-negative. Thus there is always a vertex of valence 2 or 3. □

Lemma 3. *Either the tiling of* **X** *has a vertex of valence* 2, *or else there is an isotopy of* **X** *to a new splitting* 2*-sphere* **X**′, *such that the tiling of* **X**′ *has a vertex of valence* 2. *Moreover* $c(\mathbf{X}', \mathbf{H}) \leqq c(\mathbf{X}, \mathbf{H})$.

Proof of Lemma 3. If the tiling of **X** already has a vertex of valence 2 there is nothing to prove. If not, then by Lemma 2 the tiling has a vertex p of valence 3. There are 3 tiles which meet at this vertex (see Fig. 8a) and two of them (call them T

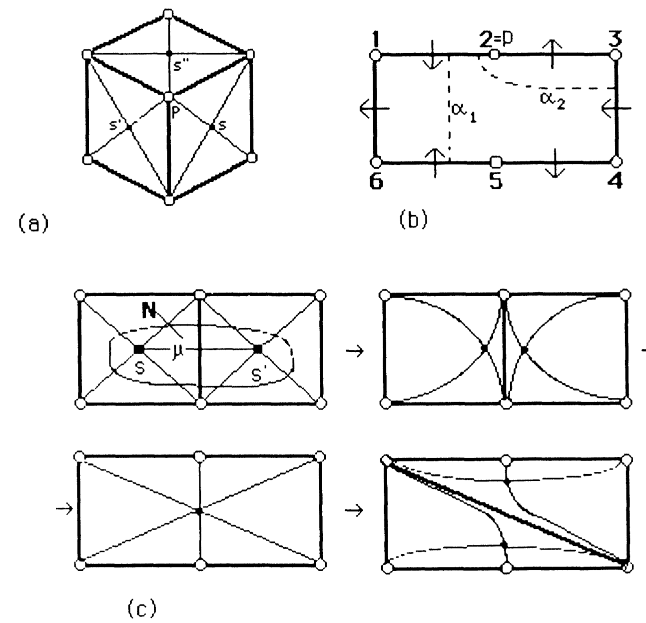

Fig. 8. **a** Vertex p has valence 3; **b** flow $T \cup T'$; **c** proposed changes in tiling $T \cup T'$

and T') necessarily have the same sign. Let their singular points be s and s'. Label the six vertices of $T \cup T'$ **1**, **2**, . . . , **6** in cyclic order, so that the vertex $p = \mathbf{2}$, as in Fig. 8b. Assume the origin to have been chosen so that the singularity at s occurs before the singularity at s'.

The future singularities have been indicated in Fig. 8b by two "joining arcs", labeled α_1 and α_2. We now explain what these mean. See Fig. 9, which shows the surface **X** and a fiber $\mathbf{H}_t$ of **H** just before one of the saddle point tangencies. We see two arcs, $\alpha \subset \mathbf{X}$ and $\beta \subset \mathbf{H}_t$. Together they bound a disc Δ in 3-space. These are our *joining arcs*. As $\mathbf{H}_t$ is pushed up the arcs α and β will shrink, vanishing at the instant of tangency. Now, in *our* situation there will be *two* singularities (at s and s'), and so there will be *two* pairs of joining arcs: α_1 and α_2 on **X** (illustrated in Fig. 8b) and β_1 and β_2 on $\mathbf{H}_t$ (illustrated in Fig. 10a). The idea of our proof is to show that we can deform **X** so that the two singularities are very close together both on **X** and in the fibration (we will make this precise below), and then interchange the order in which the singularities occur. The sequence of four pictures in Fig. 8c illustrate what we intend to do, from the point of view of an observer who is looking at the changing foliation of **X** during the deformation. After the deformation the tiling of $T \cup T'$ will have been changed so that the vertex p of valence 3 becomes a vertex of valence 2, also the tiling will be unchanged in $\mathbf{X} - T \cup T'$. We now show that these changes are always possible.

The first thing we want to prove is that the six tile vertices **1**, **2**, . . . , **6**, which are points on both **X** and **A**, must also have cyclic order **1**, **2**, . . . , **6** on **A**, as illustrated in the Fig. 10a, which depicts a fiber of **H** before the two singularities. There is no natural choice of a normal bundle on **X**, so we assign one arbitrarily so that **A** (oriented either way) pierces **X** from the $+$ side at the vertex **1**. Note that this determines how **A** pierces **X** at every other vertex, because if the flow viewed from the $+$ side of **X** is clockwise about any one vertex, then it is anticlockwise about any vertex which is joined to the given one by a leaf in the foliation. So, proceeding from vertex to vertex, the sense of the flow is determined everywhere. This, in turn, determines whether the oriented axis **A** pierces **X** from the $+$ side or the $-$ side at each vertex.

Let $\mathbf{ij}(t)$ denote a leaf in the foliation which joins the vertices **i** and **j**. In the given tiling of **X** (Fig. 8b) the first singularity is between $\mathbf{12}(t)$ and $\mathbf{56}(t)$. This shows that **1**

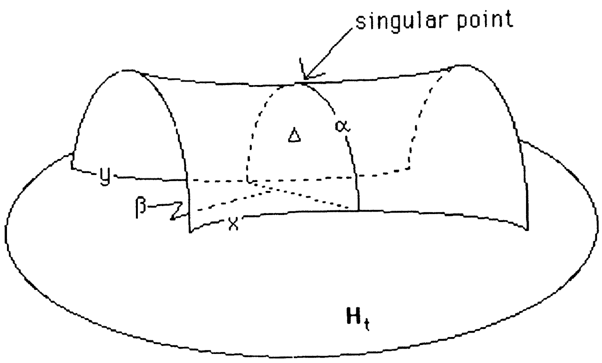

Fig. 9

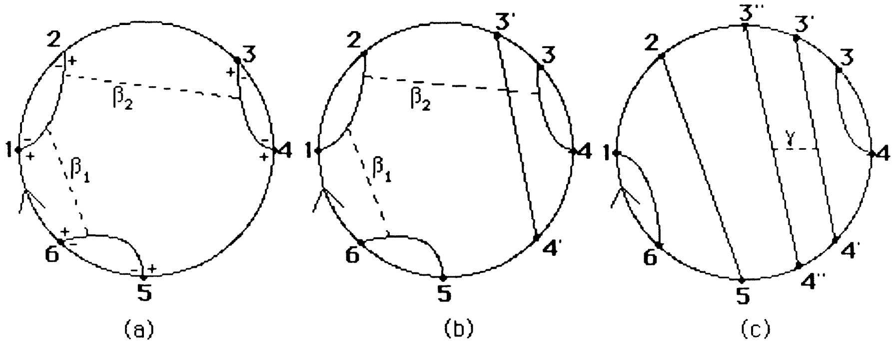

Fig. 10

and **2** cannot separate **5** and **6** on $\mathbf{A} = \partial \mathbf{H}_t$. Thus the cyclic order of the four vertices **1**, **2**, **5**, **6** on the oriented curve **A** is **1256**, **2165**, **1265** or **2156**. However, the latter two are impossible because in the future surgery the sides of $\mathbf{12}(t)$ and $\mathbf{56}(t)$ which coalesce must have the same sign, and the sign is determined by our condition on the pierce-points. The first two choices are equivalent up to symmetry, and determined by whether the surgery is to be "+" or "−". We choose the order **1256**, dictating that the first surgery is to be a "+" surgery. Similar considerations apply to the pairs **2**, **5** and **3**, **4** giving two possible cyclic orders for the 6 points: **123456** and **125436**. However, with the latter choice the two singularities will have opposite signs, so the order must be **123456** as illustrated in Fig.10a.

We now have enough information to determine how $T \cup T'$ is embedded in 3-space. Choose 6 cyclically ordered points on **A** and declare them to be the tile vertices **1**, **2**, . . . , **6**. Then choose two fibers of **H** and declare them to be the fibers $\mathbf{H}_{\phi_1}$ and $\mathbf{H}_{\phi_2}$ which are to contain the singularities s and s' respectively. Then choose points s and s' in the fibers and declare them to be s and s'. Up to a homeomorphism of 3-space and a reparametrization of the fibers these choices are all arbitrary. Finally, join up s to the points **1**, **2**, **5**, **6** by arcs in $\mathbf{H}_{\phi_1}$ which have that cyclic order when $\mathbf{H}_{\phi_1}$ is viewed from the + side, and similarly join s' to the points **2**, **3**, **4**, **5** by arcs in $\mathbf{H}_{\phi_2}$, using that cyclic order when $\mathbf{H}_{\phi_2}$ is viewed from the + side. Thus we have embedded the singular leaves through s and s' in 3-space. The next thing to do is to adjoin little discs to the points **1**, **2**, . . . ,**6**. These discs are to be the intersections of **X** with a neighborhood of **A** in 3-space. Since **A** pierces **X** from the + side (resp. − side) at **1**, **3**, **5** (resp. **2**, **4**, **6**), we know how to embed these discs as oriented surfaces. There is then a unique way to extend the embedding to a neighborhood of the singular leaves in 3-space. Finally, let μ be an arc which joins s to s' in **X**, and let **N** be a foliated neighborhood of μ, as in Fig. 8c. Then there is a unique way to adjoin **N** to the part of $T \cup T'$ already constructed. Finally, extend the embedding to all of $T \cup T'$. The result is depicted in Fig. 11a. (We have removed a tubular neighborhood of the axis **A** from 3-space, so the points **1**, **2**, . . . , **6** are replaced by circles, seen here as vertical arcs.) For comparison we have also shown the embedding when the signs disagree, in Fig. 11b.

As noted earlier, we want to push s and s' together on **X** and in the fibration and then to interchange their order in the fibration. There is clearly no obstruction to

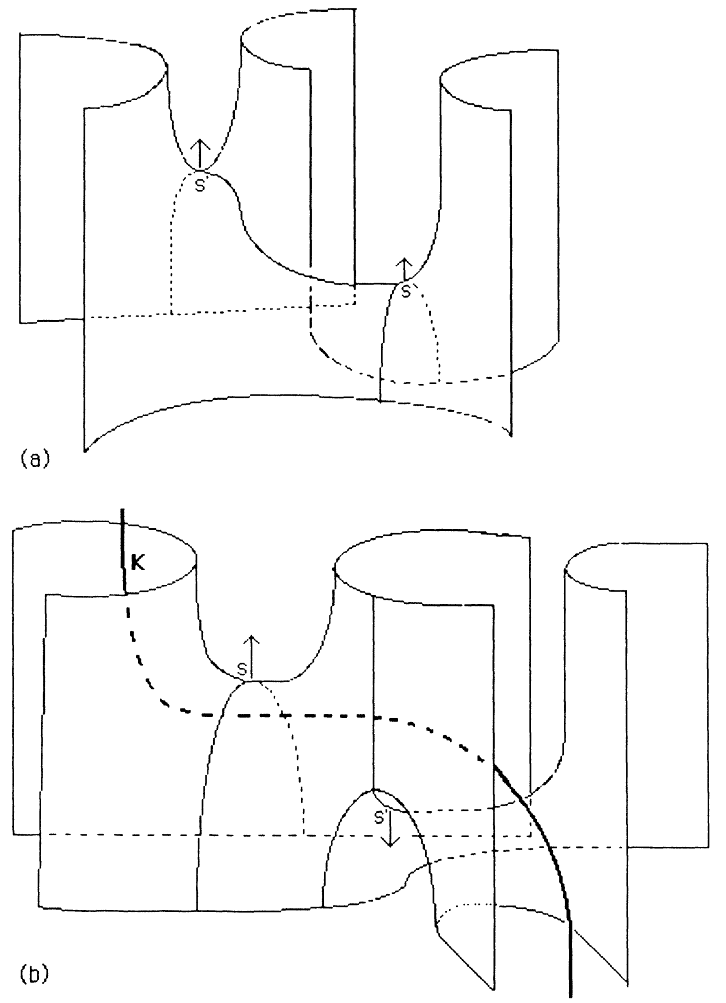

Fig. 11. a Non-degenerate saddles at s and s' signs at s and s' agree; **b** non-degenerate saddles at s and s' signs at s and s' do not agree

pushing s and s' close together on **X**. We now investigate whether there are obstructions in fibers of **H**. We first ask whether parallel sheets of **X** could interfere with the project (as they clearly might if the signs at s and s' disagree). Inspecting Fig. 10a, one potential difficulty becomes clear. It is possible that there is another leaf in the foliation of **X**, e.g. the leaf $\mathbf{3^{*}4^{*}}(t)$ illustrated in Fig. 10b, which is parallel to the leaf $\mathbf{34}(t)$ and obstructs the change we wish to make. It is not hard to see that, since we know that the singularity between $\mathbf{12}(t)$ and $\mathbf{56}(t)$ is possible, this is the only obstruction which can occur. After the first singularity the picture will be as in Fig. 10c. Now, since by hypothesis the singularity between $\mathbf{25}(t)$ and $\mathbf{34}(t)$ is

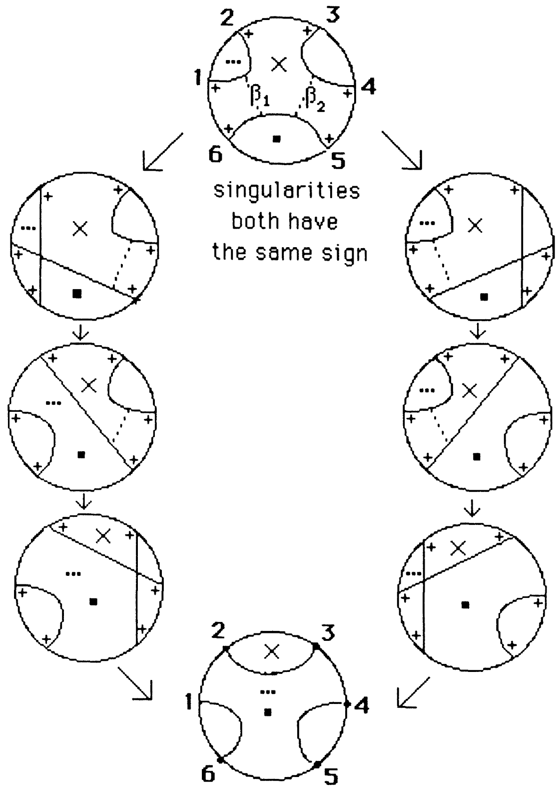

Fig. 12

possible, there must in fact be a *pair* of leaves which are parallel to $\mathbf{34}(t)$ in Fig. 10b, i.e. $\mathbf{3}^{\bullet}\mathbf{4}^{\bullet}(t)$ and $\mathbf{3}^{-}\mathbf{4}^{-}(t)$, and a singularity (indicated by the joining arc γ) between them which occurs *before* the singularity between $\mathbf{25}(t)$ and $\mathbf{34}(t)$. Now notice that on the surface **X** the four tile vertices $\mathbf{3}^{\bullet}$, $\mathbf{3}''$, $4^{\bullet}$, $\mathbf{4}''$ are disjoint from T and T' (and so not visible in Fig. 11a). Thus there is no obstruction to changing the fibration so that the singularity between $\mathbf{3}^{\bullet}\mathbf{4}^{\bullet}(t)$ and $\mathbf{3}^{-}\mathbf{4}^{-}(t)$ occurs *before* that between $\mathbf{12}(t)$ and $\mathbf{56}(t)$, removing the obstruction.

Since **K** intersects every fiber of **H**, we must also show that **K** cannot obstruct our proposed change in the order of the singularities. (Notice that **K** could be an obstruction in the situation of Fig. 11b, because if we were to push s down we would force **K** to be tangent to a fiber of **H**, so that it would no longer be a closed braid). Refer now to the sequence of pictures in Fig. 12. The top one is a repeat of Fig. 10a, with a little new data added. Initially, **K** will pierce one of the four regions of $\mathbf{H}_t$ split along the three arcs of $\mathbf{H}_t \cap \mathbf{X}$. By symmetry it suffices to consider three such choices. We have indicated these by a square black dot, a horizontal slash and an X in the three regions. We follow the time-evolution of the dot. It can wander anywhere within a region as t varies, but it cannot cross from one region to another because **K** does not intersect **X**. After each singularity, new regions open up to it. We have shown the final position of the square black dot, if the left sequence is used. We now see that we could have achieved the same end position if, instead, the right route had been used. The same argument applies to the slash and to the X. For every possible position it is possible to proceed either via the left or instead via the right sequence. Thus **K** is not an obstruction to the change in order of the singularities.

The proof is essentially complete, however we can improve our understanding of the changes by showing precisely how to do the deformation of $\mathbf{N} \subset T \cup T'$. We may regard **N** as the graph of the smooth function $f(x, y) = y^3 - 3y(x^2 - 1)$. This function has two non-degenerate saddle-type critical points of the same sign, one at $(x, y) = (1, 0)$ and the other at $(x, y) = (-1, 0)$. Now, our function $f(x, y)$ belongs to a parametrized family of functions f_ε, with $f_\varepsilon(x, y) = y^3 - 3y(x^2 - \varepsilon^2)$ and $f = f_1$. If $\varepsilon \neq 0$, the function $f_\varepsilon(x, y)$ has a pair of critical points at $(\pm\,\varepsilon, 0)$. Both critical points are saddles because at $(\pm\,\varepsilon, 0)$ the Hessian is $-36\varepsilon^2 < 0$. Deform **X** to a new surface **X**$'$ by letting the parameter ε pass to zero. This has the effect of pushing s and s' toward one-another, until at $\varepsilon = 0$ the two saddles coalesce to a monkey saddle, which has a single degenerate critical point at $(0, 0)$. The graph **X**$'$ of $f_0(x, y)$ has three "hills" and three "valleys". The singular leaves (i.e. the level sets $f_0^{-1}(0)$) are three lines in the $x - y$ plane through $(0, 0)$, of slope $0, \sqrt{3}$ and $-\sqrt{3}$. The surface **X**$'$ has local 3-fold rotational symmetry about the origin. We can then further deform **X**$^\bullet$ by rotating it through, say $-\pi/\sqrt{3}$, allowing ε to increase again, and then rotating it back again, i.e. replace f_ε by $g^{-1} f_\varepsilon g$, where g is the rotation. This has the effect of splitting the critical point into two new critical points along the line of slope $\sqrt{3}$. □

Lemma 4. *If the tiling contains a vertex of valence* 2, *then* **K** *admits an exchange. Moreover, after the exchange the complexity can be reduced by removing two or more of the points where the axis pierces* **X**.

Proof of Lemma 4. There are two tiles which meet at our vertex of valence 2, and the first step in the proof is to examine the signs of the singularities in these two tiles. Let p be the vertex in the statement of the lemma, and let q and s be the singular points of the two tiles which meet at p. Let p' and p'' be the other vertices of the two tile edges which meet at p. The singular leaves through q and s occur at t_1 and t_5, respectively. Portions of these singular leaves fit together to cut off a disc δ on **X**, and the region depicted in Fig. 13 is a neighborhood Δ of δ on **X**, chosen so that $\partial\Delta$ is everywhere transverse to the foliation. We have labeled leaves of the foliation as occurring at times $t = t_1, t_2, \ldots, t_8$.

We now examine corresponding leaves as they would occur on a sequence of fibers of **H**. See the sequence of five pictures in Fig. 14, which are labeled t_3, t_4, t_5, t_6, t_7 to correspond to the labels in Fig. 13. The full cycle is obtained from this one by adding to it the corresponding sequence of pictures run backwards. The singularities which occur at t_1 and t_5 have opposite signs. (This is a direct consequence of the fact that there are exactly two singular leaves among the leaves which are incident at p). Thus the normal bundles to **X** and **H** agree (say) at s, and disagree at q. This is the first fact which we need, in order to see how Δ is embedded in 3-space.

The next thing to notice is that (from Fig. 5) the gradient flow on **X** is always oriented in opposite senses around the two endpoints of each leaf which joins two

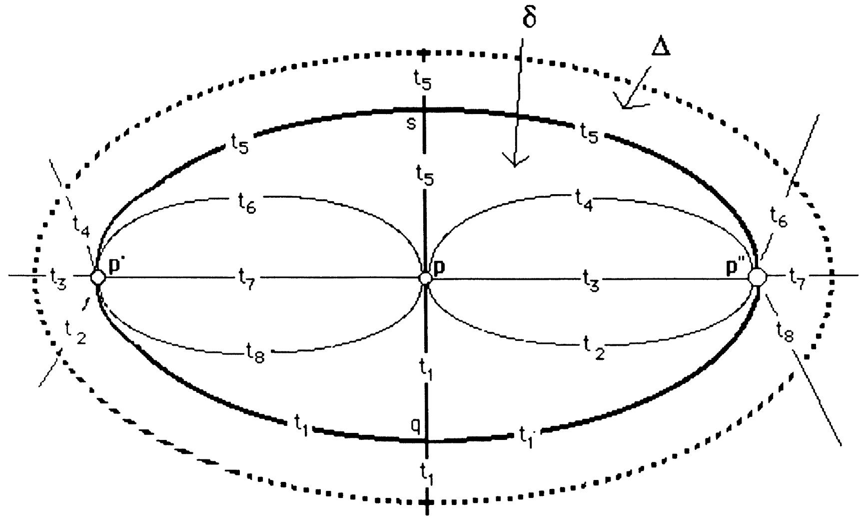

Fig. 13

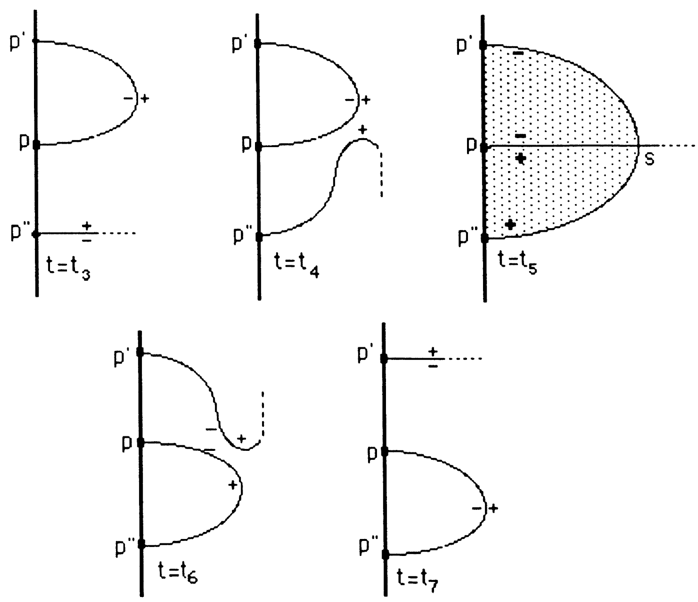

Fig. 14. Pictures in a cycle of fibers of **H**

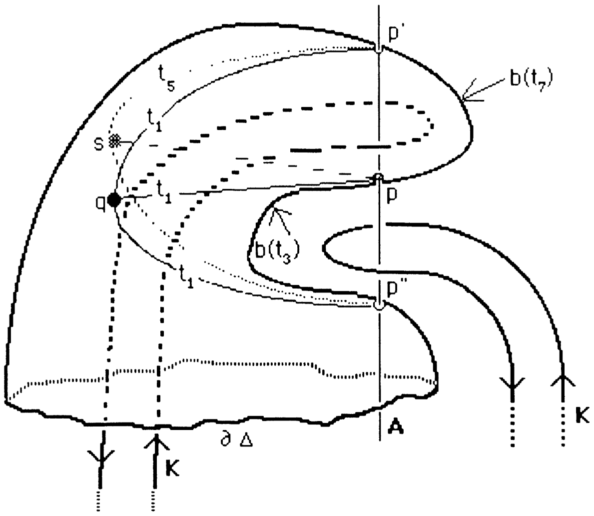

Fig. 15

points of $\mathbf{A} \cap \mathbf{X}$. This implies that $\partial\Delta$ has algebraic rotation number 1 about the axis **A**. Since $\partial\Delta$ is everywhere transverse to the foliation, it then follows that $\partial\Delta$ is a 1-braid, and also that the axis **A** pierces Δ from alternating sides at the 3 pierce-points, say the positive side, the negative side, and then the positive side again at p, p' and p''.

We now have enough information to determine how Δ is embedded in 3-space. See Fig. 15. The first step is to choose three points on the axis **A** and declare them to be p', p and p''. Up to a homeomorphism of S^3 which fixes **A** and each fiber it will not matter where we place them. Next, we attach small discs transverse to the axis with their centers at p', p and p''. These discs are the intersections of our oriented surface **X** with a neighborhood of the axis, and by the argument we have just given, if the disc at p' has its positive side up, then the discs at p and p'' will have negative and positive sides up, respectively.

The next step is to choose points on the fibers at t_1 and t_5 and declare them to be q and s. Up to a homeomorphism of S^3 which preserves **A** and each fiber it will not matter where we place them. Now, each singular point is the intersection of four singular leaves of the foliation, which go out from the singular point like the spokes on a wheel. Three of the leaves which meet at q and also at s terminate at p', p, p'', so we can extend them (in their fibers) to p', p and p''. The fact that the singularities have opposite signs shows that if the cyclic order (looking down onto the positive side of **X**) at q is p', p, p'', then it must be p'', p, p' at s. The fourth leaf at s and also at s' goes to the boundary of Δ, which (as noted earlier is a little 1-braid about **A**. Since the surface Δ is transverse to the foliation in the complement of the singular leaves, there is now a unique way to complete the embedding of Δ. In this way we obtain Fig. 15.

We study the non-singular leaf $b(t_7)$ in Figs. 12 and 13. By hypothesis, every leaf in the foliation is essential. This means that **K** itself is necessarily an obstruction to their removal, as illustrated in Fig. 15. We propose to remove this obstruction by

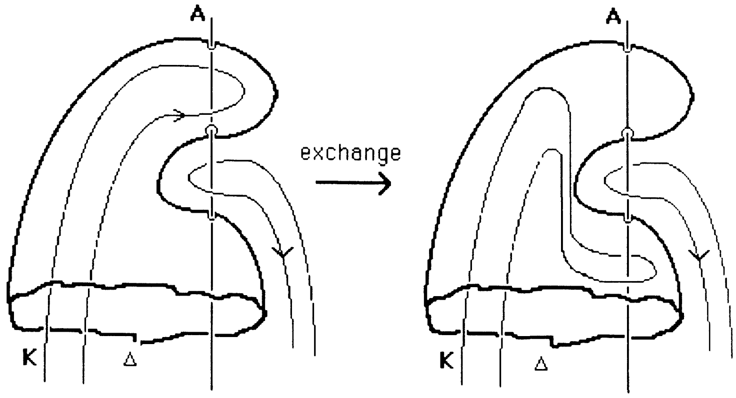

Fig. 16

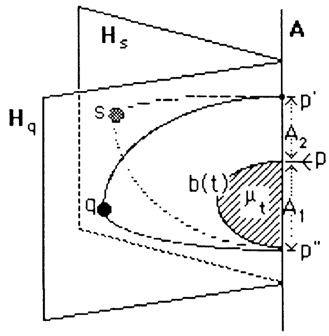

Fig. 17

an exchange, as illustrated in Fig. 16. After the exchange we will obtain a new tiling with the same complexity, however in the new tiling the leaf $b(t_7)$ will be inessential.

The remaining problem is to show that we can actually realize this exchange so that it takes braids to braids and leaves the tiling unchanged, except for reordering of the p_i's. If we can do this, then our exchange move will enable us to reduce the complexity. We need to create a region in 3-space in which we can make the exchange move in a controlled fashion.

Each non-singular $\mathbf{H}_t$ meets $\mathbf{X}$ in a unique arc $b(t)$ which joins p to p' or p'' and cobounds with a part of the axis $\mathbf{A}$ two discs in the fiber $\mathbf{H}_t$. If $b(t)$ joins p to p' (resp. p''), let μ_t be the disc which does not contain p'' (resp. p'). We have sketched in one such arc $b(t)$ and shaded in the disc μ_t in Fig. 17. If we think of the subsurface Δ of Fig. 15 as a boxing glove, the discs μ_t for $t_1 < t < t_5$ will appear to be "outside" the glove, whereas those for $t_5 < t < t_1$ will be "inside" the glove. The disc δ, which is on the glove, is a limiting position for both families of discs. Thus the closure of the union of all of the discs μ_t, $t \in [0, 2\pi]$ will be two 3-balls $\mathbf{B}_1$ and $\mathbf{B}_2$, which intersect along a single arc in the disc δ of Fig. 13, i.e. the arc which runs from q to s through p. The shaded disc in Fig. 14 is in the boundary of $\mathbf{B}_1 \cup \mathbf{B}_2$.

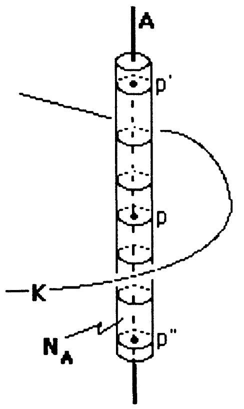

Fig. 18

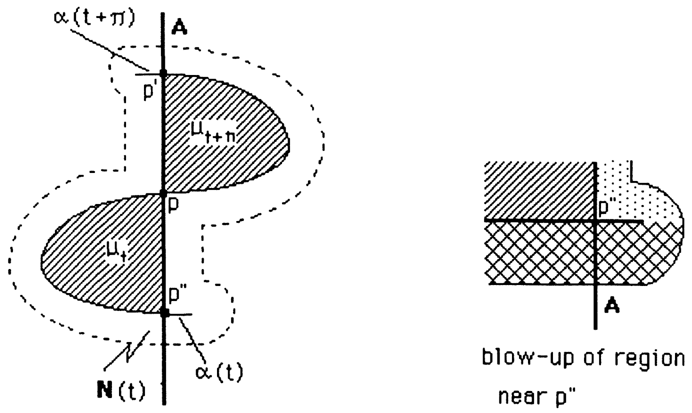

Fig. 19

We may assume that $\mathbf{B}_1 \cup \mathbf{B}_2$ intersects the link $\mathbf{K}$ in two unbraided weighted arcs. This is clearly possible, since all braiding may be pushed out of $\mathbf{B}_1 \cup \mathbf{B}_2$. Notice that the surface $\mathbf{X}$ may meet $\mathbf{B}_1$ and $\mathbf{B}_2$ in some number of sheets which are locally parallel to the embedded disc Δ which we depicted in Fig. 15. In particular, $\mathbf{X}$ may meet a neighborhood $\mathbf{N}_\mathbf{A}$ of $\mathbf{A}$ in some number of radially foliated discs between the discs at p' and p'', as in Fig. 18.

Now we need to thicken $\mathbf{B}_1 \cup \mathbf{B}_2$ a little bit. With this in mind, reparametrize the interval $[0, 2\pi]$ so that the singularities occur at $t_1 = 0$ and $t_5 = \pi$. We can then pair the discs μ_t and $\mu_{t+\pi}$, so that μ_t is in $\mathbf{B}_1$ and $\mu_{t+\pi}$ is in $\mathbf{B}_2$. See Fig. 19, which depicts subsets of $\mathbf{H}_t \cup \mathbf{H}_{t+\pi}$. Now notice (see Fig. 13) that if $\mathbf{H}_t$, $0 < t < \pi$, is non-singular, then $\mathbf{H}_t$ contains two leaves in the foliation of Δ: the leaf $b(t)$, which joins p to p'', and also a small arc $\alpha(t)$ which runs out from p'' to the boundary of Δ.

Similary, $\mu_{t+\pi}$ is a union of $b(t+\pi)$, which joins p to p', and a small arc $\alpha(t+\pi)$ which runs out to the boundary from p''. See Fig. 19. Let $\mathbf{N}(t)$ be a neighborhood of $\mu_t \cup \mu_{t+\pi} \cup \alpha(t) \cup \alpha(t+\pi)$ in $\mathbf{H}_t \cup \mathbf{H}_{t+\pi}$, chosen so that $\mathbf{N}(t) \cap \mathbf{K} = (\mu_t \cup \mu_{t+\pi}) \cap \mathbf{K}$ and $\mathbf{N}(t) \cap \mathbf{X}$ is the union of $(\mu_t \cup \mu_{t+\pi}) \cap \mathbf{X}$ and $(\mathbf{H}_t \cup \mathbf{H}_{t+\pi}) \cap \mathbf{N_A}$. Finally, choose the neighborhoods $\mathbf{N}(t)$ so that they vary smoothly as t is varied between 0 and π. Let $\mathbf{N}$ be the closure of the union of all of the $\mathbf{N}(t)$, $t \in [0, \pi)$. Notice that $\mathbf{N}$ is not a 3-ball (it is a 3-ball with holes), because $\mathbf{N}(t)$ is not a disc for t-values close to $t = 0$. Chosse $\varepsilon > \mathbf{0}$ so that $\mathbf{N}(t)$ *is* a disc for $t \in [\varepsilon, \pi - \varepsilon]$.

We can now describe precisely the exchange move which was indicated earlier in Fig. 16. See the top left picture in Fig. 20, which is intended to correspond to the left picture in Fig. 19, at some $t \in (\varepsilon, \pi - \varepsilon)$. The plane of the paper is divided by $\mathbf{A}$ into two half-planes, and we will assume the left one to be $\mathbf{H}_t$ and the right one to be $\mathbf{H}_{t+\pi}$. We want to describe an isotopic deformation of $\mathbf{X} \cup \mathbf{K}$, and shall do so by the series of pictures in Fig. 20, all of which correspond to the same fixed value of $t \in [\varepsilon, \pi - \varepsilon]$.

Choose a subdisc $\mathbf{d}$ of $\mu_{t+\pi}$ such that $\mathbf{d} \cap \mathbf{K} = (\mu_{t+\pi}) \cap \mathbf{K}$ and also $\mathbf{d} \cap \mathbf{X} = (\mu_{t+\pi}) \cap \mathbf{X} - b(\mathrm{t}+\pi)$. Our isotopy is to be supported in $\mathbf{N}(t)$ and is to be the identity on μ_t and on $\mu_{t+\pi}$ minus a neighborhood of $\mathbf{d}$ on $\mu_{t+\pi}$. The isotopy pushes the disc $\mathbf{d}$ (and the points of $\mathbf{K}$ and $\mathbf{X}$ which meet it) across $\mathbf{A}$ and then down and eventually into the cross hatched area below p'' which was illustrated in the blow-up in Fig. 19. The isotopy is to be defined on pairs of fibers, and is to be defined so that it varies continuously as we vary t. At the end of the isotopy the link $\mathbf{K}$ is to encircle the axis below p'' instead of between p' and p. Also, each sheet of $\mathbf{X}$ which intersected $\mathbf{A}$ between p' and p is to intersect $\mathbf{A}$ the same number of times below p''. (One of the authors likes to think of this change as accomplished by "putting your hand into your pocket and emptying it".)

The only remaining problem is to ask what happens in the t-interval $[0, \varepsilon]$? Let u be the isotopy parameter. The first thing to notice that when u is in the interval $[.25, .75]$ the isotopy is supported in the left half-plane. We may then assume that the deformed discs $\mathbf{d}_u$, $u \in [.25, .75]$, have empty intersection with $\mathbf{H}_t \cup \mathbf{H}_{t+\pi}$ if $t \notin [\varepsilon, \pi - \varepsilon]$. The second thing to notice is that when $u \in [0, .25]$ and $u \in [.75, 1]$ the isotopy is supported in a neighborhood of the axis $\mathbf{A}$. Since that neighborhood can be chosen to be disjoint from a neighborhood of the singular points (see Fig. 14) it follows that Fig. 20 actually tell us everything we need to describe the isotopy of $\mathbf{X} \cup \mathbf{K}$ completely.

At the end of the isotopy we may reposition $\mathbf{X}$ so that the tiling of $\mathbf{X}$ will be exactly as it was before the change. The only change will be in the order of the points of $\mathbf{A} \cap \mathbf{X}$ on $\mathbf{A}$. The interchange of order has an important consequence: one or more essential b-arcs in the foliation of $\mathbf{X}$ will have been changed to inessential b-arcs. We can therefore modify $\mathbf{X}$ by an additional isotopy to a new splitting 2-sphere $\mathbf{X}'$, reducing the complexity. The proof of Lemma 4 is complete. □

We now complete the proof of the split braid theorem. We begin with a splitting 2-sphere $\mathbf{X}$ which has complexity $c(\mathbf{X}, \mathbf{H}) \geqq (2, 0)$. If the complexity is (2, 0) we are done, so assume it is $> (2, 0)$. This implies that $|\mathbf{A} \cap \mathbf{X}| > 2$, because if it were 2, then by our general position hypotheses we would also have $|\mathbf{H} \cdot \mathbf{X}| = 0$. By Lemmas 2 and 3, we can then conclude that there is a vertex of valence 2 in the foliation of $\mathbf{X}$. By Lemma 4 we conclude that after an exchange we may replace $\mathbf{X}$ by a new splitting 2-sphere $\mathbf{X}'$ with smaller complexity. After finitely many such changes we will obtain a splitting 2-sphere $\mathbf{X}''$ which intersects $\mathbf{A}$ twice, and a split n-braid representative $\mathbf{K}'$ of our link $\boldsymbol{K}$. □

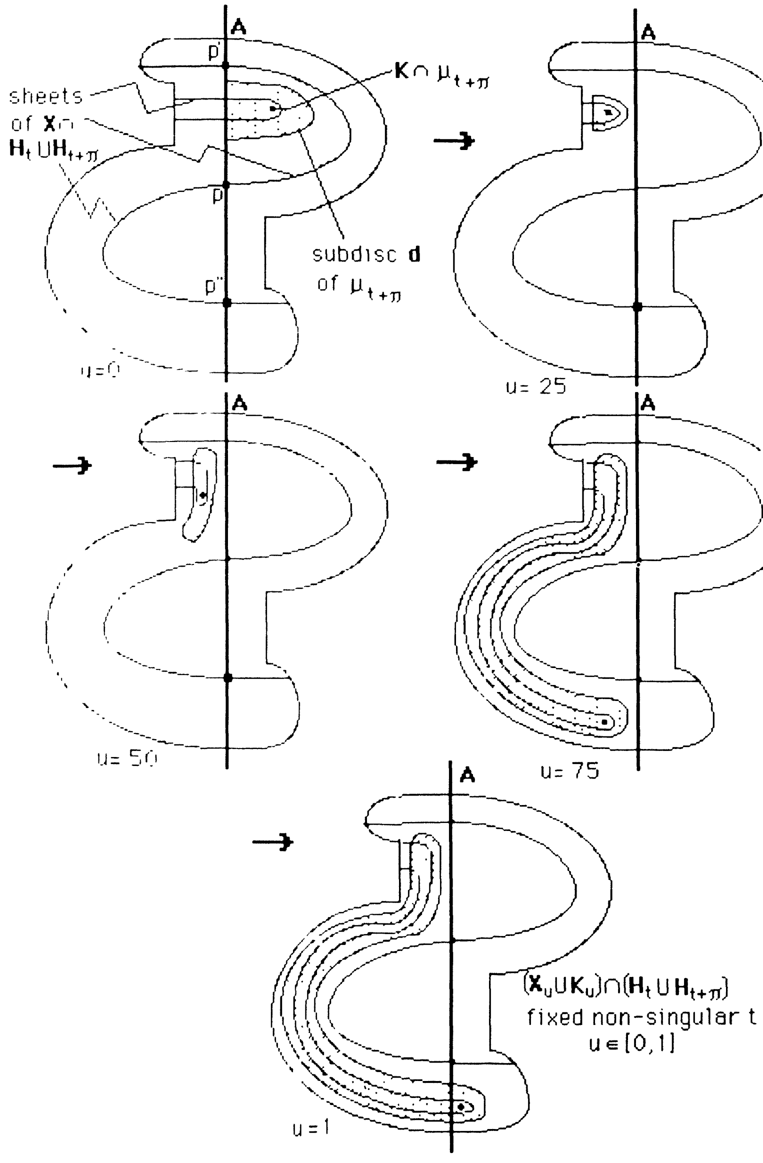

Fig. 20

3. Proof of the composite braid theorem

We begin with a closed n-braid **K** which represents a composite link $\boldsymbol{K}$. By the split braid theorem, we know that if $\boldsymbol{K}$ is a split link, we may find splitting 2-spheres which exhibit it as an obviously split closed braid. Thus we may work on one component at a time. Therefore we may assume that our representative **K** is non-split.

The proof is almost identical with that of the split braid theorem, therefore we will only discuss the points where they differ. We are given a 2-sphere **Y** (replacing **X**) which decomposes **K** into a connected sum of non-trivial link types. The only difference between the 2-sphere **X** of Sect. 2 and our 2-sphere **Y** is that **X** was disjoint from **K**, whereas **K** pierces **Y** twice. Our goal is to modify **K** and **Y** until we achieve a 2-sphere **Y**′ which also realizes the connected sum decomposition and which is pierced twice by the axis **A**.

The general position arguments given in Sect. 2 go through, but a complication occurs in the proof of Lemma 1, when we try to remove SCC's from the foliation of **Y**. The possibility exists that **K** will pierce the discs which are to be used in the surgeries. Suppose that c is a SCC in $\mathbf{H}_t \cap \mathbf{Y}$. Then c bounds a disc δ in $\mathbf{H}_t$, and also divides **Y** into two discs δ_1 and δ_2. Since **K** intersects **Y** twice, it either intersects δ_1 once and δ_2 once, or (by choosing the notation appropriately) it intersects δ_1 twice and misses δ_2. Assume the latter. Then if we surger **Y** along δ, as in the proof of lemma 2, **Y** will be split into two 2-spheres $\mathbf{Y}_1 = \delta \cup \delta_1$ and $\mathbf{Y}_2 = \delta \cup \delta_2$, and $\mathbf{Y}_2$ must be inessential because it misses **K**. Thus we can replace **Y** by $\mathbf{Y}_1$, and proceed as in the proof of Lemma 2. If, on the other hand, **K** meets both δ_1 and δ_2, then both 2-spheres $\mathbf{Y}_1$ and $\mathbf{Y}_2$ will intersect **K** twice, and at least one of them, say $\mathbf{Y}_1$, must be essential. We can then discard $\mathbf{Y}_2$, and proceed as in the proof of Lemma 2, Thus we may assume that **Y** is foliated without SCC's.

Our 2-sphere **Y** admits a tiling. Call a tile *good* if it is not pierced by **K**. Call a vertex of valence 2 or 3 *good* if it is adjacent to a good tile. As in the proof of the split braid theorem, we have to prove there is a vertex of valence 2 or 3 in the tiling, but now we need a little more: we must be sure that the vertex of valence 2 or 3 is good. The equation.

$$2V_2 + V_3 = 8 + V_5 + 2V_6 + 3V_7 + \dots.$$

still holds. In this equality every entry is non-negative. Now, we can assume (after a small isotopy of **K**) that the two points where **K** pierces **Y** are in the complement of the set of singular leaves, i.e. in a region **R** like the one depicted in Fig. 6. Call **R** a *bad* region if it is pierced by **K**. Now, a region **R** is adjacent to exactly two tile vertices, and since there are at most 2 bad regions this means that there are at most 4 bad vertices. Thus, if $V_2 + V_3 \geqq 5$, we can always find a good vertex, so assume $V_2 + V_3 \leqq 4$. This means that $V_2 \leqq 4$. However, $V_2 + (V_2 + V_3) \geqq 8$. Therefore the only case where we could fail to have a good vertex of valence 2 or 3 is in the very special situation when $2V_2 = 8$ and $V_3 = 0$. This means that **Y** is a union of exactly two tiles, also the complement of the set of singular leaves is a union of four regions, and two of these are pierced by **K**.

We can now proceed, as in the proof of Lemma 4 above, to construct **Y** as an embedded surface in S^3. Let p_1, p_2, p_3, p_4 be the 4 points where **A** pierces **Y**, in their natural cyclic order on **A**. Each of these vertices has valence 2, so there are 2 tiles and 2 singular points in the foliation, say s and q. See the top pictures in Fig. 21. Since all of our vertices are bad, we know that **K** must pierce a pair of opposite regions, say R_1 and R_3. By the first part of the proof of Lemma 4, we know that the signs of s and q are opposite, so the clockwise cyclic order of the tile vertices will be p_1, p_2, p_3, p_4 in, say, the tile which contains s and p_4, p_3, p_2, p_1 in the tile which contains q, when viewed by an upright observer on the positive side of the tile. The points s and q lie in distinct fibers, and the singular leaves will be a union of four arcs which go out from s(resp. q), in a single fiber, like the spokes in a wheel, ending at the four cyclically ordered vertices, on **A** in a neighborhood **U** of the singular leaves the 2-sphere **Y** will be transverse to each fiber, so there is a unique way to extend the embedding of **S** to an embedding of **U**. With **U** in place, we can attach the four non-singular regions R_1, R_2, R_3, R_4. The embedding of **Y** is then as in the bottom picture of Fig. 21.

By hypothesis each arc in the foliation of **Y** is essential, so our link **K** must be an obstruction to its removal. See Fig. 22. Thus **K** must encircle the axis $n \geqq 1$ times between p_1 and p_2 and also $m \geqq 1$ times between p_3 and p_4, inside **Y**. These strands will in general be braided; we have indicated the braiding as occurring inside two boxes, labeled A and B. Similarly, **K** must encircle the axis $j \geqq 1$ times between p_2 and p_3 and $k \geqq 1$ times between p_4 and p_1, outside **Y**, with braiding occurring in the boxes which are labeled C and D. In addition, one strand of **K** pierces **Y** in the region R_1 and another in the region R_3, to join up the part of **K** which is inside **Y** with the part which is outside **Y**, as illustrated. Thus **K** is a connected sum of a link which is represented as an $(n + m)$-braid and a link which is represented by a $(j + k)$-braid.

It is now clear that we can slide the strands of **K** which pierce **Y** into, say, Region 2. After such a slide we will have a good vertex of valence 2. A complexity-

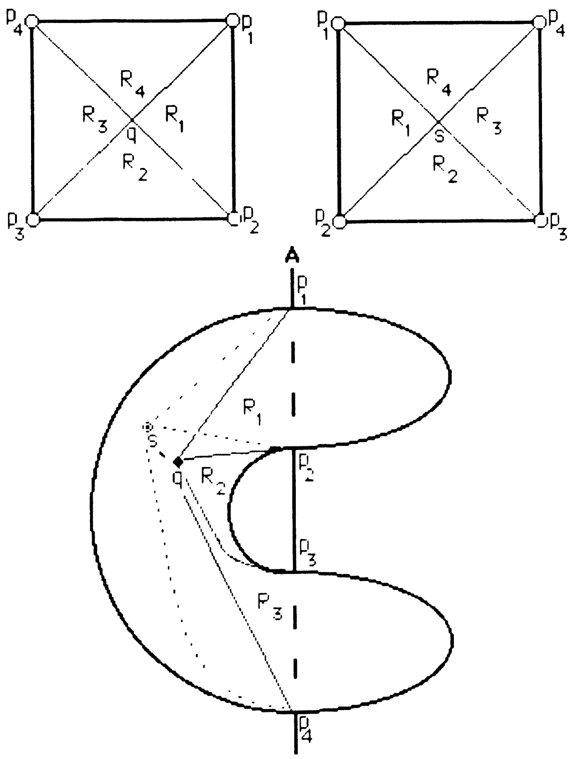

Fig. 21

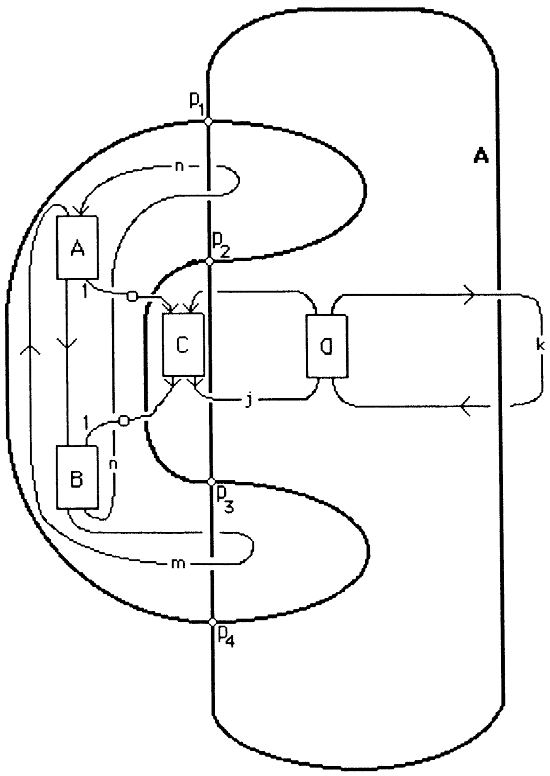

Fig. 22

reducing exchange move is applicable. The proof of the composite braid theorem is complete. □

4. Proof of the 3 and 4-braid theorem

In this section we prove special versions of the split and composite braid theorems which hold when the braid index is 3 or 4. See Sect. 1 above for the statement of the theorems.

Two of the four cases have been settled elsewhere, vis: In [Mu] Murasugi proved, in effect, that every closed 3-braid representative of a split link is conjugate to a split 3-braid representative. The analogous situation for composite links was settled by Morton in [Mo, 1]. A third case, that of split braids of braid index 4, has not been treated elsewhere, however the examples which we gave earlier in Fig. 1 show that there are split links of braid index 4 which are *not* conjugate to split closed braids. Thus the only case which requires proof is the case of composite braids of braid index 4.

Remark. The reader will have no difficulty in supplying new proofs for the three other cases, with the techniques which we will use to settle the fourth.

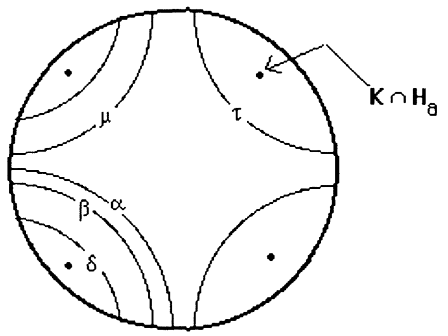

Fig. 23

We begin, as in the proof of the composite braid theorem, with a composite link and a 2-sphere **Y** which realizes the connected sum decomposition. We assume, as usual, that **Y** is in a nice position relative to **H**, so that **Y** has a tiling. Each tile has 4 vertices and 4 edges and 1 singularity in its foliation. The first step will be to examine the components of $\mathbf{H}_t \cap \mathbf{Y}$ as they occur in a sequence of non-singular fibers at times $t = a, b, c, \ldots$. See Fig. 23. Let α be a component of $\mathbf{H}_a \cap \mathbf{Y}$. Then α separates $\mathbf{H}_a$, and since α is essential (see Sect. 2) both components of $\mathbf{H}_a$ split along α must be pierced by **K**. On the other hand **K** is a closed 4-braid, so we know that **K** pierces $\mathbf{H}_a$ in exactly 4 points. This implies that α is parallel to one of the 4 arcs types illustrated in Fig. 23.

We assume that there are n_1, n_2, n_3, n_4 parallel arcs of each type, where each $n_j \geqq 0$. If we follow any one arc as we sweep through the fibration, we know it must ultimately be modified by a surgery. However, arcs in the position of β and δ in Fig. 24 cannot be surgered with one-another because such a surgery would create an inessential arc. Similarly, arcs like α and μ cannot be surgered with one-another, because that too would create an inessential arc. Thus the only possibility is that arcs in the position of α and τ are surgered with one-another. This has the effect of removing α and τ from a pair of opposite groups and adding new arcs α' and τ' to the other pair. This process must continue until there are no more arcs left in one of the pairs, at which point the process must be reversed, because every arc must both be modified by a surgery and also ultimately go back to its initial position. However, this means that eventually α' will be surgered again with τ'. Thus, in fact, if there were more than two arcs our surface **Y** would be disconnected. Hence there are exactly two arcs in $\mathbf{H}_a \cap \mathbf{Y}$, and two singular fibers.

The tiling of **Y** thus contains exactly two tiles, each with four vertices of valence two. This is precisely the situation which was described in the proof of the Composite Braid Theorem in Sect. 3, so we can conclude that the picture is exactly that in Fig. 22, with $n = m = j = k = 1$. Note that this implies immediately that our link $\boldsymbol{K}$ is actually the connected sum of *three*, not just two links, each of which can then only be a type $(2, r_i)$ torus link, for some $r_i \geqq 1$. Thus we have proved that if $\boldsymbol{K}$ is the connected sum of a link of braid index 3 and a link of braid index 2, then **K** is in fact conjugate to a composite 4-braid.

We go back to the case where $\boldsymbol{K}$ has three factors, and construct the 2-sphere which realizes the second connected sum decomposition of $\boldsymbol{K} = \boldsymbol{K}_1 \# \boldsymbol{K}_2 \# \boldsymbol{K}_3$.

Referring momentarily to Fig. 21, let c_1 be an arc on $\mathbf{Y}$ which joins the points s and q in the interior of R_1, and let c_2 be an arc which joins them in the interior of R_3. Then $c = c_1 \cup c_2$ is a simple closed curve which divides $\mathbf{Y}$ into two discs, a disc $\mathbf{Y}_2$ which contains R_2 and a disc $\mathbf{Y}_4$ which contains R_4. (Thus p_2 and p_3 are in $\mathbf{Y}_2$, while p_1 and p_4 are in $\mathbf{Y}_4$). We assume that we chose c_1 and c_2 so that the two points where $\mathbf{K}$ pierces $\mathbf{Y}$ are in $\mathbf{Y}_2$. Then $\mathbf{Y}_4$ is a disc with boundary c which is disjoint from $\mathbf{K}$, and is pierced twice by $\mathbf{A}$, at p_1 and at p_4. Keeping the curve c in mind, we now pass to Fig. 22, where we see that there is another disc, $\mathbf{Z}$ which also has boundary c, which is chosen so that the interior of $\mathbf{Z}$ lies "outside" the 3-ball bounded by $\mathbf{Y}$. In fact, we can choose $\mathbf{Z}$ so that it intersects $\mathbf{K}$ in two points, along the strands which join the blocks labeled C and D, and so that $\mathbf{Z} \cap \mathbf{A}$ is empty. Let $\mathbf{Y}' = \mathbf{Y}_4 \cup \mathbf{Z}$. Then $\mathbf{Y}'$ is also a 2-sphere which realizes the connected sum decomposition of $\mathbf{K}$, and $\mathbf{Y}' \cap \mathbf{A} = 2$ points. Thus, except for the very special case when the block D consist of a single crossing, we see that $\mathbf{K}$ is an obviously composite closed braid.

In the case where D contains a single crossing, our closed braid $\mathbf{K}$ has a trivial loop. That is, $\boldsymbol{K}$ is a connected sum of type $(2, r_1)$ and type $(2, r_2)$ torus links, $r_i \geqq 1$, and it is represented as a closed 4-braid. In this case there actually may be a non-composite 4-braid which represents $\boldsymbol{K}$. The proof is complete. □

References

[Be] Bennequin, D.: Entrelacements et equations de Pfaff. Asterisque **107–108**, 87–161 (1983)

[Bi] Birman, J.S.: Braids, links and mapping class groups. Ann. Math. Stud. **82** (1974)

[B-M, I] Birman, J.S., Menasco, W.W.: Studying links via closed braids I: A finiteness theorem. Preprint

[B-M, II] Birman, J.S., Menasco, W.W.: Studying links via closed braids II: On a theorem of Bennequin, Topology and its Applications (to appear)

[B-M, III] Birman, J.S., Menasco, W.W.: Studying links via closed braids III: Classifying links which are closed 3-braids. Preprint

[B-M, V] Birman, J.S., Menasco, W.W.: Studying links via closed braids V: The unlink. Trans. AMS.

[G] Garside, F.: The braid groups and other groups. Q. J. Math. Oxford **20** (No. 78), 235–254

[Ma] Markov, A.A.: Uber die freie aquivalenz der geschlossenen Zopfe. Rec. Soc. Math. Moscou **43**, 73–78 (1936)

[Mo, 1] Morton, H.R.: Closed braids which are not prime knots. Math. Proc. Camb. Philos. Soc. **86**, 421–426 (1979)

[Mo, 2] Morton, H.R.: Threading knot diagrams. Math. Proc. Camb. Philos. Soc. **99**, 247–260 (1986)

[Mo, 3] Morton, H.R.: Infinitely many fibered knots with the same Alexander polynomial. Topology **17**, 101–104 (1978)

[Mu] Murasugi, K.: On closed 3-braids. Mem. AMS **151** (1976)