Available online at www.sciencedirect.com
ScienceDirect
Topology and its Applications 154 (2007) 2962–2975
Topology and its Applications
www.elsevier.com/locate/topol

Modifying a branched surface to carry a foliation

Sue Goodman [a], Sandi Shields [b,*]

[a] *University of North Carolina at Chapel Hill, USA*
[b] *College of Charleston, 66 George street, Charleston, SC 29424, USA*

Received 6 July 2006; accepted 15 June 2007

Abstract

We consider C^1 nonsingular flows on a closed 3-manifold under which there is no transverse disk that flows continuously back into its own interior. We provide an algorithm for modifying any branched surface transverse to such a flow ϕ that terminates in a branched surface carrying a foliation F precisely when F is transverse to ϕ. As a corollary, we find branched surfaces that do not carry foliations but that lift to ones that do.
Published by Elsevier B.V.

MSC: 57R30; 57M50; 57M10; 57M20; 57N10

Keywords: Branched surface; Foliation; Transverse flow

Introduction

One of the open problems in foliation theory is to determine whether a nonsingular vector field on an arbitrary closed 3-manifold has a transverse foliation. Classical results by Fried [4] and Schwartzman [11] give conditions for any given flow to have a transverse section, hence a transverse foliation. Milnor [8] and Wood [14] showed that for circle bundles over a surface of genus at least one, there is a foliation transverse to the fibers precisely when the Euler number of the bundle is no larger than the negative of the Euler characteristic of the surface. This illustrates the subtlety of the problem. For example, one can have a circle bundle of sufficiently small Euler number finitely covering one with a large Euler number.

Naimi found necessary and sufficient conditions for the existence of a foliation transverse to the foliation by circles of a Seifert fibered 3-manifold [9]. Goodman showed that for a large (C^0 dense) class of flows, a simple linking property is both necessary and sufficient for the existence of a transverse foliation [7]. However, given the above example, any property of a flow which is preserved under finite covers cannot, for general flows, be both necessary and sufficient for the existence of a transverse foliation.

Here we consider C^1 nonsingular flows on a closed 3-manifold for which there is no transverse disk that flows continuously into its own interior. We provide an algorithm for modifying any branched surface transverse to such a flow ϕ that terminates in a branched surface carrying a foliation F precisely when F is transverse to ϕ. Our algorithm

* Corresponding author.
E-mail address: shieldss@cofc.edu (S. Shields).

0166-8641/$ – see front matter Published by Elsevier B.V.
doi:10.1016/j.topol.2007.06.014

differs from the Agol–Li algorithm [1] for determining if a 3-manifold admits an essential lamination, and it has the advantage that any resulting foliation is transverse to the given flow ϕ. We know of no dual algorithm, as exists for Li and Agol's, that can be run in tandem and terminates when ϕ is not transverse to a foliation. However, even in cases where our algorithm does not terminate in one that carries a foliation, it sometimes yields an interesting branched surface. In particular, the algorithm can be applied equivariantly to a branched surface and any of its lifts to a finite cover. In other words, the branched surface we get at each stage of the algorithm below could also be obtained by applying the algorithm to a lifted branched surface and projecting down. So in the case that ϕ has no transverse foliation yet lifts to a flow that does, our procedure for modifying W transverse to ϕ eventually yields a branched surface that does not carry a foliation but that lifts to one that does.

In Gabai's well-known problem list [5], he posed the question: Are there are useful branched surfaces which carry nothing? This question was answered affirmatively by Calegari [2] who showed that there exist branched surfaces contained in circle bundles over a surface which do not carry a lamination but that have finite covers that do. It follows that any property of a branched surface W which is preserved under finite covers cannot be an obstruction to W carrying a foliation. As described above, we use our algorithm to obtain an alternative proof of this result.

1. Preliminaries

Throughout this paper, F will be a C^1 codimension one foliation of a closed orientable 3-manifold M and $\phi : M X \Re \to M$ will be a C^1 nonsingular flow on M that is transverse to F.

Formally, a curve in M is a continuous map from a connected subset of $\Re$ into M. However, we shall consider a *curve* as the image of such a map, where the map parameterizes the curve. If a curve has a negative (positive) boundary point, according to the orientation induced by the parameterization, we refer to this point as the *initial point* (*terminal point, respectively*) of the curve. An *integral curve of F* is a curve contained in a leaf of F.

An *orbit segment* of ϕ will be a curve $\phi(x,t)_{t\in[a,b]}$, where $x \in M$ and $[a,b]$ is a closed interval in R. The *forward* (*backward*) *orbit* of a point $x = \phi(x,0)$ in M under ϕ will be the set of points $\phi(x,t)_{t>0}$ ($\phi(x,t)_{t<0}$, respectively).

We say the flow ϕ is *Reebless* if there is no *self-return disk for* ϕ, i.e. there is no disk transverse to orbits of ϕ that flows strictly inside itself in future or past time.

1.1. Branched surface construction

The branched surfaces we associate with a nonsingular flow ϕ are in the class of regular branched surfaces introduced by Williams [13]. Since the construction we use is in an unpublished paper of Christy and Goodman [3] and is a variation of the one in [6], we describe it here, including all details necessary for this article. This description is very similar to the one in [12]. However, it allows for the more general situation where there is no foliation transverse to ϕ and establishes notation that we shall used throughout.

We begin with a nonsingular flow ϕ transverse to F and a finite *generating set for* ϕ, $\Delta = \{D_i\}_{i=1,\dots,n}$, consisting of pairwise disjoint embedded disks satisfying the following general position requirements:

(i) each D_i is transverse to ϕ,
(ii) the forward and backward orbit of every point, under ϕ, meets $\operatorname{int}\Delta = \bigcup_{i=1}^n \operatorname{int} D_i$,
(iii) the set of points in $\partial\Delta = \bigcup_{i=1}^n \partial D_i$ whose orbit, forward or backward, meets $\partial\Delta$ before meeting $\operatorname{int}\Delta$ is finite, and
(iv) the forward orbit of any point in $\partial\Delta$ meets $\partial\Delta$ at most once before meeting $\operatorname{int}\Delta$.

Note that we can always find a generating set for ϕ. In particular, cover M with flow boxes for ϕ, and select a transverse slice from each box. Then, modify each slice slightly so that the resulting collection of disks satisfies the general position requirements above.

After choosing a generating set Δ, we cut M open along the interior of each element of Δ to obtain a closed submanifold M^* which is embedded in M so that its boundary contains $\partial\Delta$. This can be thought of as blowing air into M to create an air pocket at each generating disk. By requirement (ii) above, the restriction of ϕ to M^* is a flow ϕ^* with the property that each orbit is homeomorphic to the unit interval [0, 1]. We then form a quotient space by identifying points that lie on the same orbit of ϕ^*. That is, we take the quotient $M^*/\sim$, where $x \sim y$ if x and y lie on

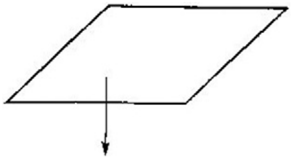

Fig. 1.

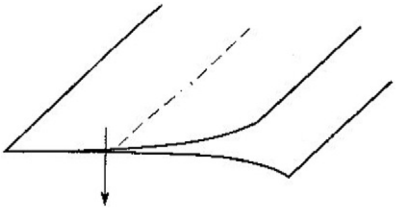

Fig. 2.

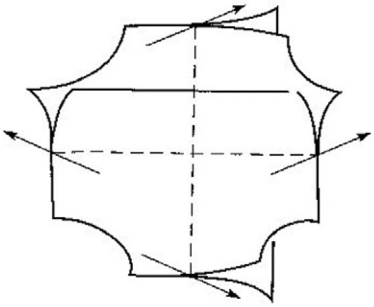

Fig. 3.

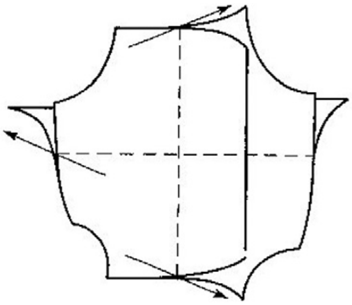

Fig. 4.

the same interval orbit of ϕ^*. This quotient space can be embedded in M so that it is transverse to ϕ. Specifically, we can view the quotient map as enlarging the components of $M - M^*$ until each interval orbit of ϕ^* is contracted to a point in M. We refer to such an embedded copy of the quotient space as *the branched surface W constructed from ϕ*. (Although there are many embeddings of the quotient $M^*/\sim$ that are transverse to ϕ, the complement of each is a union of open 3-balls. So any two embeddings of $M^*/\sim$ are *diffeomorphic in M*; that is, there is a diffeomorphism of M that maps one onto the other. Consequently, we only distinguish between branched surfaces transverse to ϕ up to diffeomorphism of M.)

Note that the branched surface W could have many generating sets. For example, if we flow a disk in Δ forward or backward slightly to another sufficiently close disk, the quotient space described above does not change.

The general position requirements for a generating set imply that the branched surface W is a compact connected 2-dimensional complex with a set of charts defining local orientation preserving diffeomorphisms onto one of the models in Figs. 1–4 such that the transition maps are smooth and preserve the transverse orientation indicated by the arrows. (Each local model projects horizontally onto a vertical model of R^2, so has a smooth structure induced by TR^2 when we pull back the local projection.) In particular, W is a connected 2-manifold except on a dimension one subset μ called the *branch set*. The set μ is a 1-manifold except at finitely many isolated points where it intersects itself transversely. The components of $W - \mu$ are the *sectors* of W.

We can thicken the branched surface W in the transverse direction to recover M^* which, for this reason, we shall henceforth call $N_\Delta(W)$, *the neighborhood* of W. In particular, $N_\Delta(W)$ is obtained when we replace each point x in W with the interval orbit of ϕ^* whose quotient is x. (For simplicity, we shall use $N(W)$ to represent this submanifold when the generating set is not relevant to the discussion.)

Throughout, $\pi_W : N(W) \to W$ will denote the quotient map that identifies points in the same orbit of ϕ^*. We say the image x of a point under this map is the *projection* of that point. In particular, the interval orbit of ϕ^* that projects onto x will be called the *fiber of $N(W)$ over x*. (See Fig. 5.) We shall always assume that points in the same fiber of $N(W)$ are ordered according to the orientation of ϕ^*.

1.2. Foliations carried by a branched surface

If a foliation F is transverse to ϕ and if each element of Δ is contained in a leaf of F, then *F is carried by W*. In particular, the foliation F becomes a foliation of $N_\Delta(W)$ whose leaves (some of which are branched) are transverse to the fibers, when we cut M open along Δ. The branched leaves are precisely those that contain a boundary component

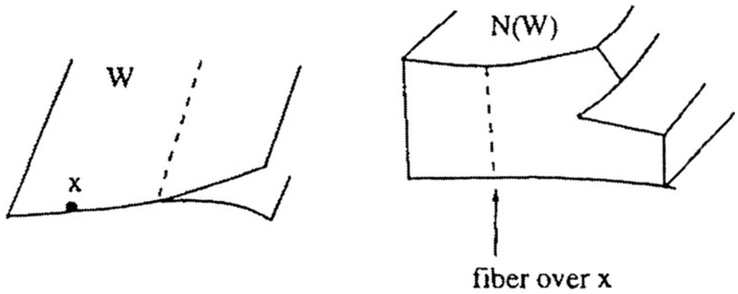

Fig. 5.

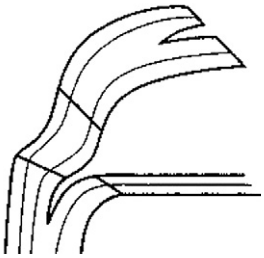

Fig. 6.

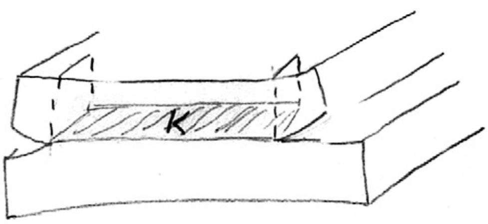

Fig. 7.

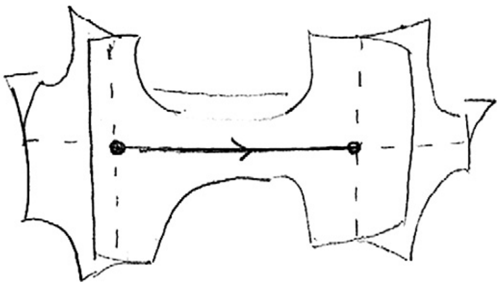

Fig. 8.

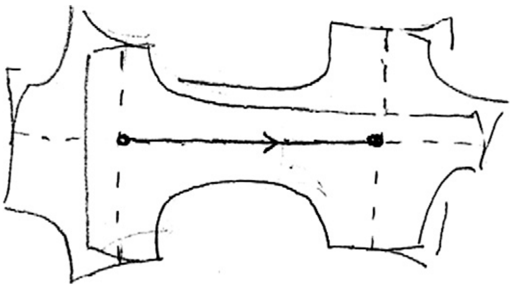

Fig. 9.

of $N_\Delta(W)$ since these are the (cut-open) leaves of F containing the elements of Δ. (They can be thought of as leaves of F with air blown into them.) Fig. 6 shows a local picture of this foliation of $N_\Delta(W)$.

In fact, each foliation of $N_\Delta(W)$ that is transverse to the fibers and whose leaves contain the boundary components of $N_\Delta(W)$ corresponds to a foliation of M that is carried by W. Specifically, when we collapse the components of $M - N_\Delta(W)$ (i.e., the air pockets), each of these foliations of $N_\Delta(W)$ yields a foliation of M that is also transverse to ϕ and whose leaves contain the elements of Δ.

1.3. Connecting curves

We say the parameterized image K of an immersion $i:[a,b] \times [-\varepsilon, \varepsilon] \to N_\Delta(W)$, $a, b, \varepsilon \in \Re$, is a *connecting strip* if its interior is transverse to the fibers, and only its ends, $i(\{a\} \times [-\varepsilon, \varepsilon])$ and $i(\{b\} \times [-\varepsilon, \varepsilon])$, are contained in $\partial N_\Delta(W) \cap \partial \Delta$. See Fig. 7.

The projected image of the curve $i(\{t\} \times \{0\})_{a \leqslant t \leqslant b}$ in W is a curve $\kappa(t)_{a \leqslant t \leqslant b}$ with two sectors branching into its initial point $\kappa(a)$ and two sectors branching out from its terminal point $\kappa(b)$. We say that κ is a *connecting curve corresponding to* K. (Since, in addition to i, there are many other immersions of $[a,b] \times [-\varepsilon, \varepsilon]$ in $N(W)$ whose image is K and whose ends are contained in $\partial N(W) \cap \partial \Delta$, there are many connecting curves corresponding to K. Likewise, there are many connecting strips corresponding to the same connecting curve.)

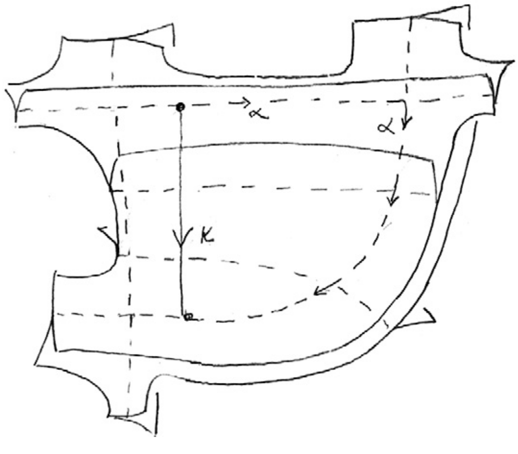

Fig. 10.

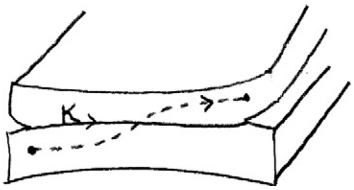

Fig. 11.

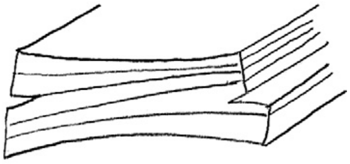

Fig. 12.

By definition, a connecting curve must correspond to a connecting strip. So not all curves in W with sectors branching from both ends are necessarily connecting curves. For example, the curve shown in Fig. 8 is a connecting curve. However, the curve shown in Fig. 9 is not, since there is no connecting strip in $N(W)$ lying over it.

A connecting curve κ is *trivial* if contains no loops and if it is fixed point homotopic, in W, to a curve α in the branch set μ along which all branchings are in the same direction (see Fig. 10); the curve α is then said to be a *μ-companion* of κ. In this case, κ^{-1} is also a trivial connecting curve and α^{-1} is one of its μ-companions.

Let $\beta = \beta(t)_{a \leqslant t \leqslant b}$ be a curve in $N(W)$ that projects onto a connecting curve $\kappa(t)_{a \leqslant t \leqslant b}$ corresponding to K; i.e., $\pi_W(\beta(t)) = \kappa(t)$ for all $t \in [a, b]$. If $\beta(a)$ lies below the point of intersection of $i(\{a\} \times [-\varepsilon, \varepsilon])$ with the fiber over $\kappa(a)$, and if $\beta(b)$ lies above the point of intersection of $i(\{b\} \times [-\varepsilon, \varepsilon])$ with the fiber over $\kappa(b)$, then β *crosses K with index* $+1$. See Fig. 11. If, instead, the curve β^{-1} crosses K^{-1} with index $+1$, then β *crosses K with index* -1. If both $\beta(a)$ and $\beta(b)$ are contained in ∂K, then β *crosses K with index* 0. Note that if F is a foliation of $N(W)$, all integral curves of F crossing K will do so with the same index. So we say a *foliation F crosses K* with a particular index if it contains an integral curve that crosses K with that index. For example, the foliation in Fig. 12 crosses the connecting strip K in Fig. 11 with index $+1$. Note that it is possible that no integral curve of F projects onto a connecting curve corresponding to K.

1.4. Staircase curves

Given a generating set Δ for a nonsingular flow ϕ, let $\gamma = \tau_1 * \sigma_2 * \cdots * \tau_{k-1} * \sigma_k * \tau_k$ be a compact curve in M, where τ_1 has nonempty interior and for any $i \geqslant 1$, τ_i is a positively oriented orbit segment of ϕ. If we can choose this decomposition of γ so that each σ_i has nonempty interior and is nowhere tangent to ϕ, then we say γ is a *staircase curve in ϕ*. See Fig. 13. (Specifically, if such a γ begins and ends at the same point, then it is a *staircase loop in ϕ*.) We refer to the σ_i's as the *steps* of γ and, given a fixed metric on M, define the *horizontal length* of γ to be the sum of the lengths of its steps. If each step σ_i has length at most ε and each τ_i has length at least T, then we say γ is an *(ε, T)-staircase curve*. (The idea is that if ε is large relative to T, then these staircase curves look almost like orbits of the flow ϕ.) If each step is contained in an element of some generating set Δ for ϕ, then we say γ is a *staircase curve in (Δ, ϕ)*. Likewise, if each step is contained in a leaf of some foliation F, then γ is a *staircase curve in (F, ϕ)*. (Note

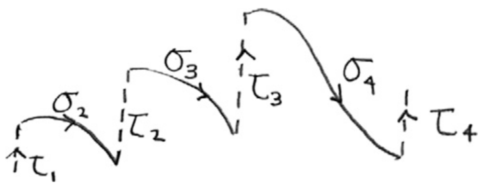

Fig. 13.

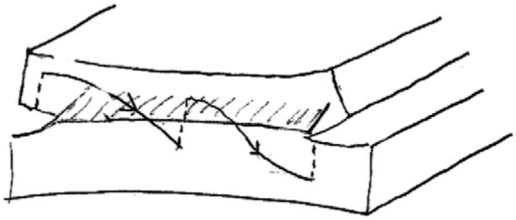

Fig. 14.

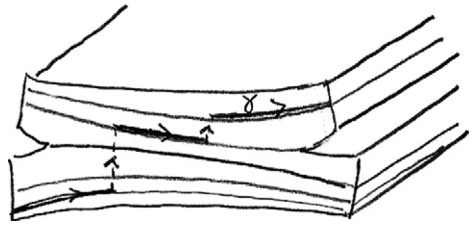

Fig. 15.

that any orbit of ϕ that is transverse to a foliation F is a $(0, T)$-staircase curve in (F, ϕ) with no steps, where T is its length.)

In Section 2, we shall want to consider staircase curves in ϕ that are contained in the interior of $N(V)$, for some branched surface V constructed from ϕ. It is important to note that if there is a staircase curve γ in (F, ϕ) that crosses a connecting strip K with index -1, then F either crosses K with index -1 or not at all. For example, if the stairs of the curve γ shown in Fig. 14 are contained in leaves of a foliation F, then γ modifies to a positively oriented transversal to F, so F cannot cross K with index $+1$ or 0. The situation is *not* analogous if γ crosses K with index $+1$. In fact, whenever a foliation F crosses K with index -1 or 0, there is a staircase curve γ in (F, ϕ) that crosses K with index $+1$. See Fig. 15.

2. Main results

In this section, we show that if a nonsingular Reebless flow ϕ is transverse to a foliation F, then we can break any generating set Δ for ϕ into smaller and smaller disks so that eventually we get a generating set for a branched surface that carries F (Theorem 2.3). We then use this to show that there are branched surfaces that do not carry foliations but have finite covers that do (Corollary 2.4). For example, one can reduce the Euler number of a circle bundle via a finite cover by unwinding in the fiber direction, as described by Calegari. Hence, one can produce a flow ϕ_2 with no transverse foliation, yet covered by a flow ϕ_1 which has one [8,14,2]. In this case, our procedure for modifying a generating set for ϕ_2 results in a branched surface transverse to ϕ_2 that does not carry a foliation, yet lifts to a branched surface transverse to ϕ_1 which does. It follows that any property of a branched surface W which is preserved under finite covers cannot be an obstruction to W carrying a foliation.

To prove Theorem 2.3, we shall modify a generating set Δ for ϕ until each of its elements slides injectively, along orbit segments of ϕ, into a leaf of the foliation F. This is possible provided that the modification makes all the disks in Δ sufficiently small. However, when we slide these disks into leaves of F, we also wish to preserve their relative positions (in the ϕ-direction) so that they still generate the same branched surface. Since it is not always possible to do so when there are staircase loops in (Δ, ϕ) (for example, if all leaves of F are compact), we first modify Δ to eliminate all such loops. Specifically, we describe a way to break the elements of Δ so that all staircase loops in (Δ, ϕ) are destroyed.

Our modification of Δ will also ensure that the remaining staircase curves in (Δ, ϕ) have a very small horizontal length. So when we embed Δ in a fiber neighborhood $N(V)$ of some branched surface V carrying F, the only connecting curves crossed by these staircase curves are trivial. From this we argue that if some connecting strip is crossed with index -1 by a staircase curve γ in (Δ, ϕ) and that connecting strip is also crossed with a different index by F, then γ is part of a null homotopic (ε, T)-staircase loop, where ε is very small relative to T. In Lemma 2.1, we

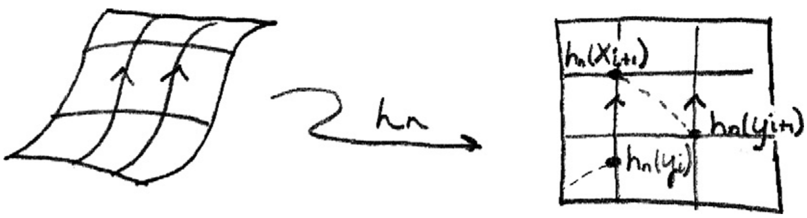

Fig. 16.

show that such a loop necessarily bounds a disk containing a self-return disk for ϕ, contradicting our assumption that ϕ is Reebless. So no connecting strip in $N(V)$, that is crossed by a staircase curve in (Δ, ϕ) with index -1 is crossed with a different index by F. In this case, the elements of Δ slide, along fibers in $N(V)$, into leaves of F as desired. (This is shown in Lemma 2.2.)

Lemma 2.1. *Suppose a nonsingular flow ϕ is transverse to a foliation. There exists a $J > 0$ and $R > 0$ such that if $\varepsilon < J$ and $T/\varepsilon > R$, then every null homotopic (ε, T)-staircase loop in ϕ bounds a disk containing a self-return disk for ϕ.*

Proof. Given a nonsingular flow ϕ transverse to a foliation F, we show that when ε is sufficiently small and T/ε is sufficiently large, any null-homotopic (ε, T)-staircase loop in ϕ modifies to a null homotopic loop τ, transverse to F; hence by Novikov's theorem [10] τ bounds a disk containing a self-return disk for ϕ.

Cover the manifold M by boxes which are both flow boxes for ϕ and foliation boxes for F and take a finite subcover $\{U_n\}_{n=1,\dots,N}$. Let λ be a Lebesque number for this cover. For each $n \leqslant N$, there is a homeomorphism $h_n : U_n \to I^3$ where $I = (0, 1)$ and the local orbits of $h_n(\phi)$ are $I^2 \times \{i\}$. There exists a constant K such that $\frac{1}{K}(d(h_n(x), h_n(y)) \leqslant d(x, y) \leqslant K(d(h_n(x), h_n(y))$, for all $n \leqslant N$ and all $x, y \in M$.

Let γ be an (ε, T)-staircase loop where $\varepsilon < \lambda$ and $T/\varepsilon > K^2$. Since the length of every step in γ is less than the Lebesgue number for the covering, each is contained in a foliation box. In addition, $T/\varepsilon > K^2$ ensures that, as we travel along γ, we never jump backward along some step more than we flow forward along ϕ. More precisely, let $\gamma = \tau_1 * \sigma_2 * \cdots * \tau_{k-1} * \sigma_k * \tau_k$, and for each $i \leqslant k$, let x_i and y_i be the initial and terminal points, respectively, of σ_i. Then $Kd(h_n(y_i), h_n(x_{i+1})) \geqslant d(y_i, x_{i+1}) \geqslant T > K^2\varepsilon$, while $d(h_n(x_i), h_n(y_i)) \leqslant Kd((x_i), (y_i)) \leqslant K\varepsilon$. Therefore, we can construct a positively oriented (with respect to the transverse orientation induced by ϕ) transversal to F from y_i to y_{i+1} that is homotopic to $\tau_i * \sigma_{i+1}$ and as close to $\tau_i * \sigma_{i+1}$ as we like. In this manner, we can modify γ to a null homotopic loop τ that is transverse to F. $\square$

In the following lemma, we show that if a branched surface V carries a foliation F and if a generating set Δ for a branched surface W is embedded in the interior of $N(V)$ so that it is transverse to the fibers, then there are only two obstructions to W carrying F. We identify these obstructions and later, in the proof of Theorem 2.3, show that they can be avoided by breaking Δ into small enough pieces in a prescribed manner.

Lemma 2.2. *Let V be a branched surface constructed from a nonsingular flow ϕ that carries a foliation F, and let $\Delta = \{D\}_{i=1,\dots,n}$ be a generating set for ϕ that is contained in the interior of $N(V)$. If there are no staircase loops in (Δ, ϕ) contained in $N(V)$, and if no connecting strip in $N(V)$ that is crossed with index -1 by a staircase curve in (Δ, ϕ) is crossed with index 0 or $+1$ by F, then F is also carried by the branched surface W generated by Δ.*

Proof. Let V, Δ and W be as in the hypotheses and let X be a generating set for V contained in F such that $X \cap \Delta = \emptyset$.

Given two elements D_i and D_j of Δ, let $D_i < D_j$ if there exists a staircase curve in (Δ, ϕ) contained in $N_X(V)$ that begins in D_i and ends in D_j. If $D_i < D_j$ and $D_j < D_i$ for some i and j, we have a staircase loop in (Δ, ϕ) contained in $N_X(V)$, which contradicts our assumption. So the relation $\leqslant$, where $D_i = D_j$ precisely when $i = j$, is a partial ordering of the elements of Δ.

It also follows that no disk in Δ can intersect a fiber of $N_X(V)$ in more than one point, so each projects injectively onto a disk in V. In fact, for each disk $D \in \Delta$ there exists a closed maximal stack Σ_D of disks contained in leaves of F that project onto the same subset of V as does D. We shall show that we can assign, to each disk $D \in \Delta$,

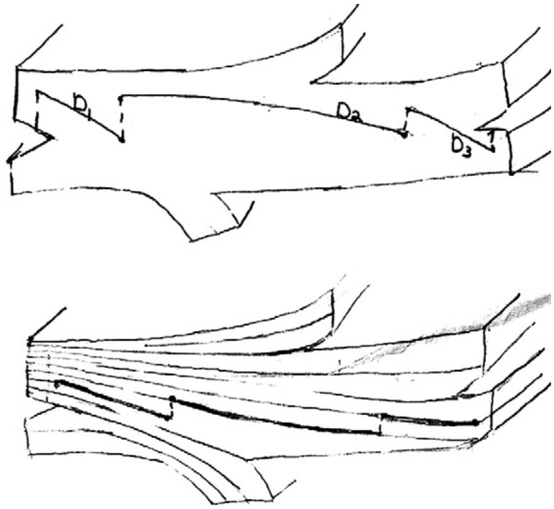

Fig. 17.

a corresponding disk in Σ_D in such a way that the partial ordering of Δ is preserved by the correspondence. See Fig. 17. In this manner, we find a collection of disks in leaves of F that generates the same branched surface W as does Δ and, consequently, W carries F.

To begin, let γ be a *maximal* staircase curve in (Δ, ϕ) contained in $N_X(V)$; that is, there is no other staircase curve in (Δ, ϕ) contained in $N_X(V)$ that contains γ and has more steps than does γ. Let Γ be the union of those disks in Δ containing the steps of γ. Take the disk $D_0 \in \Delta$ containing the lowest step in γ and isotope it, along fibers of $N_X(V)$, onto the lowest disk in Σ_{D_0}. After the isotopy, consider the next lowest disk D_1 in Γ (with respect to the partial order). By the way we chose the new D_0, it contains a point x in $\partial N_X(V)$. So if $\Sigma_{D_1} \cap D_0 \neq \emptyset$ and if the highest disk in Σ_{D_1} met by D_0 also meets some fiber through D_0 below its intersection with D_0, then there is an integral curve from x through this fiber that enters the fiber neighborhood $N(D_1)$ of D_1 above Σ_{D_1} (before it leaves D_0) and later branches away from it. In fact, this curve must exit $N(D_1)$ from above and in order to do so, it must first leave $\partial N(V)$. However, this means there is a connecting strip that is crossed with index -1 by a staircase curve contained in (Δ, ϕ) that is also crossed by F with index $+1$ or 0, contradicting our hypothesis. So if $\Sigma_{D_1} \cap D_0 \neq \emptyset$, then the highest disk in Σ_{D_1} met by D_0 does not meet any fiber through D_0 below its intersection with D_0. In this case, we let L_0 be the union of this disk with D_0. Continuing inductively, suppose that D_k is a generating disk in Γ and that all generating disks in Γ that are below D_k in the partial order can be isotoped, along orbit segments in $N_X(V)$, into a connected subset L_0 of the leaf through D_0 in a way that does not reverse the partial ordering of any two disks. Note that some of these generating disks might intersect after the isotopy. However, if there is a positive orbit segment in $N_X(V)$ from one generating disk to another after the isotopy, then such an orbit segment existed before the isotopy as well. We can also assume that for any point in ∂L_0 there exists a curve in L_0 from D_0 to this point that corresponds to some staircase curve in Γ under the isotopy. Then if $\Sigma_{D_k} \cap L_0 \neq \emptyset$, we can argue, as above, that the highest disk in Σ_{D_k} met by L_0 does not meet any fiber through L_0 below its intersection with L_0. Hence, when we extend L_0 to include this disk, the union of D_k with all generating disks in Γ that are lower than D_k can be isotoped, along orbit segments in $N_X(V)$, onto L_0 in a way that does not reverse the partial ordering of any two disks. In some cases, it is possible to continue to extend L_0 as above so that all of Γ isotopes onto it. See Fig. 18.

If not, extend L_0, in this manner, as far as possible and let Γ_0 be the corresponding subset of Γ which projects onto L_0 and consists of disks through consecutive steps of γ. If D_j is the lowest disk in Γ that is not contained in Γ_0, then $\Sigma_{D_j} \cap L_0 = \emptyset$. In this case, L_0 meets the fiber neighborhood $N(D_j)$ of D_j (since D_{j-1} meets $N(D_j)$ and isotopes, along fibers of $N_X(V)$, into L_0). Hence, we can extend L_0 in its leaf to contain an integral curve from $D_0 \cap \partial N_X(V)$ that enters $N(D_j)$ before it leaves L_0 and later branches away from $N(D_j)$. Suppose this curve exits $N(D_j)$ from above. Since it must first leave $\partial N(V)$, this means there is a connecting strip that is crossed with index -1 by a staircase curve contained in (Δ, ϕ) that is also crossed by F with index $+1$ or 0, contradicting our hypothesis. See Fig. 19. Hence, the extended L_0 must exit $N(D_j)$ from below. So after isotoping the generating disks in Γ_0 into

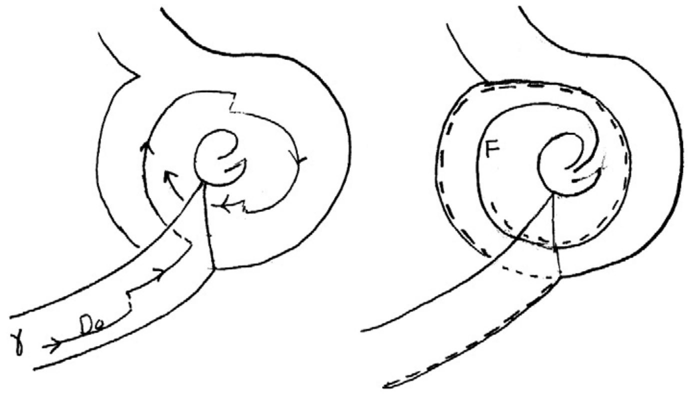

----- indicates L_0

Fig. 18.

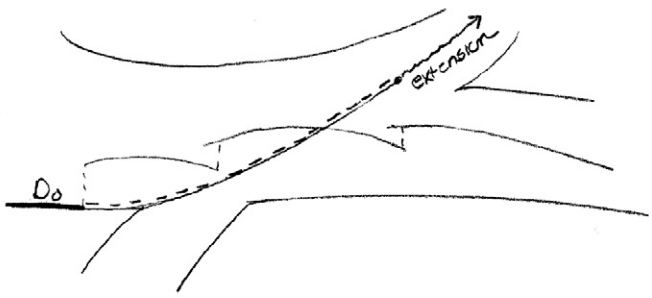

----- indicates L_0

Fig. 19.

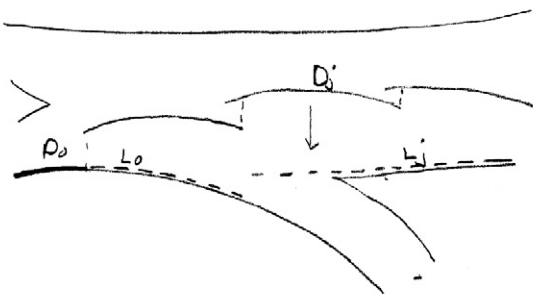

Fig. 20.

L_0, we can isotope D_j, along fibers of $N_X(V)$, onto the lowest disk in Σ_{D_j} without reversing the order of any two generating disks. See Fig. 20. Continue by repeating this procedure for the rest of Γ, this time beginning with D_j rather than D_0. We eventually get a staircase curve in (F, ϕ) corresponding to γ in the sense that each of its steps is the union of steps in γ after we move the latter into leaves of F. Furthermore, if there is a positive orbit segment in this new staircase curve from the image of one generating disk to the image of another, then such an orbit segment existed in γ as well.

We use the procedure above to assign a destination to each generating disk in Γ. However, this assignment is temporary since the destination we designate for a particular generating disk D may have to be adjusted slightly when γ and some other staircase curve γ' in (Δ, ϕ) both contain steps in D.

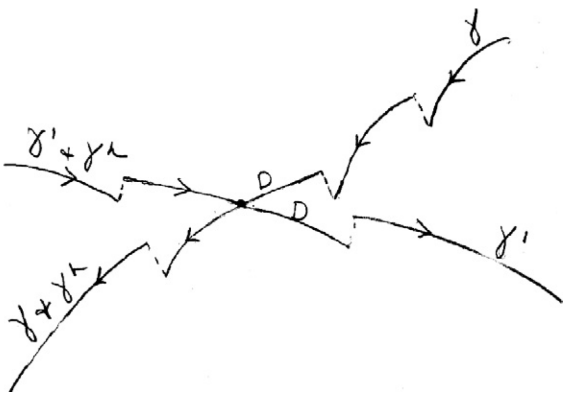

Fig. 21.

Specifically, let γ' be another maximal staircase curve in (Δ, ϕ), and let Γ' be the corresponding union of generating disks. Suppose that some generating disk $D \in \Delta$ is contained in both Γ' and Γ. If we choose γ above carefully, we can ensure that there is no disk lower than D in the partial order that is contained in more than one staircase curve in (Δ, ϕ). The procedure above assigns each generating disk contained in Γ to a temporary destination in a leaf of F. However, it is possible that the temporary destination of D is different when the procedure is applied to Γ' rather than to Γ. In the former (latter) case, let $L \subseteq N(\Gamma)$ ($L' \subseteq N(\Gamma')$, respectively) be the portion of a leaf containing the temporary destination of D. Without loss of generality, assume that the isotopic image of D in L' lies above its image in L. Then consider a "hybrid" staircase curve γ^h in (Δ, ϕ) which is a composition of γ', up until its intersection with D, followed by γ after that point. See Fig. 21. Let Γ^h be the union of disks in Δ containing the steps of γ^h and let D^h be the next generating disk above D in Γ^h; in particular, D^h is contained in $\Gamma \cap \Gamma^h$. As argued above, L' cannot be extended so that it enters the fiber neighborhood $N(D^h)$ of D^h and then exits it from above or we would have a contradiction to our hypothesis. Moreover, if such an extension exits $N(D^h)$ from below, then L could be extended to do the same. So we can reassign each generating disk in $\Gamma \cap \Gamma^h$ to its temporary destination as determined by γ^h, rather than by γ, without reversing the order of any two disks. Note that this only changes the temporary destination of disks that are contained in Γ and lie above D in the partial order.

It is also worth noting that there might be another maximal staircase curve γ'' containing D such that the temporary destination of D as determined by γ'' lies below its destination in L' as determined by γ' and γ^h. In this case, we can use a hybrid of γ' with γ'' to assign new temporary destinations for those disks met by γ'' that lie above D in the partial order. Continuing in this manner, for each of the finite number of maximal staircase curves in (Δ, ϕ) that meet D, we eventually find consistent temporary destinations in leaves of F for D and those generating disks in Δ that lie below D in the partial order.

We next proceed, as above, to find a consistent temporary destination in F for the next lowest disk E in the partial order that is shared by two or more maximal staircase curves in (Δ, ϕ). (It is possible that there is more than one choice for E, and in that case neither $E < D$ nor $D < E$.) This will not change the temporary destination of D assigned in the previous step. So, continuing inductively and in this manner, we eventually find consistent temporary destinations in leaves of F for each of the finitely many generating disks in Δ. Then, whenever the temporary destinations of two disks intersect, we bump these destinations forward into distinct nearby leaves in such a way that the partial ordering of Δ is preserved. After doing so, the final destination for the elements of Δ is defined and respects the partial ordering of Δ, hence we can carry out the entire isotopy. □

We now describe an algorithm for breaking up a generating set into smaller pieces to produce a modified branched surface. This will prove the following:

Theorem 2.3. *Let ϕ be a nonsingular Reebless flow and let W be a branched surface constructed from ϕ. Then ϕ is transverse to a foliation F if and only if W can be modified to carry F while staying transverse to ϕ.*

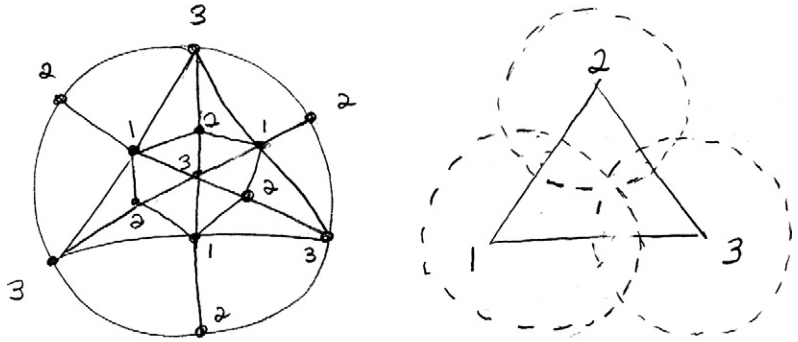

Fig. 22.

Proof. Suppose ϕ is transverse to some foliation F and, without loss of generality, assume that the orbits of ϕ are parameterized by arc length. Let $\Delta = \{D_i\}_{1 \leqslant i \leqslant n}$ be a generating set for a branched surface W constructed from ϕ. If W carries F, then we are done. So suppose this is not the case. Since F is transverse to ϕ, we can construct a branched surface V carrying F using another generating set X for ϕ contained in leaves of F such that $X \cap \Delta = \emptyset$. Then, when we cut M open along X to obtain $N_X(V)$, each element of Δ becomes embedded in the interior of $N_X(V)$ and is transverse to the fibers. Extending or contracting elements of Δ, we can ensure that at most finitely many fibers of $N_X(V)$ meet $\partial\Delta$ more than once. After doing so, let T be the maximal time it takes for a point in $X \cup \Delta$ to flow back into $X \cup \Delta$. By change of metric, if necessary, we can assume that $T < 1$.

Now, the projection of $\partial\Delta$ along fibers of $N_X(V)$ onto V produces a finite graph. Each staircase loop in (Δ, ϕ) contained in $N_X(V)$ corresponds to a cycle of generating disks which, when projected, gives a (possibly self-intersecting) annulus in V. There exists a minimal length generating loop for that annulus contained in its boundary, hence contained in the finite graph produced above. It follows that there exists a minimal horizontal length H on staircase loops in (Δ, ϕ) contained in $N_X(V)$.

We shall later wish to consider staircase curves in (Δ, ϕ) with the property that the minimal distance between any two steps lying in the same element D of Δ is at least $\min\{\frac{H}{2}, \frac{\operatorname{diam} D}{2}\}$. So we note here that since each generating disk $D \in \Delta$ can be covered with finitely many disks of diameter less than $\min\{\frac{H}{2}, \frac{\operatorname{diam} D}{2}\}$, there exists an upper bound P on the number of steps in such a staircase curve.

Now for each natural number k, consider those trivial connecting curves in V with the property that over any μ-companion curve, there exists either a staircase in (X, ϕ) with a step of length at least T^{k+1} or a curve in X with length at least T^{k+1}. Let L_k be a lower bound on the horizontal length of staircase curves in (Δ, ϕ) that project onto such a connecting curve.

Now find $\varepsilon_k > 0$ with the property that flowing any disk D embedded in $\bigcup_{i=1}^{n} D_i$ with diameter less than ε_k forward or backward for time at most $\frac{T}{3}$ gives a disk of diameter less than T^{k+1}. Cover each element of Δ by disks of diameter less than $\min\{\frac{L_k}{3P}, \varepsilon_k\}$ in the following manner: For each $D_i \in \Delta$, triangulate D_i with a graph of even valence (except along ∂D_i) so that each edge has length less than $\frac{1}{2}[\min\{\frac{L_k}{3P}, \varepsilon_k\}]$. Also ensure that for every vertex v, the union of all faces whose boundary contains v meets each fiber of $N_X(V)$ at most once. Cover each vertex of the graph with a disk of diameter less than $\min\{\frac{L_k}{3P}, \varepsilon_k\}$ so that any point $x \in D_i$ is contained in at least one and at most three disks and so that no disk meets a fiber of $N_X(V)$ more than once. Next, number the disks covering each $D_i \in \Delta$ 1, 2 and 3 so that no two disks of the same number meet (see Fig. 22). Then lift all disks numbered 1 forward along the flow for time $\frac{T}{3}$ and push all disks numbered 3 backward along the flow for time $\frac{T}{3}$. (Leave those labeled 2 fixed.)

We may assume that the new collection Δ_k of disks satisfies the conditions for a generating set transverse to ϕ; specifically, Δ_k generates a branched surface W_k. (Note that by the way we chose T, each element of Δ_k is contained in the interior of $N_X(V)$ and any orbit segment between points in $\Delta_k \cup X$ has length at least $\frac{T}{3}$.)

If at the kth step, the branched surface W_k carries the foliation F, then we are done. So assume that for every k, W_k does not carry F. We show that then the flow must admit self-return disks, contradicting our hypothesis. By Lemma 2.2, W_k not carrying F means that either there exists a staircase loop in (Δ_k, ϕ) contained in $N_X(V)$ or some connecting strip K_k in $N_X(V)$ is crossed by a staircase curve in (Δ_k, ϕ) with index -1, yet is crossed with a different index by F.

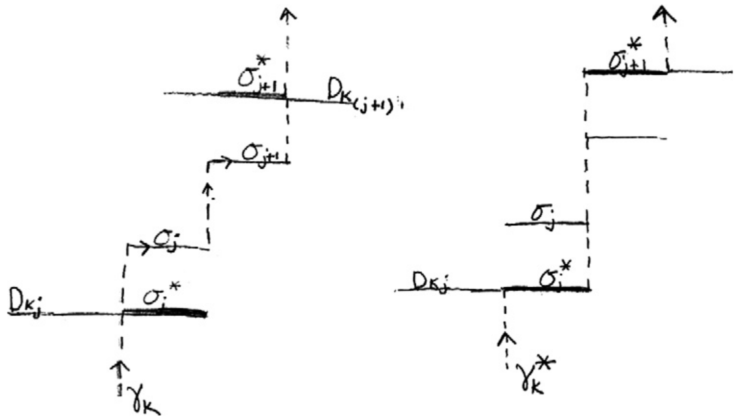

Fig. 23.

Now for every k, each staircase curve in (Δ_k, ϕ) corresponds to a staircase curve in (Δ, ϕ). Specifically, suppose that $\gamma_k = \tau_1 * \sigma_2 * \cdots * \tau_{m-1} * \sigma_m * \tau_m$ is a staircase in ϕ, where each step σ_j is contained in an element D_{k_j} of Δ_k. For each $j \leqslant n$, consider the preimage σ_j^* of σ_j in $\bigcup_{i=1}^n D_i$ before we flow the disk embedded in $\bigcup_{i=1}^n D_i$ that corresponds to D_{k_j} forward or backward onto D_{k_j}. (If these disks are the same, then $\sigma_j^* = \sigma_j$.) By the way we chose T and Δ_k, no orbit segment from σ_j^* to σ_j intersects $\partial N(V)$. In other words, each orbit segment between σ_j^* and σ_j is contained in a fiber of $N(V)$. Furthermore, for each $j < m$, the interiors of any pair of orbit segments of ϕ between $\text{int}(\sigma_j^*)$ and $\text{int}(\sigma_j)$ and between $\text{int}(\sigma_{j+1}^*)$ and $\text{int}(\sigma_{j+1})$, respectively, cannot intersect. So each τ_j can be extended and/or contracted to obtain a positive orbit segment of ϕ from σ_j^* to σ_{j+1}^* whose interior misses Δ. This gives a curve γ_k^*, homeomorphic to γ_k, that either is contained in an element of Δ or is a positive staircase curve in (Δ, ϕ) corresponding to γ_k. (See Fig. 23, for example.) Furthermore, if γ_k is contained in $\text{int}\, N_X(V)$, then so is γ_k^* and both curves project onto the same curve in V; i.e., $\pi_V(\gamma_k) = \pi_V(\gamma_k^*)$. So if γ_k crosses a connecting strip in $N_X(V)$ with index -1 then, by its construction, γ_k^* does so as well. It should be noted that if γ_k is a loop, then γ_k^* is also a loop.

If γ_k is contained in $N_X(V)$, then by our construction of Δ_k at most three consecutive steps in γ_k have preimages in the same element of Δ. Hence, each step in γ_k^* corresponds to at most three steps in γ_k. It follows that the length of each step in γ_k^* cannot exceed $3\varepsilon_k$. Now suppose that there are two steps of γ_k^* contained in some $D \in \Delta$. We can, in fact, choose D and these steps so that there is a subcurve of γ_k^* from the terminal point of the lower step to the initial point of the higher that does not meet any other generating disk more than once and whose interior does not meet D. If we then take any arc σ_k in D from the higher step to the lower, it is part of a staircase loop in (Δ, ϕ) whose steps are contained in distinct elements of Δ. In particular, we choose this loop so that one of its steps is a composition of σ_k with two arcs contained in distinct steps of γ_k^*, and the rest of it is the subcurve of γ_k^* described above. So the length of σ_k must be at least $H - (2 + P)(3\varepsilon_k)$. (This follows from the way we chose H and P.) In such cases, suppose there is another staircase curve contained in γ_k^* that begins and ends in the same generating disk D'. We can choose this curve so that if it contains a shorter staircase curve with the same property, the latter must be the subcurve of γ_k^* considered above, with ends in D. Since $\varepsilon_k \to 0$ as $k \to \infty$, for large enough k the length of σ_k is at least $\frac{H}{2}$. So this (possibly longer) subcurve of γ_k^* also has at most P steps, and any arc σ_k' between the steps of γ_k^* containing its ends has length at least $H - (2 + P)(3\varepsilon_k)$. Continuing in this manner, we can argue that for k sufficiently large, γ_k^* contains at most P steps and its horizontal length cannot exceed $3P\varepsilon_k$. In particular, for large enough k the staircase curve γ_k^* cannot be a loop; i.e. there can be no staircase loops in (Δ_k, ϕ) contained in $N_X(V)$.

Hence, for k sufficiently large, there exists a connecting curve κ_k corresponding to a connecting strip K_k that is crossed with index -1 by a staircase curve γ_k in (Δ_k, ϕ), and is also crossed with a different index by F. Moreover, if we choose k large enough, the connecting curve κ_k is so short that it must be trivial.

Now, let α_k be a μ-companion of κ_k. Without loss of generality, assume that α_k is parameterized by $[0, 1]$, and let ∂_1 and ∂_2 be the points of intersection of the ends of K_k with the interior of the fibers over $\alpha_k(0)$ and $\alpha_k(1)$, respectively. By definition, the branching along α_k is in only one direction. (In this case, α_k contains no connecting

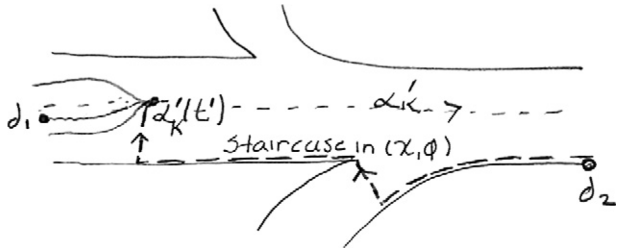

Fig. 24.

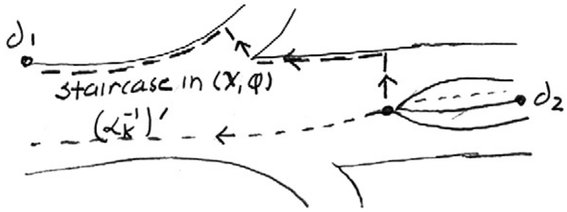

Fig. 25.

curves.) So either there exists an integral curve of F over α_k beginning at ∂_1 or there exists an integral curve of F over α_k^{-1} beginning at ∂_2. Since F crosses K_k and since κ_k is homotopic to α_k in V, there are three possibilities:

(1) F crosses K_k with index 0,
(2) F crosses K_k with index $+1$ and there exists an integral curve α_k' of F over α_k that begins at ∂_1 and ends above ∂_2, or
(3) F crosses K_k with index $+1$ and there exists an integral curve $(\alpha_k^{-1})'$ of F over α_k^{-1} that begins at ∂_2 and ends below ∂_1.

In case 1, there exists an integral curve of F over α_k that is contained in $\partial N_X(V)$ and joins ∂_1 and ∂_2 (since no curve over α_k can leave $\partial N_X(V)$ and reenter it later). So ∂_1 and ∂_2 are contained in the same component of $\partial(M - N_X(V))$; hence, they can be joined by a curve in the corresponding element of X.

In case 2, there exists a t', $0 < t' < 1$, such that $\alpha_k'(t)_{0 \leqslant t \leqslant t'} \subseteq \partial N_X(V)$ and $\alpha_k'(t)_{t' < t \leqslant 1} \subseteq \operatorname{int} N_X(V)$. So the only branchings of $N_X(V)$ over $\alpha_k(t)_{t'<t\leqslant 1}$ are in the direction shown in Fig. 24. Hence, there exists a staircase curve in (X, ϕ) over α^{-1} that begins in ∂_2 and ends in ∂_1, as shown.

In case 3, as in case 2, there exists a t', $0 < t' < 1$, such that $(\alpha_k^{-1})'(t)_{0 \leqslant t \leqslant t'} \subseteq \partial N_X(V)$ and $(\alpha_k^{-1})'(t)_{t'<t\leqslant 1} \subseteq \operatorname{int} N_X(V)$. Furthermore, the only branchings of $N_X(V)$ over $\alpha_k^{-1}(t)_{t'<t\leqslant 1}$ are in the direction shown in Fig. 25. So, again, there exists a staircase curve in (X, ϕ) over α_k^{-1} that begins in ∂_2 and ends in ∂_1, as shown.

In any case, it follows that then there exists a null homotopic staircase loop in $(X \cup \Delta, \phi)$ over $\kappa_k^*(\alpha_k)^{-1}$. Specifically there exists a staircase curve γ_k in (Δ_k, ϕ) over κ_k and the ends of this curve are connected by a staircase curve in (X, ϕ) over α_k^{-1} or by a curve in some element of X over α_k^{-1} (e.g., case 1). By our construction of Δ_k, each orbit segment in this staircase loop has length at least $\frac{T}{3}$, and each step that is contained in Δ_k has length at most T^{k+1}. Furthermore, either we can choose α_k so that the steps in X have length less than T^{k+1} or the horizontal length of the staircase curve γ_k^* in (Δ, ϕ) corresponding to γ_k (in the manner described above) is at least L_k. In the latter case, each step in γ_k^* corresponds to at most three steps in γ_k (which are contained in distinct elements of Δ_k). So each step in γ_k^* has length less than $3[\frac{L_k}{3P}]$. However, there are at most P steps in γ_k^*, so this means the horizontal length of γ_k^* is less than $3P[\frac{L_k}{3P}]$ or L_k, a contradiction. So all steps in the staircase loop in $(\Delta \cup X, \phi)$ over $\kappa_k^*(\alpha_k)^{-1}$ have length less than T^{k+1} for each k. It follows that for all k sufficiently large, there exists a null homotopic $(T^{k+1}, \frac{T}{3})$-staircase loop in ϕ. Now $T^{k+1} \to 0$ and $(T/3)/T^{k+1} \to \infty$. Furthermore, by assumption, ϕ is transverse to a foliation F. So by Lemma 2.1, there is a self-return disk for ϕ, contradicting our hypothesis that ϕ is Reebless.

Hence, W_k carries F, for all sufficiently large k. $\square$

Corollary 2.4. *Suppose ϕ is a nonsingular Reebless flow that is not transverse to any foliation. If ϕ lifts to a flow in a finite cover that is transverse to some foliation F, then there exists a branched surface W constructed from ϕ that does not carry any foliation yet lifts to one that does.*

Proof. Let ϕ be a nonsingular Reebless flow, and let $\hat{\phi}$ be the lift of ϕ to some finite cover $\hat{M}$ of the ambient manifold M. Suppose that ϕ is not transverse to a foliation, and let $\{\Delta_k\}$ be a sequence of generating sets for ϕ constructed by modifying some generating set Δ, as in the proof of Theorem 2.3. Each Δ_k lifts to a generating set $\hat{\Delta}_k$ for a branched surface $\hat{W}_k$ transverse to $\hat{\phi}$. In fact, the lifted generating set $\hat{\Delta}_k$ can also be obtained by modifying the lift $\hat{\Delta}$ of Δ, in the manner described in the proof of Theorem 2.3. So if $\hat{\phi}$ is transverse to a foliation F, then the branched surface generated by $\hat{\Delta}_k$ carries F, for k sufficiently large. However, this branched surface is the lift of the branched surface generated by Δ_k, which, by assumption, does not carry a foliation. □

Acknowledgement

The authors would like to thank the referee for many helpful comments that improved this paper.

References

[1] I. Agol, T. Li, An algorithm to detect laminar 3-manifold, Geom. Topol. 7 (2003) 287–309.
[2] D. Calegari, Useful branched surfaces which carry nothing, The Mathematics ArXiv GT/0010146.
[3] J. Christy, S. Goodman, Branched surfaces transverse to codimension one foliations, preprint.
[4] D. Fried, Cross sections to a flow, Topology 21 (1982) 265–282.
[5] D. Gabai, Problems in foliations and laminations, in: Geometric Topology, Athens, Georgia, 1993, in: AMS/IP Stud. Adv. Math., vol. 2.2, 1993, pp. 1–33.
[6] D. Gabai, U. Oertel, Essential laminations in 3-manifolds, Ann. Math. 130 (1989) 41–73.
[7] S. Goodman, Vector fields with transverse foliations, II, Ergodic Theory Dynam. Systems 6 (1986) 193–203.
[8] J. Milnor, On the existence of a connection with curvature zero, Comment. Math. Helv. 32 (1958) 215–223.
[9] R. Naimi, Foliations transverse to fibers of Seifert manifolds, Comment. Math. Helv. 69 (1994) 155–162.
[10] S.P. Novikov, Topology of foliations, Trudy Moskov. Mat. Obshch. 14 (1965) 248–278 (in Russian); Trans. Moscow Math. Soc. (1967) 268–304.
[11] S. Schwartzman, Asymptotic cycles, Annals of Math. 66 (1957) 270–283.
[12] S. Shields, Classifying foliations of 3-manifolds via branched surfaces, Topology Appl., submitted for publication.
[13] R.F. Williams, Expanding attractors, Inst. Haute Études Sci. Publ. Math. 43 (1973) 473–487.
[14] J. Wood, Bundles with totally disconnected structure group, Comment. Math. Helv. 46 (1971) 257–273.