

THE SIMPLEST BRANCHED SURFACES FOR A FOLIATION

SANDRA SHIELDS

Given a foliation F of a closed 3-manifold and a Smale flow ϕ transverse to F , we associate a “simplest” branched surface with the pair (F, ϕ) , which is unique up to two combinatorial moves. We show that all branched surfaces constructed from F and ϕ can be obtained from the simplest model by applying a finite sequence of these moves chosen so that each intermediate branched surface also carries F . This is used to partition foliations transverse to the same flow into countably many equivalence classes.

Introduction

In this paper we study branched surfaces constructed from transversely orientable foliations of closed orientable Riemannian 3-manifolds.

From the time of their introduction by Bob Williams [1969], branched manifolds have been powerful tools in the study of the dynamical systems, foliations and laminations. The one-dimensional case, branched 1-manifolds or *train tracks*, was introduced earlier to study Anosov diffeomorphisms [Williams 1970] and was used by Thurston [1988] to describe the dynamics of surface automorphisms. Branched surfaces were constructed by Williams [1973] to study the dynamics of hyperbolic expanding attractors for C^1 diffeomorphisms of compact 3-manifolds and have since been used to obtain many important results in the theory of foliations and laminations of 3-manifolds. For example, Gabai [1987] used branched surfaces to construct taut foliations, which allowed him to identify the minimal genus spanning surface for a wide class of knots. Brittenham [1995] showed that a 3-manifold contains an essential lamination if and only if it contains one that is carried (with full support) by one of a finite collection of normal branched surfaces. Agol and Li [2003] used branched surfaces to develop an algorithm for determining if a manifold contains such a lamination. (For more on the use of branched surfaces in the study of laminations and foliations, see also [Gabai and Kazez 1998; Gabai and Oertel 1989; Oertel 1988].)

There are various ways in which a branched surface can be constructed from a foliation. They are all similar in nature and involve cutting the ambient manifold

MSC2000: primary 57R30, 57M50; secondary 57M10, 57M20, 57N10.

Keywords: branched surface, foliation, transverse flow.

open along a subset of the leaves and then modding out by some quotient. The branched surfaces we consider are constructed according to a technique introduced by Christy and Goodman [1987]. (Details are given in Section 1.) This technique utilizes a nonsingular flow ϕ transverse to the foliation F , and we say that the resulting branched surface W carries (F, ϕ) .

A branched surface carrying (F, ϕ) often reflects the topology of F . For example, topological properties such as tautness or the R -covered property can often be detected from it [Goodman and Shields 2000; Shields 2004; 1997]. Indeed, a primary motivation for using branched surfaces to study foliations is that a large number of foliations can be approximated by the same branched surface. However, there is great variation among branched surfaces carrying the same foliation, and our ability to extract information about a foliation from a branched surface often depends on which branched surface we use. In fact, a central issue when using branched surfaces to study foliations is the search for the right branched surface.

For example, if we wish to show C^1 stability of a topological property for some foliation F , one approach is to show that F is carried by a *standard* branched surface (defined in Section 1) that carries only foliations with that property [Shields 1991]. (Here we are using the C^1 metric defined by Hirsch [1973], where a nearby foliation is obtained by perturbing the tangent bundle to the leaves to another integrable plane field.) When we are unable to find such a branched surface, it is often unclear whether or not one exists.

We show that for any foliation F and any transverse flow ϕ that meets a certain criterion, there is a natural choice for a branched surface carrying (F, ϕ) . (The condition we impose on the flow ϕ is harmless since it is satisfied by a dense set of flows in the C^0 topology.) Specifically, we define an equivalence relation on the set of branched surfaces transverse to ϕ . Like Penner's equivalence relation on measured train tracks [Penner and Harer 1992], our relation on branched surfaces is defined using two combinatorial moves that modify them. We associate a *simplest* branched surface $W_{F,\phi}$ with the pair (F, ϕ) , which is unique up to equivalence, and show that any standard branched surface carrying (F, ϕ) is obtained by modifying $W_{F,\phi}$ in a very restricted way (Theorem 3.4). (For example, each intermediate branched surface obtained during the modification process also carries (F, ϕ) .) Thus, if $W_{F,\phi}$ cannot be modified so that it carries only foliations with a certain topological property, no standard branched surface carrying (F, ϕ) has this property. We then use Theorem 3.4 to partition all foliations transverse to ϕ into countably many equivalence classes, each corresponding to a distinct simplicial complex.

We begin with a review of the construction of a branched surface from a pair (F, ϕ) . In Section 2 we discuss the relationship of the branched surface to the foliation F and the techniques for modifying it that we later use to define our equivalence relation. The main results are proved in Section 3.

1. Branched surfaces constructed from foliations

Throughout this paper, F will be a C^1 codimension one foliation of a closed orientable 3-manifold M and $\phi : M \times \mathbb{R} \rightarrow M$ will be a C^1 nonsingular flow on M that is transverse to F . We shall often refer to the *forward orbit* of a point $x = \phi(x, 0)$ in M under ϕ , meaning the set of points $\{\phi(x, t) : t > 0\}$. Likewise, the *backward orbit* is $\{\phi(x, t) : t < 0\}$.

Formally, a curve in M is a continuous map from a connected subset of \mathbb{R} into M . However, we shall consider a *curve* to be the image of such a map, and the map as a parametrization of the curve. If a curve has a negative boundary point, according to the orientation induced by the parametrization, we call this the *beginning* of the curve; a positive boundary point is the curve's *end*. A curve contained in a leaf of F is an *integral curve of F* .

1.1. Branched surface construction. The branched surfaces we associate with a foliation F are in the class of regular branched surfaces introduced by Williams [1973]. Since the construction we use is in an unpublished paper [Christy and Goodman 1987] and is a variation of the one in [Gabai and Oertel 1989], we describe it here, including all details necessary for this article.

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

- (i) Each D_i is contained in a leaf of F (hence is transverse to ϕ) and has finitely many boundary components.
- (ii) The forward and backward orbit of every point, under ϕ , meets $\text{int } \Delta = \bigcup_{i=1}^n \text{int } D_i$.
- (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 $\text{int } \Delta$ is finite.
- (iv) The forward orbit of any point in $\partial \Delta$ meets $\partial \Delta$ at most once before meeting $\text{int } \Delta$.


We can always find a generating set consisting of embedded disks. In particular, cover M with foliation boxes for F that are also flow boxes for ϕ , and select a slice from each box. Then, modify each slice slightly so that the resulting collection of disks satisfies the general position requirements above. If a generating set Δ for (F, ϕ) consists of embedded disks, we say that it is *standard*.

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 the leaves

of F to create an air pocket at each element of the generating set. 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 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 the embedded copy of the quotient space as *the branched surface W carrying F and ϕ* (or *carrying (F, ϕ)*). (Although the embedding of the quotient M^*/\sim is not unique, any two embedded copies 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 up to diffeomorphism of M . We emphasize that this notion of equivalence for branched surfaces differs from the usual equivalence up to diffeomorphism isotopic to the identity.) If Δ is standard, then we say that W is *standard*. The complement of any standard W in M is the union of finitely many open 3-balls.

A branched surface W can 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 two-dimensional complex with a set of charts defining local orientation preserving diffeomorphisms onto one of these models:



and further satisfying the condition 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, called *crossings*, 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, $\pi_W : N(W) \rightarrow W$ 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. Accordingly, we say points in the preimage of x *lie over* x . In particular, the interval orbit of ϕ^* that projects onto x will be called the *fiber of $N(W)$ over x* , as illustrated:



1.2. Foliations carried by a branched surface. The foliation F clearly gives rise to a foliation of $N_\Delta(W)$ whose leaves (some of which are branched) are transverse to the fibers. The branched leaves are precisely those that contain a boundary component of $N_\Delta(W)$. Specifically, the branched leaves in this foliation 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.) The figure shows a local picture of this foliation of $N_\Delta(W)$.

There are, of course, many foliations that are transverse to the fibers of $N_\Delta(W)$ with the property that each boundary component of $N_\Delta(W)$ is contained in a leaf. 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 each element




of Δ is contained in a leaf of this foliation. In other words, each such foliation of $N_\Delta(W)$ corresponds to a foliation of M that is carried by W .

2. Modifications of W


In this section, we describe techniques for changing a branched surface by modifying its generating set. We will use these techniques in Section 3 to define an equivalence relation on the set of branched surfaces transverse to a flow ϕ .

Given a branched surface W carrying (F, ϕ) with generating set Δ , we can modify W by enlarging an element D of Δ to include some compact integral surface D' of F such that $\partial D' \cap \partial D \neq \emptyset$, $\text{int } D' \cap \text{int } D = \emptyset$, $\partial(D' \cup D) \neq \emptyset$ and $\partial(D' \cup D)$ has finitely many components. This, in turn, enlarges the component B of $M - N(W)$ corresponding to D . We refer to this modification of D as an F -extension. If the new Δ is, in fact, another generating set for (F, ϕ) , then the extension corresponds to a *splitting* of W along the projection $\pi_W(D')$ of D' :



Clearly, we can extend $D \in \Delta$ to include *any* compact surface D' embedded transverse to the fibers of $N_\Delta(W)$ such that $\partial D' \cap \partial D \neq \emptyset$ and $\text{int } D' \cap \text{int } D = \emptyset$. If $\partial(D' \cup D)$ is nonempty and has finitely many components, and if $(\Delta - \{D\}) \cup \{D' \cup D\}$ satisfies conditions (ii)–(iv) for a generating set, then this set can be used to construct another branched surface W' transverse to ϕ . In this case the extension of D to include D' corresponds to a splitting of W along the projection $\pi_W(D')$ of D' to yield W' . However, if D' is not an integral surface of F , then there is no guarantee that F is carried by W' . We illustrate with a lower-dimensional example.

The branched 1-manifold W at the top of the next page carries a foliation F of T^2 with 2 Reeb components and 2 compact leaves. Yet, when we modify W by a



W embedded in a planar model of T^2 W' obtained by splitting W along γ

splitting along the curve γ (indicated by the dashed line) we obtain the branched 1-manifold W' which does not carry F .

Henceforth, a splitting of a branched surface carrying F that corresponds to an F -extension in its generating set shall be called an F -splitting.

We can also modify an element D of Δ by replacing it with a proper subset of itself. If this subset is connected and has finitely many boundary components, and if the new Δ also satisfies condition (ii) for a generating set, then we refer to this modification of D as a *contraction*. The connectedness condition ensures that a contraction does not change the cardinality of the generating set. This is also true for F -extensions provided that the elements of Δ are contained in distinct leaves of F . In such cases, each F -extension can be reversed by a contraction.

If a contraction of some $D \in \Delta$ yields another generating set, then it corresponds to a *pinching* of W . Specifically, suppose that the contraction deletes some open subset S of D , and let B be the component of $M - W$ corresponding to D . There exist two subsets S^+ and S^- of ∂B corresponding to S such that $\text{int } S^-$ flows injectively onto $\text{int } S^+$. We may identify each point of S^- with the point it flows to in S^+ , to partially collapse B . In other words, we can pinch these pieces of W together to obtain the branched surface that is generating by Δ after the contraction.


We can also modify Δ by adding another compact integral surface of F , chosen so that the requirements for a generating set are still satisfied. This, in turn, adds another component to $M - W$ by a move that we call an F -cutting of W .

At times, we shall want to move an element of the generating set Δ . This will involve flowing that element, either forward or backward along orbits of ϕ , onto some other compact integral surface of F . We describe this in detail below.

Definition 1. Given surfaces C and D embedded transverse to the flow (for instance, embedded in leaves of F), we say that C ϕ -covers D if there exists a continuous mapping of D along forward orbits of ϕ into C .

We realize this terminology can be slightly misleading since we are not claiming C is a covering space for D in the usual sense. For example, certain points in C might not flow backwards into D and points in the boundary of D might not contain

regularly covered neighborhoods:



Suppose D lies in Δ and C is a compact integral surface of F with boundary. According to the following proposition, if C ϕ -covers D , then replacing D with C in Δ yields another set satisfying condition (ii) for a generating set.

Proposition 2.1. *Given a foliation F and a nonsingular flow ϕ transverse to F , let $X = \{C_j\}_{1 \leq j \leq m}$ be a set of disjoint compact surfaces with boundary that are embedded in leaves of F . If there exists a generating set $\Delta = \{D_i\}_{1 \leq i \leq n}$ for a branched surface carrying (F, ϕ) such that the forward orbit of each point in $\bigcup_{i=1}^n D_i$ meets $\bigcup_{j=1}^m \text{int } C_j$, then X satisfies condition (ii) for a generating set.*

Proof. Without loss of generality, assume the orbits of ϕ are parametrized by arc length. We verify that, under the hypotheses, the orbit $\phi(x, t)_{t \in \mathbb{R}}$ of any $x \in M$ meets the interior of some element of X in both positive and negative time.

Since Δ satisfies condition (ii), for every $x \in M$ there exist $t > 0$ and $i \leq n$ such that $\phi(x, t) \in D_i$. It follows that there exists a $t_0 > t$ and $j \leq m$ such that $\phi(x, t_0) \in \text{int } C_j$. That is, the forward orbit of any $x \in M$ meets the interior of an element of X .

Since $\bigcup_{i=1}^n D_i$ is compact, there exists a time $T > 0$ with the property that every point in $\bigcup_{i=1}^n D_i$ flows into $\bigcup_{j=1}^m \text{int } C_j$ within time T . That is, for every $x \in \bigcup_{i=1}^n D_i$, there exists a time t , $0 \leq t \leq T$, and an integer $j \leq m$ such that $\phi(x, t) \in \text{int } C_j$. Since Δ is a generating set, the backward orbit of any $x \in M$ meets $\bigcup_{i=1}^n D_i$ infinitely many times. In particular, there exists a monotonically decreasing sequence $\{t_k\} \rightarrow -\infty$ such that for every k , $\phi(x, t_k) \in \bigcup_{i=1}^n D_i$. So for all k such that $t_k < -T$, the forward orbit of ϕ from $\phi(x, t_k)$ to $\phi(x, 0)$ meets $\bigcup_{j=1}^m \text{int } C_j$. That is, $\phi(x, t) \in \bigcup_{j=1}^m \text{int } C_j$ for some t such that $t_k + T \geq t \geq t_k$. So the backward orbit of any $x \in M$ meets the interior of some element of X . \square

Now suppose C is a compact integral surface of F with boundary that ϕ -covers some $D \in \Delta$. In particular, the interior of D flows continuously into the interior of C . Since Δ is a generating set, each point in Δ flows into $\text{int } \Delta$. It follows that each

point in Δ flows into $\text{int } C \cup \text{int}(\Delta - \{D\})$. So by [Proposition 2.1](#), if we replace D with C in Δ , we still have a set satisfying condition (ii) for a generating set. In the case that D flows continuously and injectively onto another integral surface C of F such that $(\Delta - \{D\}) \cap \{C\} = \emptyset$, the set $(\Delta - \{D\}) \cup \{C\}$ also satisfies the remaining conditions for a generating set. In this case, we say that C is a *vertical translate* of F and refer to this type of modification of Δ as a *vertical translation of D* .

As noted in [Section 1](#), there are vertical translations that do not change W . (If a generating set X is obtained after such a vertical translation in Δ , then there is a diffeomorphism of M mapping $N_\Delta(W)$ onto $N_X(W)$ that preserves fibers. However, it might not be possible to choose this diffeomorphism so that it preserves leaves of the foliations of $N_\Delta(W)$ and $N_X(W)$, respectively, that are induced by F .) Since these types of translations usually involve moving a generating surface to a nearby leaf, any vertical translation of a generating surface D that does not change W is called a *bumping of D* .

3. Standard minimal branched surfaces carrying (F, ϕ)

In this section, we consider only those branched surfaces carrying a pair (F, ϕ) for which some generating set is *standard minimal for (F, ϕ)* ; that is, there is a generating set consisting of disks embedded in leaves of F and no other branched surface can be constructed from F and ϕ using a generating set consisting of fewer disks (although it is possible that some branched surface could be constructed from F and ϕ using a generating set that contains fewer embedded surfaces, some of which are not simply connected). A *branched surface is standard minimal for (F, ϕ)* if it has a generating set with this property. (Clearly, if some generating set for W is standard minimal for (F, ϕ) , then all generating sets for (F, ϕ) that generate W have this property.) We let $\Omega_{F, \phi}$ denote the set of all branched surfaces that are standard minimal for (F, ϕ) . Since we can always find a generating set consisting of embedded disks (see [Section 1](#)), this set is nonempty for every (F, ϕ) . It is worth noting that all elements in a standard minimal generating set for (F, ϕ) are contained in distinct leaves of F (since otherwise, we could extend some element Δ in its leaf so that it merges with another to form one large generating disk.)

Using the density of Smale flows in the C^0 topology of nonsingular flows [[Oliveira 1976](#)], a frequent hypothesis in this section will be that ϕ is Smale. (Recall that a nonsingular flow ϕ on a manifold is a *Smale flow* provided that the chain recurrent set \mathcal{R} of ϕ has hyperbolic structure and topological dimension one, and that for any two points x and y in \mathcal{R} , the stable manifold of x and the unstable manifold of y intersect transversely. For a general discussion of Smale flows, see [[Franks 1982](#)]. Sullivan [[2000](#)] also gives a clear and detailed description of the

dynamics of these flows.) However, the only property of Smale flows that we shall use is the following: there exists a closed invariant one-dimensional subset \mathcal{R} of M such that each orbit of ϕ contains, in its limit set, some orbit in \mathcal{R} . (When ϕ is Smale, we can choose \mathcal{R} to be the chain recurrent set.) Our main result will be to show that for flows with this property, any branched surface $W \in \Omega_{F,\phi}$ can be modified, by F -extensions, contractions and bumpings in its generating set, to yield any other $V \in \Omega_{F,\phi}$ (Theorem 3.4). First, we shall need the following:


Lemma 3.1. *Let F be a foliation of M and ϕ be a Smale flow that is transverse to F . Any generating set Δ for (F, ϕ) can be modified by F -extensions and contractions to obtain a standard generating set with the same number of elements. In particular, if Δ is standard minimal for (F, ϕ) , then no generating set for (F, ϕ) has fewer elements than does Δ .*

Proof. Given a generating set $\Delta = \{D_i\}_{i=1,\dots,n}$ for (F, ϕ) , suppose that for some $i \leq n$, D_i is not an embedded disk. Since the chain recurrent set \mathcal{R} for ϕ has topological dimension one, we can take an arbitrarily small extension of D_i within its leaf so that its boundary misses \mathcal{R} ; \mathcal{R} being closed implies that after the extension there exists an open collar neighborhood of ∂D_i missing \mathcal{R} . General position arguments then allow us to modify D_i within this neighborhood so that the conditions for a generating set are still satisfied by Δ . We can then choose a subset K of D_i consisting of finitely many compact connected 1-manifolds, each missing \mathcal{R} , with the property that $D_i - K$ is connected and simply connected. In fact, there exists an open collar neighborhood $U(K)$ of K in its leaf whose closure misses \mathcal{R} (again, since \mathcal{R} is closed). For every point $x \in U(K)$, the forward orbit $\phi_{t>0}(x, t)$ of x meets $\text{int } \Delta - \{\text{cl } U(K) \cap \text{int } \Delta\}$ since it limits on some orbit in \mathcal{R} with this property. So by Proposition 2.1, we can remove $U(K)$ from D_i and still have a set that satisfies condition (ii) for a generating set. Since there exists an open collar neighborhood of $\partial U(K)$ in D_i that misses \mathcal{R} , general position arguments allows us to perturb $D_i - U(K)$ to a generating disk. It follows that there exists a generating set for (F, ϕ) consisting of embedded disks that has the same cardinality as does Δ . \square

Lemma 3.2. *Given a foliation F and a nonsingular flow ϕ transverse to F , let $X = \{C_j\}_{1 \leq j \leq m}$ and $\Delta = \{D_i\}_{1 \leq i \leq n}$ be generating sets for (F, ϕ) such that the elements of X are contained in distinct leaves of F and Δ is standard. There exists a function $C : \{1, \dots, n\} \rightarrow X$ with the property that for every $i \leq n$, $C(i)$ can be extended to an integral surface that ϕ -covers D_i . In addition, every forward orbit from D_i will meet the extended $C(i)$ before meeting an element of $X - \{C(i)\}$.*

Proof. For each $i \leq n$, the forward orbit of each point in D_i meets $\bigcup_{j=1}^m C_j$. So we can define a mapping $\phi_i : D_i \rightarrow \bigcup_{j=1}^m C_j$ such that for every $x \in D_i$, $\phi_i(x)$ is the first point in $\bigcup_{j=1}^m C_j$ met by the forward orbit of x .

Given $i \leq n$, we choose an element $C(i)$ in X that is *closest* to D_i along forward orbits of ϕ in the following sense: If $D_i \cap C_j \neq \emptyset$ for $j \leq m$, then $C(i) = C_j$; otherwise, we require that there exists an $x \in D_i$ with the properties that $\phi_i(x) \in C(i)$ and for every $y \in D_i$ there is *no* integral curve beginning at $\phi_i(y)$ and ending in the interior of the orbit from x to $\phi_i(x)$ that projects continuously, along segments of orbits of ϕ^{-1} , onto a curve in D_i from y to x . In other words, in the latter case the following situation *must not* occur:




(The vertical arrows indicate orbits of ϕ , the dotted curve an integral curve of F .) In the former case, C_j and D_i are contained in the same leaf of F , so C_j can be extended to contain D_i and $C(i)$ is unique (since the elements of X are in distinct leaves of F). In the latter case, we verify that there is also only one choice for $C(i)$ and that it extends to ϕ -cover D_i .

Suppose $D_i \cap C_j = \emptyset$ for every $j \leq m$ and that $x \in D_i$ has the properties above. Let x' be any other point in D_i . It suffices to show that $C(i)$ can be F -extended, over a curve in D_i , to meet the orbit of ϕ from x' to $\phi_i(x')$. For this, choose any curve $\alpha(s)_{0 \leq s \leq 1}$ from x to x' whose interior is contained in $\text{int } D_i$. Given any $s \in [0, 1]$, we can find an open integral disk U_s containing $\alpha(s)$ that flows continuously and injectively onto another open integral disk V_s containing $\phi_i(\alpha(s))$. In particular, we can embed a cylinder $D^2 \times [0, 1]$ in M so that the image of $D^2 \times \{0\}$ is $\text{cl } U_s$, the image of $D^2 \times \{1\}$ is $\text{cl } V_s$, the image of each copy of $[0, 1]$ is contained in an orbit of ϕ and the image of each copy of D^2 is contained in a leaf of F . By choosing U_s sufficiently small, we can ensure that for every $z \in \text{cl } U_s$, the orbit from z to $\phi_i(z)$ meets $\text{cl } V_s$. Now choose a finite subcover $\{U_0, U_{s_1}, \dots, U_{s_N}\}$ of α and, without loss of generality, assume that $U_0 \cap U_{s_1} \neq \emptyset$ and $U_{s_k} \cap U_{s_{k+1}} \neq \emptyset$ for all $k < N$. Since the disk $C(i)$ contains $\phi_i(\alpha(0))$, it can be F -extended to contain $\text{cl } V_0$. In particular, $C(i)$ can be extended, over $\text{cl } U_0$ to meet the orbit from $\alpha(s)$ to $\phi_i(\alpha(s))$, for all s such that $\alpha(s) \in \text{cl } U_0$.

Next, consider the embedded cylinder with base $\text{cl } U_{s_1}$, as described above, and choose $s_0 \in [0, 1]$ such that $\alpha(s_0) \in \text{cl}(U_{s_1}) \cap U_0$. If $C(i)$ does not meet the boundary of this cylinder after we extend it to meet the orbit of ϕ from $\alpha(s_0)$ to $\phi_i(\alpha(s_0))$,

then the extended $C(i)$ meets this orbit after its intersection with $\text{cl } V_{s_1}$. But this means there is an integral curve over $\alpha(s)_{0 \leq s \leq s_1}$ that begins in the interior of the orbit from x to $\phi_i(x)$ and ends at $\phi_i(\alpha(s_1))$, contradicting the way we chose x and $C(i)$. So $C(i)$ can be extended further, over $\text{cl } U_{s_1}$, to contain the image of $D^2 \times \{t\}$ in our cylinder for some $t \in [0, 1]$:



In this case, it meets the orbit of ϕ from $\alpha(s)$ to $\phi_i(\alpha(s))$ for all s such that $\alpha(s) \in \text{cl } U_{s_1}$. In this manner, we can argue inductively that $C(i)$ can be F -extended, over α , to meet the orbit of ϕ from x' to $\phi_i(x')$. Since the elements of X are contained in unique leaves of F , $C(i)$ defined as above is unique for each i .

So we have a function $C : \{1, \dots, n\} \rightarrow X$ such that for every $i \leq n$, $C(i)$ is the closest element in X to D_i along forward orbits of ϕ , and we can extend $C(i)$ in its leaf so that every positive orbit from D_i will meet the extended $C(i)$ before meeting any other element of X . Since D_i is a disk, we can choose this extension so that it ϕ -covers D_i . (If D_i were not a disk, then it would not, in general, be possible to extend $C(i)$ so that D_i flows continuously into it.) \square

For the following lemma, the map $C : \{1, \dots, n\} \rightarrow X$ is as defined above.

Lemma 3.3. *Given a foliation F and a Smale flow ϕ transverse to F , let $X = \{C_j\}_{1 \leq j \leq m}$ and $\Delta = \{D_i\}_{1 \leq i \leq n}$ be generating sets for (F, ϕ) such that the elements of X and Δ , respectively, are contained in distinct leaves of F and Δ is standard.*


If $C_j = C(i)$ for some $j \leq m$ and $i \leq n$, we can modify D_i by contractions and F -extensions so some vertical translate of D_i is contained in the same leaf as C_j .

The proof of this lemma is quite long. So before we begin, we give a brief description of the argument, together with an illustrating example

By [Lemma 3.2](#), each generating disk D_i in Δ maps continuously, along orbits of ϕ , onto some integral surface D_i' intersecting $C(i) \in X$. Our approach will be to consider subsets of D_i that map onto the same subset of D_i' , which we shall refer to as “layers of a stack” in D_i . (A precise definition is given in the proof below.) We modify each disk D_i in stages to eliminate all but the top layer of each stack (the layer closest to D_i'). When deleting the lower layers of a particular stack, we must ensure that D_i stays connected; otherwise this deletion is not a contraction. So, at each stage, it is usually necessary to delete more of D_i than just the lower layers of a stack. This can be done by contractions provided that we first extend D_i from the top layer to catch all orbits from the portion we intend to delete. (By [Proposition 2.1](#), this ensures that condition (ii) for a generating set is satisfied after our deletions). Our ability to do this relies on our assumption that D_i is simply connected. So before moving on to the next stack, we cut (i.e., take a small contraction of) the new generating surface so that it is also simply connected. We then repeat the modification process described above and eventually get D_i to have the desired property.


Since it might be helpful to follow through an example, we illustrate one stage of the modification process.

Suppose D_i is as shown in the figure and that D_i' is its planar projection. (We assume the transverse flow ϕ is perpendicular to the page and oriented toward the reader.) The shaded region indicates a stack in D_i consisting of 3 layers.




We wish to delete the two lower layers of this stack. However, when we delete the lowest layer, we disconnect D_i into two components. So since we want our deletion to be a contraction, we must also delete one of these components. In particular, we delete the component that does not contain the uppermost layer of our stack. The following extension ensures that when we do so, condition (ii) for

a generating set is still satisfied.




After the contraction, we have this surface:



It is not simply connected since the deletion of the middle layer of our stack creates a hole.

So we contract D_i further (as in the proof of [Lemma 3.1](#)) so that it is simply connected. The process is then repeated using the new D_i shown here.




Proof of [Lemma 3.3](#). Let $X = \{C_j\}_{1 \leq j \leq m}$ and $\Delta = \{D_i\}_{1 \leq i \leq n}$ be generating sets for branched surfaces carrying (F, ϕ) such that the elements of X and Δ , respectively, are contained in distinct leaves of F and Δ is standard.

For the map $C : \{1, \dots, n\} \rightarrow X$ defined in the proof of [Lemma 3.2](#), suppose that $C_j = C(i)$ for some $j \leq m$ and $i \leq n$. Since C_j can be extended to ϕ -cover D_i , the disk D_i flows continuously, along orbits of ϕ , onto some integral surface D_i'

that intersects C_j . Our concern will be those parts of D_i that flow onto the same subset of D_i' .

Let Φ be the collection of orbit segments of ϕ that define this continuous mapping of D_i onto D_i' . (In particular, the terminal point of each orbit of Φ is the image, in D_i' , of its initial point.) Without loss of generality, we can assume that $\partial D_i \cap \mathcal{R} = \emptyset$ and that there are only finitely many orbits of Φ containing more than one point in ∂D_i . Specifically, we can take an arbitrarily small F -extension of D_i , as in the proof of Lemma 3.1, so that its boundary misses \mathcal{R} . After this modification, there exists an open collar neighborhood of ∂D_i in its leaf that misses \mathcal{R} . General position arguments allow us to further modify D_i within this neighborhood to get a generating disk with the property that when we flow it continuously forward along orbits of ϕ onto an integral surface D_i' intersecting $C(i)$, at most finitely many of the orbit segments from D_i to D_i' meet ∂D_i more than once. This ensures that the image ∂^* of ∂D_i in D_i' is a closed and connected subset of $C(i)$ which is a one-manifold except at finitely many points where it self intersects. (It also ensures that D_i' has finitely many boundary components.) Consequently, ∂^* divides D_i' into finitely many connected regions whose interiors are disjoint. It follows that the preimage of ∂^* in D_i (namely, the set of points in D_i that map onto ∂^* when we flow D_i continuously forward, along orbits in the set Φ , onto D_i') divides D_i into finitely many connected regions whose interiors are pairwise disjoint. Let Γ be an open subset of $\int \int D_i$ such that each component of Γ is the interior of one of these regions. If the components of Γ all have the same image when we flow D_i onto D_i' , and if Γ is not contained in any larger open subset of $\text{int } D_i$ with these properties (i.e., one with more components), then we say Γ is a *stack* of D_i . We will refer to each component of a stack as a *layer* of that stack. (A stack may have only one layer.) Clearly, there are only finitely many stacks of D_i and they are pairwise disjoint.


For example, if D_i is as shown below, there are four stacks, only one of which has more than one layer. (The other three stacks are each connected.) The image in D_i' of the two-layer stack is labeled Γ' .



Let Γ be a stack of D_i that is not connected and let $\Gamma_1, \Gamma_2, \dots, \Gamma_p$ be the layers of Γ indexed so that for every $k \leq p$, Γ_k flows injectively onto Γ_{k+1} . We wish to delete $\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ from D_i . However, we need each orbit of ϕ to meet $\text{int } \Delta$ after this deletion. (This is necessary for our deletion to be a contraction.) This is ensured by our assumptions on ϕ and D_i . Specifically, for any $x \in \text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$, the closure of $\phi_{t>0}(x, t)$ contains an orbit in the closed one-dimensional chain recurrent set \mathcal{R} . Furthermore, this orbit in $\mathcal{R} \cap \text{cl}(\phi_{t>0}(x, t))$ meets $\text{int } \Delta - (\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1}) \cap \text{int } \Delta)$. (By assumption, it cannot meet ∂D_i , hence, it cannot meet $\partial(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$. If it meets $\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1}$, then it also meets Γ_p .) It follows that $\phi_{t>0}(x, t)$ also meets $\text{int } \Delta - (\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1}) \cap \text{int } \Delta)$. So by [Proposition 2.1](#), the closure of this set satisfies condition ii) for a generating set. That is, its interior is met by the forward and backward orbit of each point in M . So since $\Delta - \text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ has the same interior, condition ii) is still satisfied by Δ after we delete $\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ from D_i .

However, if $D_i - \text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ is not connected, then this deletion is still not a contraction. In this case, we delete a larger portion of D_i . In particular, we also delete the union R_1 of those components of $D_i - \text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ that do not contain Γ_p . (Each point in ∂R_1 that is not contained in ∂D_i is contained in $\partial(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$.) However, we must again make certain that requirement ii) for a generating set is still satisfied after the deletion. This can be ensured by first extending D_i .

Specifically, any $x \in \partial(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ flows into some $y \in \partial(\Gamma_p)$. Furthermore, since the entire disk D_i flows continuously forward, along orbits in the set Φ , onto another integral surface of F containing y in such a way that x maps to y , this is also true for any portion of D_i . So since the boundary of each component of R_1 intersects $\partial(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$, by flowing $\text{cl } R_1$ forward, we can find an integral (not necessarily connected) surface R_2 of F that meets all forward orbits in the set Φ from $\text{cl } R_1$ such that $R_2 \cup \text{cl } \Gamma_p$ is compact, connected and has finitely many boundary components. (It is possible that $R_2 \cap D_i$ contains points not in $\partial \Gamma_p$.) In fact, we can choose R_2 so that $\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1}) \cap R_2 = \emptyset$. Consequently, we can F -extend D_i to include all of R_2 (see figure) and then delete



$\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1}) \cup (R_1 - (R_1 \cap R_2))$ to obtain another connected integral surface \tilde{D}_i (which is not compact since it does not contain all its boundary points). We have then deleted all but the top layer of the stack Γ . (By assumption, the elements of Δ are contained in distinct leaves of F . So since \tilde{D}_i is contained in the same leaf of F as D_i , it does not intersect any other element of Δ . Hence, $(\Delta - \{D_i\}) \cup \{\tilde{D}_i\}$ is a collection of pairwise disjoint surfaces with boundary and compact closure.) By [Proposition 2.1](#) and the way we chose R_2 , condition (ii) for a generating set is still satisfied; that is, each orbit of ϕ meets the interior of some element of $(\Delta - \{D_i\}) \cup \{\tilde{D}_i\}$ in forward and backward time. We now want to delete the lower layers of some stack in \tilde{D}_i in this manner, but first we must modify \tilde{D}_i so that it is simply connected.

For this, we find a finite collection K of compact arcs in \tilde{D}_i such that $\tilde{D}_i - K$ is simply connected. In particular, we choose each arc in K to be the continuous image of some arc in ∂D_i (as we flow it partially forward or backward along the orbits in the set Φ). This ensures that $K \cap \mathcal{R} = \emptyset$, which allows us to remove K from \tilde{D}_i and still have condition (ii) for a generating set satisfied. (For details, see proof of [Lemma 3.1](#)).

After these deletions, \tilde{D}_i flows continuously, along orbits in the set Φ , onto a surface \tilde{D}_i' that is contained in the same leaf of F as C_i . We can think of \tilde{D}_i' as the first integral surface intersecting D_i' as we flow \tilde{D}_i forward along orbits of ϕ .

Just as we observed in the initial situation, $\partial \tilde{D}_i \cap \tilde{D}_i$ flows continuously (along orbits in the set Φ) onto a subset of \tilde{D}_i' which partitions it into finitely many regions whose interiors are pairwise disjoint. So we can again define stacks for \tilde{D}_i .

It is possible that \tilde{D}_i has more multilayer stacks than does D_i . However, if we consider the set of points in $D_i(\tilde{D}_i)$ that are contained in orbits in the set Φ through ∂D_i ($\partial \tilde{D}_i$ respectively), we see this set divides $D_i(\tilde{D}_i)$ into finitely many regions whose interiors are disjoint, and the closure of each stack of D_i (\tilde{D}_i respectively) is the union of such regions. Furthermore, our construction of \tilde{D}_i ensures that it has fewer regions of this type than does D_i . So, continuing in this manner, we would eventually get some \tilde{D}_i with only one such region, hence one (single-layer) stack.

At each stage of our modification process, \tilde{D}_i is simply connected but not compact. So to make it an embedded disk we technically need to delete the intersection of \tilde{D}_i with open neighborhoods U and $U(K)$ in D_i of $\text{cl}(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1})$ and K respectively. (The neighborhoods U and $U(K)$ can be chosen so that the conditions for a generating set are still satisfied since $\partial(\Gamma_1 \cup \Gamma_2 \cup \dots \cup \Gamma_{p-1}) \cap \mathcal{R} = \emptyset$ and $K \cap \mathcal{R} = \emptyset$. The argument here is again analogous to that used for [Lemma 3.1](#).) It is possible that when we do this, we shrink or subdivide some of the regions in \tilde{D}_i discussed above, as well as some stacks. However, we continue to modify, as above, to decrease the number of layers in some stack of \tilde{D}_i , rather than consider stacks of the resulting generating disk. Specifically, we continue by modifying \tilde{D}_i

to delete all but the top layer of one of its stacks (as we did for D_i), keeping in mind that at any stage \tilde{D}_i could be contracted to a generating disk. We eventually obtain \tilde{D}_i with only one (single-layer) stack, and it can be contracted to a generating disk that flows continuously and injectively onto a disk in the leaf containing C_j . \square

Although there are many branched surfaces carrying a foliation F , it is often desirable to find one with certain properties. In particular, certain structural criteria on a standard branched surface W carrying F are sufficient to guarantee that every foliation carried by W , including F , has some topological property such as tautness or the R -covered property; see [Goodman and Shields 2000; Shields 2004; 1997; 1996]. In such cases, we know that topological property is C^1 stable for F since all foliations sufficiently close to F are also carried by W [Shields 1991]; that is, each foliation within some $\varepsilon > 0$ of F , in the C^1 metric defined by Hirsch [1973], is carried by W . (There may also be foliations carried by W that are not within ε of F .) So a central issue in using branched surfaces to study foliations is the search for the right branched surface. According to the following theorem, if we cannot modify any given standard minimal branched surface W carrying (F, ϕ) , in a very restricted way, to obtain one with a desired property, then a standard branched surface with that property and carrying (F, ϕ) does not exist.

Theorem 3.4. *Let F be a foliation and ϕ be a nonsingular Smale flow transverse to F . Any branched surface $W \in \Omega_{F, \phi}$ can be modified to obtain any other standard branched surface V carrying (F, ϕ) by F -splittings and pinchings, followed by finitely many F -cuttings. In particular, if $V \in \Omega_{F, \phi}$, then any generating set for (F, ϕ) that generates W can be modified by F -extensions, contractions and bumpings to get any generating set for (F, ϕ) that generates V .*

Proof. Let $\Delta = \{D_i\}_{1 \leq i \leq n}$ be a generating set that is standard minimal for (F, ϕ) . Take another standard generating set $X = \{C_i\}_{1 \leq i \leq m}$ for (F, ϕ) , with $m \geq n$; without loss of generality, we can assume that the elements of X are contained in distinct leaves of F , since we can bump elements in any standard generating set to nearby leaves.

Consider the function $C : \{1, \dots, n\} \rightarrow X$ as defined in the proof of Lemma 3.2. Using Lemma 3.3, we modify Δ by contractions and F -extensions so that for all $i \leq n$, some vertical translate D_i' of D_i is contained in the same leaf of F as $C(i)$.

Suppose that for some $i \leq m$, there exist distinct $j, k \leq n$ such that $C_i = C(j) = C(k)$. That is, C_i is the closest element of X to both D_j and D_k along orbits of ϕ . We can extend D_j' and D_k' in their leaf so that they both intersect C_i . Then, by Proposition 2.1, a slight modification of $D_j' \cup D_k' \cup C_i$ can be used to replace both D_j and D_k in Δ . Although the resulting generating set might not be standard, by Lemma 3.1 we have a contradiction to our assumption that our original Δ was

standard minimal for (F, ϕ) . It follows that the function $C : \{1, \dots, n\} \rightarrow X$ is injective. So without loss of generality, assume $C(i) = C_i$ for all $i \leq n$.

If x is a point along an orbit of ϕ from D_i to D_i' for some $i \leq n$, then D_i flows continuously, along orbits of ϕ , onto an integral surface S_x of F through x . Since Δ is minimal, S_x cannot intersect $\Delta - \{D_i\}$ (As argued above, this follows from [Proposition 2.1](#) and [Lemma 3.1](#)). So no orbit from D_i to D_i' meets $\Delta - \{D_i\}$. That is, the vertical translation from D_i onto D_i' is a bumping, for all $i \leq n$.

For each $i \leq n$, we can use F -extensions and contractions of D_i' to obtain C_i . It follows that $\Delta = \{D_i\}_{1 \leq i \leq n}$ can be modified by F -extensions, contractions and bumpings to obtain $\{C_i\}_{1 \leq i \leq n}$; this corresponds to modifying the branched surface W generated by Δ using F -splittings and pinchings. If $n < m$, we can then add $\{C_i\}_{n+1 \leq i \leq m}$ to obtain X . By definition, these additions correspond to finitely many cuttings of W . \square

Given a nonsingular flow ϕ on some manifold M , we can use our modification techniques to define an equivalence relation on the set of branched surfaces that carry foliations transverse to ϕ . Specifically, given two branched surfaces W and V transverse to ϕ , we say W is *equivalent* to V if W can be modified to obtain V (up to some diffeomorphism of M) by splittings and pinchings, each of which does not change the number of complementary components in M . We use $[W]$ to represent the equivalence class of W under this relation.

If ϕ is Smale and F is transverse to ϕ , then by [Theorem 3.4](#) any two branched surfaces in $\Omega_{F,\phi}$ are equivalent. So given a foliation F and a transverse flow ϕ that is Smale, we can associate a *simplest* branched surface $W_{F,\phi} \in \Omega_{F,\phi}$ with (F, ϕ) , which is unique up to equivalence.

Theorem 3.5. *Let ϕ be a nonsingular Smale flow on M . The set of foliations transverse to ϕ can be partitioned into countably many equivalence classes so that there exists an injective function from the set of all such classes into a countable collection of simplicial complexes. Specifically, each equivalence class can be associated with a distinct branched surface K and each standard minimal branched surface for a foliation in that class can be obtained by modifying K .*

Proof. We can define an equivalence relation on the set of foliations transverse to ϕ by letting F be equivalent to G precisely when some $W_{F,\phi} \in \Omega_{F,\phi}$ is equivalent to some $W_{G,\phi} \in \Omega_{G,\phi}$ (By [Theorem 3.4](#), equivalence of two foliations F and G is independent of our choices for $W_{F,\phi}$ and $W_{G,\phi}$.) We can therefore associate the equivalence class for F with any standard branched surface in $[W_{F,\phi}]$ that carries a foliation. So it suffices to show that the set of branched surfaces that can be constructed from foliations of M and generated by disks is countable (up to diffeomorphism of M).

For any branched surface W carrying a foliation of M , the intersection W_ε of W with a small regular neighborhood of the branch set in the ambient manifold M is obtained by piecing together local neighborhoods of the crossings, each of which is modeled on either of the two lower diagrams on page 104. (We glue these local models together along the Y -shaped components of their boundaries in a manner dictated by the branch set.) The branched surface W can then be constructed (up to diffeomorphism) by gluing the boundaries of surfaces homeomorphic to the sectors of W to ∂W_ε . Since the branched surfaces we consider have finitely many crossings (that is, the branch set for each is a finite graph) and no boundary, there are only countably many possibilities for W , up to diffeomorphism.

By definition, our branched surfaces are embedded in the ambient manifold in a particular manner determined by the construction; moreover, we only distinguish between them up to diffeomorphism of M . When we restrict to branched surfaces generated by disks, the complement of each is simply connected. In particular, if W and V are diffeomorphic branched surfaces constructed from foliations of M and generated by disks, then the complement of each in M must be a disjoint union of open 3-balls. Hence, the diffeomorphism from W onto V extends to a diffeomorphism of M . It follows that the collection of all standard branched surfaces constructed from foliations of M is countable up to diffeomorphism of M . \square

The branched surface K that we associate with the foliation F might not carry F . However, it is equivalent to many that do. So although the topological behaviors of leaves in two equivalent foliations F and G can differ substantially, there is often a branched surface $W \in [K]$ carrying both F and G . In this case, if the intersection of the set $\{\gamma : \gamma = \pi_W(\gamma_F) \text{ for some integral curve } \gamma_F \text{ of } F\}$ with the set $\{\gamma : \gamma = \pi_W(\gamma_G) \text{ for some integral curve } \gamma_G \text{ of } G\}$ is sufficiently large, then F and G will share many important topological properties. See [Shields 2004; 1991; 1997; 1996].

Acknowledgements

The author thanks Sue Goodman for very helpful conversations and the referee for many useful suggestions and corrections.

References

- [Agol and Li 2003] I. Agol and T. Li, “An algorithm to detect laminar 3-manifolds”, *Geom. Topol.* **7** (2003), 287–309. [MR 2004e:57007](#) [Zbl 1037.57008](#)
- [Brittenham 1995] M. Brittenham, “Essential laminations and Haken normal form”, *Pacific J. Math.* **168**:2 (1995), 217–234. [MR 96c:57028](#) [Zbl 0838.57011](#)
- [Christy and Goodman 1987] J. Christy and S. E. Goodman, “Branched surfaces transverse to codimension one foliations”, preprint, 1987, Available at <http://www.math.unc.edu/Faculty/seg/construction.pdf>.

- [Franks 1982] J. M. Franks, *Homology and dynamical systems*, CBMS Regional Conference Series in Mathematics **49**, Amer. Math. Soc., Providence, 1982. [MR 84f:58067](#) [Zbl 0497.58018](#)
- [Gabai 1987] D. Gabai, “Foliations and the topology of 3-manifolds. III”, *J. Differential Geom.* **26**:3 (1987), 479–536. [MR 89a:57014b](#) [Zbl 0639.57008](#)
- [Gabai and Kazez 1998] D. Gabai and W. H. Kazez, “Group negative curvature for 3-manifolds with genuine laminations”, *Geom. Topol.* **2** (1998), 65–77. [MR 99e:57023](#) [Zbl 0905.57011](#)
- [Gabai and Oertel 1989] D. Gabai and U. Oertel, “Essential laminations in 3-manifolds”, *Ann. of Math.* (2) **130**:1 (1989), 41–73. [MR 90h:57012](#) [Zbl 0685.57007](#)
- [Goodman and Shields 2000] S. Goodman and S. Shields, “A condition for the stability of \mathbf{R} -covered on foliations of 3-manifolds”, *Trans. Amer. Math. Soc.* **352**:9 (2000), 4051–4065. [MR 2000m:57042](#) [Zbl 0965.57026](#)
- [Hirsch 1973] M. W. Hirsch, “Stability of compact leaves of foliations”, pp. 135–153 in *Dynamical systems* (Salvador, Brazil, 1971), edited by M. M. Peixoto, Academic Press, New York, 1973. [MR 48 #12555](#) [Zbl 0272.57015](#)
- [Oertel 1988] U. Oertel, “Measured laminations in 3-manifolds”, *Trans. Amer. Math. Soc.* **305**:2 (1988), 531–573. [MR 89d:57011](#) [Zbl 0652.57006](#)
- [Oliveira 1976] M. Oliveira, *C^0 -density of structurally stabled dynamical systems*, Thesis, University of Warwick, 1976.
- [Penner and Harer 1992] R. C. Penner and J. L. Harer, *Combinatorics of train tracks*, Annals of Mathematics Studies **125**, Princeton University Press, Princeton, NJ, 1992. [MR 94b:57018](#) [Zbl 0765.57001](#)
- [Shields 1991] S. Shields, *Branched surfaces and the stability of compact leaves*, Thesis, University of North Carolina, Chapel Hill, 1991.
- [Shields 1996] S. L. Shields, “Branched surfaces and the stability of compact leaves”, *Houston J. Math.* **22**:3 (1996), 591–620. [MR 98f:57051](#) [Zbl 0869.57033](#)
- [Shields 1997] S. L. Shields, “Branched surfaces and the simplest foliations of 3-manifolds”, *Pacific J. Math.* **177**:2 (1997), 305–327. [MR 98f:57032](#) [Zbl 0868.57028](#)
- [Shields 2004] S. Shields, “ \mathbf{R} -covered branched surfaces”, *Pacific J. Math.* **217**:2 (2004), 303–330. [MR 2005j:57003](#) [Zbl 1067.57027](#)
- [Sullivan 2000] M. C. Sullivan, “Visually building Smale flows in S^3 ”, *Topology Appl.* **106**:1 (2000), 1–19. [MR 2001g:37025](#) [Zbl 0964.37019](#)
- [Thurston 1988] W. P. Thurston, “On the geometry and dynamics of diffeomorphisms of surfaces”, *Bull. Amer. Math. Soc. (N.S.)* **19**:2 (1988), 417–431. [MR 89k:57023](#) [Zbl 0674.57008](#)
- [Williams 1969] R. F. Williams, “Expanding attractors”, pp. 79–89 in *Colloque de Topologie Différentielle* (Mont-Aigoual, 1969), Université de Montpellier, Montpellier, 1969. [MR 44 #4784](#) [Zbl 0208.25801](#)
- [Williams 1970] R. F. Williams, “The DA'' maps of Smale and structural stability”, pp. 329–334 in *Global Analysis* (Berkeley, 1968), vol. 1, edited by S.-S. Chern and S. Smale, Proc. Sympos. Pure Math. **14**, Amer. Math. Soc., Providence, R.I., 1970. [MR 41 #9296](#) [Zbl 0213.50303](#)
- [Williams 1973] R. F. Williams, “Expanding attractors”, *Inst. Hautes Études Sci. Publ. Math.* **43** (1973), 169–203. [MR 50 #1289](#) [Zbl 0279.58013](#)

Received May 4, 2004. Revised November 5, 2005.

SANDRA SHIELDS
COLLEGE OF CHARLESTON
66 GEORGE STREET
CHARLESTON SC
29424-1001
shieldss@cofc.edu