When incompressible tori meet essential laminations

Mark Brittenham[1] and Rachel Roberts[2]

Abstract. We show that there is only one obstruction to isotoping an essential lamination $\mathcal{L}$ in a 3-manifold M containing an essential torus T so that when both $\mathcal{L}$ and M are split open along T, the resulting lamination $\mathcal{L}|T$ is essential in the split open manifold $M|T$. This is the existence of a cylindrical component in $\mathcal{L}$. This generalizes previous results relating taut foliations and essential tori.

§0
Introduction

Roussarie and Thurston [Ro],[Th] showed that given a taut foliation $\mathcal{F}$ and an incompressible torus T in a 3-manifold M, the torus may be isotoped so that it is either transverse to the leaves of the foliation, or is equal to one of the leaves of the foliation. This result was known to be false in the more general context of a foliation $\mathcal{F}$ without Reeb components. The only obstruction, however [Ro], is the existence of a 'cylindrical component' of the foliation. This is an I-bundle over a torus or Klein bottle, which is saturated by the leaves of the foliation. The boundary component(s) of the bundle are leaves, and the interior of the bundle is foliated by open annuli and Möbius bands, which spiral in towards the boundary components in the same direction (see, e.g., [Ro,p.109]).

Even when $\mathcal{F}$ is taut and transverse to T, the foliation $\mathcal{F}|T$ (see §1 for notation) need not be taut; it may have half-Reeb components, which are solid tori in $M|T$ bounded by an annulus leaf of $\mathcal{F}|T$ and foliated as 'half' of a Reeb component (see Figure 1). However, if the foliation is $C^{(2)}$, an argument due to Hass shows that such components can be removed by a further $C^{(2)}$ isotopy of $\mathcal{F}$ (see [BNR]).

Figure 1

In this paper we show that these facts can be extended to $C^{(0)}$ foliations, and in fact to essential laminations:

Theorem:. *Given an essential lamination $\mathcal{L}$ and an incompressible torus T in a 3-manifold M, then, possibly after splitting $\mathcal{L}$ open along a finite number of leaves, either $\mathcal{L}$ can be isotoped so that it is transverse to T and $\mathcal{L}|T$ is essential in $M|T$, or $\mathcal{L}$ has a cylindrical component.*

The isotopy of $\mathcal{L}$ can be thought of as an ambient isotopy of M, and so can equivalently be thought of (by taking inverses) as an isotopy of T. It is more

Key words and phrases. essential lamination, incomressible torus.
[1] Research supported in part by NSF grant # DMS−9704811
[2] Research supported in part by an NSF Postdoctoral Fellowship

convenient, however, to think in terms of pushing $\mathcal{L}$ around, since it is $\mathcal{L}|T$ (and not $T|\mathcal{L}$!) that we will constantly be worrying about.

This result plays a central role in understanding the structure theory of taut foliations and essential laminations in graph manifolds, given in [BNR]; it implies that these objects come from their counterparts in the Seifert-fibered pieces that one obtains after splitting a graph manifold open along incompressible tori. Since essential laminations in Seifert-fibered spaces are extremely well-understood ([Br1],[Br2],[Cl],[JN],[Na]), we end up with a nearly complete picture of the existence and structure of essential laminations and taut or Reebless foliations in graph manifolds; see [BNR].

This result also gives a method for determining that a laminar 3-manifold M is atoroidal; if M contains an essential torus T, then one of three things must be true: the lamination $\mathcal{L}$ has a torus leaf, or the torus misses $\mathcal{L}$ (and so $M|\mathcal{L}$ contains an essential torus), or $\mathcal{L}$ (after splitting) can be isotoped to meet the torus tautly, as above. In the last case, $T\cap(M|B)$, where B is a branched surface carrying a splitting of $\mathcal{L}$, consists of a collection of essential bigons; disks whose boundaries each meet the branch locus of B in two points, and are not homotopic to a disk in B. Consequently, if we have an essential lamination with no torus leaves, whose complement is atoroidal (this is typically the case; see, e.g., [Br4]), and whose complement has no bigons, then our manifold must be atoroidal. This picture is very similar to that given by Wu [Wu] for detecting atoroidality after Dehn filling of an atoroidal manifold. It also gives us a way to search for incompressible tori in a laminar manifold which we might suspect has one, by 'stitching together' essential bigons in the complement of our lamination.

§1
Preliminary considerations

For basic information on essential laminations in 3-manifolds, see [GO]. For information on train tracks and incompressible laminations in surfaces, see [BC]. For convenience, we assume that our ambient manifold M is orientable; all of our arguments can be modified in a straightforward manner to handle the non-orientable case, after expanding the conclusion to include the obvious non-orientable analogue of a cylindrical component.

Given a 3-manifold M and an embedded torus T in M, we denote by $M|T$ the manifold resulting from splitting M along T, i.e., $M\backslash\mathrm{int}(\mathrm{N}(T))$. Given a foliation $\mathcal{F}$ or lamination $\mathcal{L}$ transverse to T, we denote by $\mathcal{F}|T$ or $\mathcal{L}|T$ the induced foliation or lamination in $M|T$. Similar notation also applies to 1-dimensional laminations in surfaces, split along simple closed curves.

Given an essential lamination $\mathcal{L}$ in a 3-manifold M, and an incompressible torus T in M, we can make $\mathcal{L}$ transverse to T by a small homotopy of $\mathcal{L}$. This is most easily achieved by making T transverse to a branched surface B which carries $\mathcal{L}$ (although see also [Br3]). For a foliation $\mathcal{F}$, this first requires splitting $\mathcal{F}$ open along a finite number of leaves; see [GO]. The intersection $\mathcal{L}\cap T=\lambda$ is then a 1-dimensional lamination in T. By Reeb Stability, the collection of trivial loops in λ is both open and closed in λ. By the 'usual arguments' (see, e.g., [Br1]), we can remove the trivial loops of λ by an isotopy of $\mathcal{L}$.

After this preliminary isotopy, $\mathcal{L}|T$ is almost essential in $M|T$; all but one of the properties in the definition of essential lamination [GO] are automatically satisfied:

(1) no leaf of $\mathcal{L}|T$ is a sphere, or a boundary-parallel disk, or a torus bounding a solid torus.

This follows immediately from the analogous property for $\mathcal{L}$, and because every loop in $\mathcal{L}\cap T$ is essential in T, hence in the leaves of $\mathcal{L}$.

(2) $(M|T)\backslash(\mathcal{L}|T)$ is irreducible.

If S is a 2-sphere in $(M|T)\backslash(\mathcal{L}|T)$, then since it lies in $M|T$ and $M|T$ is irreducible, S bounds a 3-ball B_1 in $M|T$. Since S also lies in $M\backslash\mathcal{L}$ and $M\backslash\mathcal{L}$ is irreducible, S also bounds a 3-ball B_2 in $M\backslash\mathcal{L}$. Therefore either $M = B_1 \cup B_2 =$ the 3-sphere (a contradiction) or $B_1 = B_2$ lies in $M|T$ and in $M\backslash\mathcal{L}$, i.e., in $(M|T)\backslash(\mathcal{L}|T)$. So $(M|T)\backslash(\mathcal{L}|T)$ is irreducible.

(3) Each leaf of $\mathcal{L}|T$ is π_1-injective in $M|T$.

If γ is a loop in a leaf L_0 of $\mathcal{L}|T$, which is null-homotopic in $M|T$, then γ is also null-homotopic in M. So since γ is contained in a leaf L of $\mathcal{L}$ (the one containing L_0) and $\mathcal{L}$ is essential, γ is null-homotopic in L. If we choose a null-homotopy $H : D \to L$ for γ, make it transverse to T, then $H^{-1}(T)$ consists of a finite collection of circles and arcs in the 2-disk D. But since γ lies in (we may assume the interior of, by pushing off) L_0, $\gamma\cap T = \emptyset$, so $H^{-1}(T)$ consists only of circles. Each circle maps to a loop in a leaf of λ. Since each circle bounds a subdisk of D, its image is a loop in T null-homotopic in M, so since T is essential, each loop is null-homotopic in T. Since λ consists of circles essential in T and lines, each circle in $H^{-1}(T)$ therefore maps to a loop which is null-homotopic in its leaf of λ. So we can redefine H on each subdisk to map into $\lambda\cap L_0$; by pushing off slightly, we can then map D into L_0. So γ is null-homotopic in L_0; so $\mathcal{L}|T$ is π_1-injective.

This property is equivalent to the two geometric properties that the boundary leaves of $\mathcal{L}|T$ are (geometrically) incompressible in $(M|T)\backslash(\mathcal{L}|T)$, and $\mathcal{L}|T$ does not contain a torus which is compressible in $M|T$.

(4) $\mathcal{L}|T$ is end-incompressible in $M|T$.

This also follows immediately from the fact that $\mathcal{L}$ is transverse to T; an end-compressing disk for $\mathcal{L}|T$ cannot be one for $\mathcal{L}$, so can be homotoped into $\mathcal{L}$; the disk in $\mathcal{L}$ can then be made disjoint from T as above.

The only properties missing, therefore, in order to make $\mathcal{L}|T$ essential in $M|T$ is for the leaves of $\mathcal{L}|T$ to be (geometrically) ∂-incompressible in $M|T$, and for $\mathcal{L}|T$ to admit no half-Reeb components. (These are equivalent to the algebraic condition of ∂-injectivity; see [GO].) It is easy to see that we cannot expect this to be the case in general. We can, for example, easily perturb something which is essential after splitting along T so that it instead has leaves which are boundary-parallel annuli, by 'folding' the lamination near a circle of intersection with T.

However, we can see that, as in the case of a compact surface in M, if a lamination fails to be essential only because one of these remaining properties fail, this is due to the presence of ∂-parallel annuli. For a half-Reeb component, this is immediate. In the other case, let D be a boundary-compressing disk for $\mathcal{L}|T$. Because $\mathcal{L}$ is essential and meets T transversely, and because we have already removed trivial circles of intersection of $\mathcal{L}$ with T, the intersection $\lambda = \mathcal{L}\cap T$ is an essential lamination in the torus T. It therefore consists either entirely of non-compact leaves, with product regions lying between, or contains compact leaves, possibly with Reeb or Kronecker-type laminations lying between (see, e.g., [Br1]). In particular, every complementary component of λ has the structure of an infinite rectangle $\mathbf{R} \times I$ or annulus $\mathbf{S}^1 \times I$. But then the leaves of λ which meet D must be compact; otherwise D together with a half-infinite rectangle in the complementary region lying between them gives an end-compressing disk for $\mathcal{L}$ (see, e.g., [Br1]).

Consequently, both points in $\partial(D \cap \mathcal{L})$ lie on compact loops of λ. But since these loops must be parallel in T, this implies that the leaf of $\mathcal{L}|T$ containing $D \cap \mathcal{L}$ must be a ∂-parallel annulus. This is because $D \cap \mathcal{L}$ splits the leaf of $\mathcal{L}|T$ into a surface with null-homotopic boundary (it is parallel to a null-homotopic loop in T), which must therefore be a disk, because it is a subsurface of a leaf of $\mathcal{L}$. (Note that a priori the leaf could be a boundary-parallel Möbius band, but this would imply that T does not locally separate M, which it does, since M is orientable.)

In what follows it will be important that these ∂-parallel annuli cannot be too 'thin'; there is a positive constant ϵ (depending on $\mathcal{L}$) so that the intersection of the associated solid torus with T has a width of at least ϵ. This constant, which we will call a *monogon number* for $\mathcal{L}$, is determined by an essential branched surface B carrying $\mathcal{L}$. After making B transverse to T, then $N(\tau) = N(B) \cap T$ is a neighborhood of the train track $\tau = B \cap T$; ϵ is half of the minimum distance between distinct horizontal components of $N(\tau)$. Consequently, any two leaves of $\mathcal{L} \cap T \subseteq N(\tau) \cap T$ (such as the edges of a ∂-parallel annulus leaf A of $\mathcal{L}|T$) which are within ϵ of one another must intersect the same I-fiber of $N(\tau)$. But this would then imply that an arc joining the edges of the annulus A is homotopic (via the solid torus it cuts off) rel endpoints to a segment in an I-fiber of $N(B)$. This, however, is impossible, by [GO, Theorem 1(d)].

§2
Warmup: laminations in surfaces

As a warmup to the proof of our main theorem, we will first prove the analogous result for 1-dimensional laminations, and essential simple closed curves, in a surface. The arguments used for the 3-dimensional case will turn out to be remarkably similar to those we will give now; many of them will in fact carry over with no changes whatsoever.

A lamination λ in a surface F is a 1-dimensional foliation of a closed subset of F. A lamination λ is *incompressible* if no compact leaf is a simple closed curve bounding a disk in F or a ∂-parallel arc, and $F|\lambda$ contains no monogons.

Theorem:. *If λ is an incompressible lamination in the orientable surface F and γ is an essential simple closed curve on F, then λ may be isotoped so that $\lambda|\gamma$ is an incompressible lamination in $F|\gamma$, unless λ contains a Reeb annulus.*

A Reeb annulus is the 2-dimensional analogue of a cylindrical component. It is an annulus whose boundary components are leaves of λ; in the interior the leaves of λ are lines which spiral in the same direction toward the two ∂-components (see Figure 2).

Figure 2

Proof: For notational convenience, we will arbitrarily choose a 'positive' side of γ, and refer to the two copies of γ in $F|\gamma$ as γ_+ and γ_-. If we do not care which one we are dealing with, we will call it $\gamma_\pm$.

As in the previous section, it is easy to see that, once λ is made transverse to γ, the only thing which keeps the lamination $\lambda|\gamma$ from being incompressible is the possible existence of ∂-parallel arcs. Because $F|\gamma$ is compact, the set $\mathcal{C}$ of compact arc leaves of $\lambda|\gamma$ is a (closed) sublamination in $\lambda|\gamma$; this is essentially the result of Reeb stability [Re]; the set of simply-connected compact leaves is both open and closed in any closed, saturated subset of a foliation. Similarly, the set $\mathcal{A}$ of ∂-parallel arcs in $\lambda|\gamma$ is a sublamination of $\lambda|\gamma$ (and hence of $\mathcal{C}$); this is most easily seen by noting that the intersection $\mathcal{A}\cap(\gamma_+ \cup \gamma_-)$ is a closed set. For if not, then (since $\mathcal{C}\cap(\gamma_+ \cup \gamma_-)$ is closed) there are ∂-parallel arcs having endpoints which limit on a non-∂-parallel arc α. But any arc having an endpoint sufficiently close to an endpoint of α must be parallel to α. This can be seen by covering α with coordinate charts running from one end to the other; since in each chart the leaves of λ look like parallel arcs, by starting at one end and looking at the arcs in each successive chart that intersect the end of the previous chart, we can stitch together successively thinner collections of arcs all the way from one end to the other. Consequently, all leaves of λ sufficiently close to α *at any point* are parallel to α. So if the end of α is limited upon by ∂-parallel arcs, then they are all eventually parallel to α, so α is also ∂-parallel. Once we know that $\mathcal{A}\cap(\gamma_+ \cup \gamma_-)$ is closed, it follows by the same argument above, using coordinate charts, that $\mathcal{A}$ is closed.

Just as important, the argument above implies that $\mathcal{A}$ and $\mathcal{C}$ are *open* in $\lambda|\gamma$, since any leaf of $\lambda|\gamma$ sufficiently close to a leaf α of $\mathcal{C}$ must be an arc, i.e., in $\mathcal{C}$. In particular, it must be parallel to α, so anything close to a leaf of $\mathcal{A}$ must also be in $\mathcal{A}$.

With these facts we can now describe how we will attempt to isotope λ so that $\lambda|\gamma$ is incompressible in $F|\gamma$. If $\lambda|\gamma$ isn't incompressible, then $\mathcal{A}$ is non-empty; these arcs fall into finitely-many families of parallel arcs. Each family is open and closed in $\mathcal{A}$, and so we can, by an isotopy of λ, push each family of arcs across γ (Figure 3a). To do this right, we should distinguish between two different kinds of families; those that hit γ on the 'left' side, and those that hit on the 'right'. Any two collections of arcs hitting γ on the same side may be pushed across γ at the same time, since they don't interfere with one another. So we can set up a sequence of isotopies of λ, attempting to make $\lambda|\gamma$ incompressible, by first pushing every ∂-parallel arc to the left of γ across γ and then pushing every ∂-parallel arc to the right of γ across γ. Let us call the resulting lamination, which is isotopic to λ, λ_1. We then start all over again, finding and pushing all ∂-parallel arcs, to produce a new lamination λ_2. In this way we obtain a sequence of laminations λ_i, all isotopic to our original λ. The effect on λ_i of each push is essentially to graft together the families of leaves which lie at either end of the ∂-parallel arcs, while erasing the ∂-parallel arcs (Figure 3b).

Figure 3

If we now assume that we cannot isotope λ so that $\lambda|\gamma$ is incompressible, then none of the λ_i is incompressible, and the process of finding and pushing ∂-parallel arcs can be continued indefinitely. What we will do now is analyse how this process can be continued; in the end, we will find that the only mechanism which can do this is a Reeb annulus.

Suppose we wish to push a collection of ∂-parallel arcs across γ. In each parallel family there is an innermost arc (since the collection is closed), and the disk it cuts off from $F|\gamma$ meets γ (and, for notational convenience, $\gamma_\pm$ as well) in an arc α. Pushing the ∂-parallel arcs across γ has the effect of erasing all the points of $\lambda\cap\gamma$ lying in α, while no other points of $\lambda\cap\gamma$ are moved; we call α the *site* of the push. Because the only effect of the isotopy is to erase points, the sets $\mathcal{C}_i = \lambda_i \cap \gamma$ form a nested sequence of closed sets in γ. Since γ is compact, either $\mathcal{C}_i = \emptyset$ for some i (implying that $\lambda_i|\gamma$ is incompressible in $F|\gamma$, because it has no compact arc leaves), or $\cap \mathcal{C}_i = \mathcal{P}$ is a non-empty collection of points that are never moved by any of the isotopies. These are the *stable points* of the sequence of isotopies. Since we assume that at every stage there are some ∂-parallel arcs to push, it follows that the inclusion of each $\mathcal{C}_i$ in its predecessors is proper.

Using a 'monogon number' approach as in the 3-dimensional case, there is a number $\epsilon > 0$ so that every site has length at least ϵ (otherwise, an arc in a leaf of λ would be homotopic, rel endpoints, to a segment of an I-fiber in the neighborhood of an incompressible train track carrying λ, implying that the train track has a monogon). Consequently, there can be only an (a priori) bounded number (less than (length of $\gamma)/\epsilon$) of disjoint sites on γ. Furthermore, since after each round of isotopies, λ_i is disjoint from the sites of all previous pushes, the site for every successive push is either disjoint from, or properly contains, the site of each previous push. This in particular implies that if we wait long enough, no 'new' sites are ever created; we cannot continually have new sites which are disjoint from all previous ones, otherwise we could pack too many intervals of length greater than ϵ into γ. After new sites stop appearing, we can then have sites that coalesce; a new site could contain two or more previously disjoint sites. But this then lowers the number of disjoint sites we can have, so eventually sites stop coalescing. Finally, after this point, sites could 'die'; a site might never be contained in a new site. But this again lowers the number of disjoint sites. So eventually the sites of our pushes stabilize, in a very strong way; beyond a certain stage λ_i, no new sites are created, each site at any stage is eventually contained in a new site, and each new site contains at most one site from any previous stage of the isotopy after λ_i.

Once we reach this point in our sequence of isotopies, if we now look at the endpoints of any set of nested sites, they form a pair of monotone sequences on γ, bounded from above or below by any point of the other sequence. Each sequence therefore converges in γ. The interval between these limit points (which contains both sequences) is in a very natural sense the limit of the nested sequence of sites. The endpoints of these limiting intervals form a finite set, which is in fact precisely the set of points of $\mathcal{P}$, the stable points on γ of the sequence of isotopies, that are limited upon by those points of $\lambda \cap \gamma$ which are removed by the isotopies. This is because they are the only points limited upon by the endpoints of the ∂-parallel arcs. It is this set which we will focus our attention on, to find our Reeb annulus.

We now assume that, in all that follows, we have waited through our sequence of isotopies long enough so that the sites of all of our pushes have stabilized in the sense given above, and further, that the endpoints of our sites are all within our monogon number ϵ of the corresponding ends of the limiting intervals. We redefine the isotoped lamination to be λ, and proceed to number subsequent laminations from this new starting point as λ_1, λ_2, etc. Note that the new lamination still has monogon number ϵ w.r.t. γ, since any arc in our new λ, with endpoints on γ, is homotopic rel endpoints to an arc in the old λ.

Each of our limiting intervals has a well-defined transverse direction; the direction (after this stability of the number and kinds of sites has occurred) that any ∂-parallel arc meeting the interval is pushed to. An interval cannot be assigned different directions at different times, since this would imply the existence of ∂-parallel arcs pushing in opposite directions across the same site. But then if we take the disk cut off by the arc occurring later in the isotopy, and allow it to flow back to the point in time where the earlier one occurs, we find that we have two arcs in λ_i, which at least start on opposite sides of γ, which are both homotopic rel endpoints to a subarc of the limiting interval (see Figure 4). But this requires the existence either of a compact, null-homotopic leaf in λ, or an arc in λ homotopic rel endpoints to an an arc of the limiting interval having length less than ϵ. This can be seen by following the ends of the innermost (i.e., earlier) ∂-parallel arc; they must either close up, giving the null-homotopic loop, or hit γ again at a point inside of the outer (i.e., later) arc, giving a null-homotopic arc cutting off an arc from γ of length less than ϵ.

Figure 4

One of the more important questions we must answer is: where do the ∂-parallel arcs that we keep needing to push across γ come *from*? If we think of a ∂-parallel arc α in some $\lambda_i|\gamma$, we can then imagine letting it flow *backwards* through our isotopies, back to λ. Its endpoints will not move, since each isotopy does not move any of the points of $\lambda\cap\gamma$ that it does not erase. Since α is a compact arc, its image in λ intersects γ only finitely-many times. The image of α therefore consists of a finite number of compact arcs in $\lambda|\gamma$. Since α is not a ∂-parallel arc for λ (otherwise we would have pushed it at the beginning), the image of α (which we will still call α, for convenience) must intersect γ at points in the interior of α. In order for these points to be removed, they must each, therefore, eventually be contained in the ends of their own ∂-parallel arcs. If we imagine how the disk Δ that α cuts off from $\lambda_i|\gamma$ changes as we flow back in the isotopy, we get our basic picture of the isotopy process; α gets cut up into finitely many arc leaves of $\lambda_{i-1}|\gamma$; every other arc is ∂-parallel. Δ is cut up by γ into a disk lying between finitely many leaves of $\lambda_{i-1}|\gamma$, and finitely many disks bounded by ∂-parallel arcs (Figure 5). Running this picture forward then give us our picture of how new ∂-compressing arcs are formed. They form from disk regions Δ lying between finitely many arc leaves of $\lambda_{i-1}|\gamma$; all but one of the arcs of $\Delta\cap\gamma$ are sites for pushes of ∂-parallel arcs of $\lambda_{i-1}|\gamma$; the remaining arc becomes the new site for the newly-created ∂-parallel arc. Put differently, the transverse orientations for all but one of the sites around the edge of Δ point into Δ; the remaining site points out of Δ and marks the site where the newly created ∂-parallel arc will reside.

Figure 5 Figure 6

Some of the leaves of $\partial\Delta$ in $\lambda_i|\gamma$ may not have existed in our original $\lambda|\gamma$; they may have been pushed there sometime after the stabilization of pushing sites occurred. This push must have occurred from *outside* of Δ, however; a push from inside of Δ would violate our monogon number condition (Figure 6). So this leaf arrived in its place after a push from a site other than those meeting Δ, allowing us to enlarge Δ by adding another, inward pointing site. Since there are only a finite number of sites, however, this process cannot continue forever. As we continue to look at the ∂-parallel arcs which appear at our outward pointing site, and look at the disks Δ which created them, and evolve the leaves in their boundaries back in time, we must eventually find ourselves always reconstructing the same disk. This is because new ∂-parallel arcs at the same site are wider than the ones that come before, so the arc leaves used to build its disk Δ must get out of the way of the arc leaves used to build earlier disks, as we flow backwards in time. In other words, the disk from time $2N$, say, evolved back in time to time N, must contain the disk, from the same site, which we find at time N. So now when we evolve *both* disks backwards in time, the $2N$-disk, evolved to time 0, contains the N-disk, evolved to 0.

So, eventually, the disk associated to each site stabilizes, at least in terms of which incoming sites it associates to its outgoing site. Any disk coming from further along in the isotopy contains, but is otherwise parallel to, this largest disk. Taken together, these disks combine along each of the sites to give a subsurface of F. But an easy counting argument shows that this subsurface must be an annulus.

Each disk has some number $n \geq 1$ of inward-pointing sites, and one outward-pointing site. Each site is the outward-pointing site of one of these disks Δ. It is not necessary for a site to be the inward-pointing site of a disk; such a site would represent a place where ∂-parallel disks go off to 'die' (Figure 7). But since each disk requires at least as many inward sites as outwards, and there are precisely the same number of inward sites as outward sites, the inescapable conclusion is that each disk has precisely one inward site, and there are no sites where arcs die. More precisely, if we have k disks (i.e., k outward-pointing sites) each having $n_i \geq 1$ inward sites, and $\Sigma n_i \leq k$ (since each inward site is associated to at most (some may die) one disk, hence at most one outward site), then each n_i=1, and $\Sigma n_i = k$, i.e., each disk has exactly one inward site, and each inward site is associated to a disk.

Figure 7 Figure 8

Each disk is therefore a rectangle, and they must be glued end-to-end to form an annulus (Figure 8). These gluings are of course approximate, since the leaves in the boundaries of the disks are not stable; they are pushed away in later rounds of isotopies. However, what this does tell us is that there is no growth of these disks as we evolve them backwards in time; the arc leaves in all of the $\partial\Delta$ are present in our original $\lambda|\gamma$. If we look at the leaves of $\lambda|\gamma$ that come out of the endpoints of our limiting intervals, since these endpoints are limited upon by the ends of compact leaves of $\lambda|\gamma$, these leaves must themselves be compact, and therefore parallel to the leaves in the boundaries of our disks Δ. These compact leaves glue together end-to-end to give compact loops, which are leaves of λ. Together they bound an annulus (which our rectangles above were trying to approximate). The leaves of λ inside of this annulus consist of ∂-parallel arcs of $\lambda|\gamma$ together with the sides of the rectangles, which together form the non-compact leaves of a Reeb lamination in the annulus.

Consequently, if λ can never be isotoped so that $\lambda|\gamma$ is incompressible, then if we follow the procedure outlined above, eventually what we see is a collection of ∂-parallel arcs spiralling around F between a pair of (stable - no isotopy moves them) parallel circle leaves of λ. This is precisely the Reeb annulus which we seek. We should note that, technically, the two boundary curves could be the same curve, if the Reeb annulus actually fills up the surface F. This would imply that F is a torus, since it is an annulus with its two ends glued together. Our lamination consists only of a Reeb annulus, whose non-compact leaves are limiting on the same closed loop, from both sides. ∎

§3
The 3-dimensional case: killing ∂-parallel annuli

We now return to the proof of our main theorem. What we shall see is that we have a nearly identical situation to the case of a lamination in a surface. The only difference is that everything is 'crossed with S^1'. The basic idea is that our main concern, ∂-parallel annuli, are ∂-parallel arcs crossed with S^1, and the phenomenon we seek, a cylindrical component, is a Reeb annulus, crossed with S^1.

It is well-known that the collection of compact leaves of a codimension-one foliation form a closed subset of the ambient manifold (see, e.g., [Co]). In fact, the set of leaves homeomorphic to a given compact surface fall into finitely-many parallel families in the ambient manifold, each of which (after including some leaves that the surface 2-fold covers, as, for example, with annulus leaves in the neighborhood of a Möbius band leaf) form a closed set as well. The arguments apply equally well to a lamination, and so we may assume that the annular and Möbius band leaves of $\mathcal{L}|T$ fall into finitely many parallel families whose union is a closed set (and therefore, the collection of ∂-parallel annuli do, too). The collection of Möbius band leaves is therefore finite. If we, a priori, split $\mathcal{L}$ along the finitely-many leaves containing these Möbius bands, we can, at the outset, assume that $\mathcal{L}|T$ has no Möbius band leaves.

What we will do now is isotope $\mathcal{L}$ by pushing these annuli, one family at a time, across the solid tori that they split off of $M|T$, and through T to the other side. The key observation that we will need to make is that while this isotopy may create new Reeb annuli, the 'slope' of the boundary curves of the annuli cannot change. This fact will allow us to use a loop in T geometrically dual to (i.e., having a single transverse intersection with each of) the boundary loops of our annuli as a bookkeeping device to keep track of these isotopies.

To see that this is so, choose such an initially dual curve γ in T, and assume that $\mathcal{L}$ has been isotoped so that $\lambda = \mathcal{L} \cap T$ meets γ as tightly as possible, i.e., γ meets all closed loops in λ in a single point, and all non-compact leaves travel around T so that as we travel along the leaf, we always meet γ on the same side, except for Reeb-type leaves, which switch sides exactly once (see Figure 9).

Now consider the effect on λ when we push a family of ∂-parallel annuli through T (Figure 10). We can think of this simply as pushing the innermost annulus A of the family through T. The effect on λ is to erase all of the leaves of λ meeting the solid torus M_0 that A cuts off of $M|T$ (including the boundary curves of A), and then splicing together any leaves of λ that limit on ∂A from outside of $M_0 \cap T = A_0$. In particular, any compact leaves of λ lying outside of A_0 remain unchanged. So if pushing A is going to result in ∂-parallel annuli of a different slope being created, then there must be no compact leaves lying outside of A_0. So there are only non-compact leaves, some of which limit on ∂A from outside of A_0. But then the new ∂-parallel annulus that is created must cross where ∂A used to run, and it is easy (see Figure 11) to splice together the new ∂-compressing disk and an infinite rectangle between the leaves that were limiting on ∂A to create an end-compressing disk for $\mathcal{L}|T$, and hence for $\mathcal{L}$, a contradiction. For a similar reason, no ∂-parallel annulus of the *same* slope can intersect the site where an old annulus was pushed, because we could just as easily build an end-compressing disk in this situation, as well. This tells us that in order for there to be new ∂-parallel annuli for us to push, there must always be more compact loops of λ lying outside of A_0. So the new leaves of λ created by the push, which were made by splicing together the leaves limiting on the ends of A_0, are all non-compact; no new compact leaves are ever created.

Figure 9 Figure 10 Figure 11

Now that we know that the slope of a ∂-parallel annulus is (for lack of a better phrase) an invariant of $\mathcal{L}$ and T, we can set up a general procedure aimed at sequentially killing off any boundary-parallel annuli that we can find. To build our isotopy, we find all of the boundary parallel annuli in $\mathcal{L}|T$; their intersection with T forms a (closed) sublamination of $\mathcal{A}$. These annuli fall into a finite collection of parallel families in M; the loop γ intersects the annulus that each family cuts off from T in an arc that has length greater than ϵ. The families in turn are of one of two types, depending upon whether it intersects T on the positive or negative side. We now, in turn, push each family of annuli across T, by pushing all of the annuli lying on the positive side across (this can be done all at once, because they do not interfere with one another), then pushing all of the annuli on the negative side, and then starting the process all over again. At each stage we create a 'new' lamination $\mathcal{L}_i$, which is isotopic to our original lamination $\mathcal{L}$. We assume that at no stage in the isotopy is $\mathcal{L}_i|T$ essential in $M|T$, i.e., at every stage we find new ∂-parallel annuli to push.

As in the case of a lamination in a surface, we can choose a loop γ in T which will track the sites of our pushes. In this case, this means that γ meets $\lambda = \mathcal{L} \cap T$ as tautly as possible; in particular, γ meets every closed loop of λ in exactly one point. If we isotope $\mathcal{L}$ by pushing a ∂-parallel annulus A through T, the loop γ will nearly be as taut w.r.t the lamination in T (which we will call λ_1, etc.) resulting from the intersection of the isotoped lamination (which we will call $\mathcal{L}_1$, etc.) with T. In particular, it will be taut w.r.t. the closed loops of λ_1, because no new closed loops have been created.

§4
Building the cylindrical component(s)

As with the case of a lamination in a surface, the most important question which we need to answer at this point is: where will the *new* ∂-parallel annuli A_i for our isotoped lamination $\mathcal{L}_i$ come from after we have killed the old ones for $\mathcal{L}_{i-1}$? Any new leaves of $\mathcal{L}_i$, since they weren't present before the isotopy, must arise from leaves of $\mathcal{L}_{i-1}$ which were spliced together in the neighborhood of ∂A when A was pushed through T. Since our new ∂-parallel annulus A_i, when we allow it to flow back to $\mathcal{L}$, is compact, it intersects T only in closed loops. These loops are essential in T and hence in A_i. The annulus A_i is therefore cut up into annuli; alternate annuli are ∂-parallel, and were pushed across T in order to create A_i. Iterating this procedure for later and later appearing ∂-parallel annuli, we can conclude that new ∂-parallel annuli come from the annulus leaves of λ, by the sequence of gluings which occur when ∂-parallel annuli are pushed. In particular, all of the annuli which together make up any ∂-parallel annulus which we will push are present in our original lamination $\mathcal{L}$. This is directly analogous to the 2-dimensional case. In particular, since a Möbius band leaf of some $\mathcal{L}_i|T$ would require a Möbius band leaf of $\mathcal{L}|T$ among the leaves glued together to create it, none of the laminations $\mathcal{L}_i|T$ have a Möbius band leaf.

Consider now the ∂-parallel annuli for the lamination $\mathcal{L}_i|T$. The innermost annulus A in each family of ∂-parallel annuli cuts off a solid torus M_0 from $M|T$ (which contains all of the annuli in its family); its intersection with the loop γ is an arc α, which we will call (in conscious imitation of §2) the *site* of the push corresponding to the family of ∂-parallel annuli. The site, together with an arc in A, bounds a (∂-compressing) disk Δ in M_0; M_0 can be thought of as $\Delta\times S^1$. This disk is the analog of the disks Δ from §2, as well.

Once a collection of annuli has been pushed, $\mathcal{L}_i\cap T = \lambda_i$ is disjoint from the site (and in fact from the annulus in T that the solid torus meets), and so, as before, all further sites are either disjoint from, or contain, the site α. Using the same monogon number argument, we can therefore wait until the collection of sites along γ stabilizes; after this point, no new sites are created, every site is eventually contained in a new site, and every new site contains at most one site from any previous isotopy. We can also then wait until the ends of each site are within the monogon number ϵ of the points that these endpoints are limiting on. We then, as before, redefine this stage of the isotopies to be our actual starting point in the collection of isotopies. As before, each of our sites has a well-defined transverse orientation, the direction in which the ∂-parallel annuli are pushed across the site. A site cannot be assigned different orientations at different times, since otherwise we arrive at the same situation as Figure 5, crossed with S^1, leading, again, to a violation of our monogon number condition.

If we now imagine evolving the solid torus M_0 cut off by a ∂-parallel annulus of $\mathcal{L}_i|T$ back through the isotopy, we see a similar pattern to the 2-dimensional case emerge: the annulus A, flowing back to $\mathcal{L}_{i-1}$, intersects T in a finite number of loops, cutting A into parallel annuli; alternate annuli are ∂-parallel, and are pushed in $\mathcal{L}_{i-1}$ to create A in $\mathcal{L}_i$. The solid torus M_0 pushed back to $\mathcal{L}_{i-1}|T$, is therefore cut into finitely-many solid tori, all but one of which is the site for a push of $\mathcal{L}_{i-1}$(Figure 12). The remaining solid torus has boundary alternately consisting of annuli in T and annulus leaves of $\mathcal{L}_{i-1}|T$, and has the structure of a disk, Δ, whose boundary alternately lies in T and in annular leaves of $\mathcal{L}_{i-1}|T$, crossed with S^1. Some number $n\geq 1$ of these annuli in T represent inward-pointing sites, where the ∂-parallel annuli of $\mathcal{L}_{i-1}|T$ are push, while the remaining annulus is the outward pointing site representing the new ∂-parallel annulus that is being created. In other words, the situation here is the same one we had for a lamination in a surface, crossed with S^1.

Figure 12

As before, these disks Δ can evolve in only one way, as we look backwards in the sequence of isotopies. The annulus leaves of $\mathcal{L}_{i-1}|T$ meeting the boundary of Δ might not belong to $\mathcal{L}|T$, but were pushed there. Such leaves could only be pushed from outside of Δ; otherwise we have the exact same situation as in Figure 6, crossed with S^1, which again violates our monogon number condition. Consequently, as we look at ∂-parallel annuli further and further along in the isotopy, and look at how the solid torus cut off from it evolves backwards to the beginning of the isotopies, we again get the same sort of stability as in the surface case; eventually, the disks Δ associated to each become 'parallel' to one another. Then the same counting argument as in the 2-dimensional case leads us to conclude that these stabilized disks Δ are all rectangles R_i, and so our backwards-evolved solid tori all have the form of $R_i\times S^1$. As before, the annulus leaves of $\mathcal{L}|T$ forming the opposite sides of these rectangles therefore limit upon a finite number of annulus leaves of $\mathcal{L}|T$, which stitch together to form a pair of torus leaves of $\mathcal{L}$. (In limiting cases, this could be a single torus leaf, limited upon on both sides, or a Klein bottle leaf, cut by T into annuli; only a single leaf would occur, because it would be one-sided, so the annuli would have to limit on it from 'both' sides.) The solid tori $R_i \times S^1$ would stitch together to form an annulus crossed with S^1, approximating the region lying between the two torus leaves. As in the 2-dimensional case, the lamination in this cylindrical region would look like a Reeb lamination in an annulus, crossed with S^1. In other words, these two torus leaves of $\mathcal{L}$ bound the cylindrical component that we seek. ∎

References

[BC] S. Bleiler and A. Casson, *Automorphisms of surfaces, after Nielsen and Thurston*, Cambridge Univ. Press, New York, 1988.

[BNR] M. Brittenham, R. Naimi, and R. Roberts, *Graph manifolds and taut foliations*, J. Diff. Geom. **45** (1997), 446-470.

[Br1] Brittenham, M., *Essential laminations in Seifert-fibered spaces*, Topology **32** (1993), 61-85.

[Br2] ———, *Essential laminations in Seifert-fibered spaces: boundary behavior*, to appear in Top. Appl..

[Br3] ———, *Essential laminations and deformations of homotopy equivalences: from essential pullback to homeomorphism*, Top. Appl. **60** (1994), 249-265.

[Br4] ———, *Essential laminations in non-Haken 3-manifolds*, Top. Appl **53** (1993), 317-324.

[Cl] W. Claus, *Essential laminations in closed Seifert-fibered spaces*, thesis, Univ. of Texas at Austin (1991).

[Co] L. Conlon, *Foliations of codimension-one*, lecture notes from course at Washington Univ., 1989.

[GO] D. Gabai and U. Oertel, *Essential laminations in 3-manifolds*, Annals of Math. **130** (1989), 41-73.

[JN] Jankins and W. Neumann, *Rotation numbers of products of circle homeomorphisms*, Math. Ann. **271** (1985), 381-400.

[Na] R. Naimi, *Foliations transverse to fibers of Seifert manifolds*, Comment. Math. Helv. **69** (1994), 155-162.

[Ro] Roussarie, R., *Plongements dans les variétés feuilletées et classification de feuilletages sans holonomie*, Prepubl. IHES **43** (1973), 101-142.

[Th] Thurston, W., *Foliations of 3-manifolds which are circle bundles*, thesis, U. Cal. at Berkeley, 1972.

[Wu] Y.-Q. Wu, *Sutured manifold hierarchies, essential laminations, and Dehn surgery*, to appear in J. Diff. Geom.

Department of Mathematics, Box 305118, University of North Texas, Denton, TX 76203-5118
E-mail address: `britten@unt.edu`

Department of Mathematics, Washington University, St. Louis, MO 63130
E-mail address: `roberts@math.wustl.edu`