Essential Laminations and Haken Normal Form

Mark Brittenham

We show that if a 3-manifold M contains an essential lamination, then for any triangulation of M there exists an essential lamination which is in Haken normal form with respect to that triangulation.

0. Introduction.

The notion of (Haken) normal form w.r.t. a triangulation of a 3-manifold traces back to Kneser's work in the 1930's on surfaces in 3-manifolds. Haken studied it extensively in the 1960's, and showed [8] how to use it to create finite algorithms for the determination of various properties of embedded surfaces. This has since culminated, in the work of Jaco and Oertel [10], in an algorithm to determine if an irreducible 3-manifold is a Haken manifold, i.e., if it contains a 2-sided incompressible surface.

In [7] a generalization of the incompressible surface, the essential lamination, was introduced. There it was shown that a 3-manifold M containing an essential lamination has some of the same desirable properties of a 3-manifold containing an incompressible surface, the most notable property being that M has universal cover $\mathbf{R}^3$. Since then, it has also been shown [6] that, in some sense, 'most' 3-manifolds contain essential laminations.

The purpose of this paper is to prove a Haken normal form result for essential laminations.

The reader is referred to [7] for definitions and basic properties concerning essential laminations. In this paper the word 'lamination' will mean a lamination which is carried by a branched surface, i.e., it has 'air' between its leaves. Since we will ultimately be interested only in the existence of an essential lamination with certain properties, this additional restriction will cause no difficulties; we can 'blow air' between leaves the leaves of a foliation (see [7]) to obtain a lamination in our sense.

Generalizing the definition for a compact surface [8], we will say that a lamination $\mathcal{L}\subseteq$M is in *Haken normal form* w.r.t. a triangulation τ of M if $\mathcal{L}$ is in general position w.r.t. τ, and for every 3-simplex $\Delta^3$of τ, $\mathcal{L}\cap\Delta^3$ is a lamination consisting of compact disks, each of which meets the 1-skeleton of Δ_i^3, and such that each disk D of $\mathcal{L} \cap \Delta^3$ meets each 1-simplex of Δ^3 at most once.

Then we have the following theorem:

Normal Limit Theorem: *If* M^3 *contains an essential lamination* $\mathcal{L}$*, and* τ *is a triangulation of* M*, then there exists an essential lamination* $\mathcal{L}_0$ *in* M *which is in Haken normal form w.r.t.* τ*.*

Corollary: M *contains an essential lamination iff it contains one which is carried with full support by one of a finite, constructible, collection of normal branched surfaces.*

In general, the lamination $\mathcal{L}_0$ is not isotopic to $\mathcal{L}$; it is, when $\mathcal{L}$ is measured [9], or, more generally, has no holonomy [3]. But it does arise out of a 'limit' of isotopic copies of $\mathcal{L}$, as described below.

Ultimately, we would like to see an algorithm found to determine if an irreducible 3-manifold contains an essential lamination, as in [10]. This paper can be taken to be a first step in that direction. In a sequel [2], we show how this result can be used to prove the same result for more general cell decompositions of 3-manifolds. From this point of view, a key step in developing an algorithm would be to replace 'normal branched surfaces' in the corollary above with 'normal *essential* branched surfaces'.

I have been told that David Gabai has also proved a version of the main result of this paper.

The proof of the theorem is in broad outline very similar to the arguments found in [1]. We shall describe an algorithm for performing a (typically infinite time) isotopy of $\mathcal{L}$, controlled along the intersection of $\mathcal{L}$ with the 1-skeleton $\tau^{(1)}$ of τ. After identifying a collection of points which are left fixed by all of the isotopies, we then study what happens when we take the 'limit' of these isotopies. At this point the proof diverges markedly from [1]; seeing that pieces of the lamination stabilize around these fixed points is easier, but in the situation encountered here, the union of these stable pieces is not a lamination - it is a disjoint union of 1-to-1 immersed surfaces, but it is not, in general, a closed set.

Making the limit into a lamination requires cutting and pasting a finite number of the leaves in the limit; then by passing to a sublamination, to avoid possible compressible tori, one finds the lamination guaranteed by the theorem.

This result, as well as those of [1] and [2], illustrates a technique for attacking problems whose solutions for incompressible surfaces relies on induction (on the number of points of intersection with a 1-dimensional object). The technique used here is, in fact, simply a different way of thinking about this induction process; instead, it is something we might call 'eventual stability under a sequence of isotopies'. For a compact surface S this amounts to the same thing as induction, since (in the case of this paper, for instance) once the number of points of $S \cap \tau^{(1)}$ is at a minimum, the entire surface must eventually stabilize around these points to a normal surface. Such reformulations of methods which, for incompressible surfaces, rely on the *compactness* of the object, rather than their *closedness*, will, it is certain, play an important role in furthering the development of the theory of essential laminations.

The author would like to thank the referee for his/her many useful comments.

1. The isotopies.

We will assume that the reader is familiar with the procedure for putting an incompressible surface in an irreducible 3-manifold M into Haken normal form w.r.t. a triangulation of M, see [8] or [11]; and confine our discussion here to dealing with the difficulties that arise when adapting this process to essential laminations. Since if $\mathcal{L}$ contains a compact leaf, we can apply Haken's algorithm to it to get an essential surface in normal form, we can assume, for convenience, that $\mathcal{L}$ contains no compact leaves. The only terminology for essential laminations which we use that cannot be found in [7] is that of a *monogon number* for a lamination w.r.t a 1-complex K; this is a number ϵ so that for some fixed branched surface B, with branched-surface neighborhood N(B) carrying $\mathcal{L}$with $K \cap N(B)$=a collection of I-fibers of N(B), any two points of $N(B) \cap K$ which are within ϵ of one another are contained in the same I-fiber of N(B). We shall routinely abuse notation by using the same symbol to represent an isotopy of a lamination and the embedding of the lamination that results at the end of the isotopy.

Given an essential lamination $\mathcal{L}$, we can put it in general position w.r.t. a triangulation τ, by putting a branched surface B carrying $\mathcal{L}$ in general position w.r.t. τ, and then embedding $\mathcal{L}$ in a fibered neighborhood N(B). Consider the lamination $\lambda=\mathcal{L}\cap\partial\Delta^3\subseteq\Delta^3$, where Δ^3 is a 3-simplex of τ. This 1-dimensional lamination consists of a finite number of parallel families of loops in $\partial\Delta^3_i=S^2$; λ cannot contain a non-compact leaf, because the existence of such a leaf would either imply the existence of spiralling (non-trivial holonomy) around a null-homotopic loop in $\mathcal{L}$, or a monogon for λ (which would give an end-compressing disk for $\mathcal{L}$). By a process entirely analogous to §2.b of [1], we can by an isotopy I of $\mathcal{L}$ arrange that $I(\mathcal{L})\cap\Delta^3_i$ is a (finite) collection of parallel families of disks, and $I(\mathcal{L})\cap\partial\Delta^3_i\subseteq\mathcal{L}\cap\partial\Delta^3_i$. This process is most easily envisioned as a surgery on all of the parallel families in λ, using a surgery disk parallel to either disk that the loops of λ bound in $\partial\Delta^3_i$; one then throws away any 2-spheres that this surgery creates. As in [1], this surgery process can be realized as an isotopy.

Rename $I(\mathcal{L})=\mathcal{L}$; $\mathcal{L}\cap\partial\Delta^3_i=\lambda$ still consists of a finite number of families λ_i of parallel loops, and each of these parallel families further break down into finite families λ_{ij} of loops which are inessentially isotopic to one another in the cell decomposition of $\partial\Delta^3_i$, i.e., any two loops cobound an annulus in $\partial\Delta^3_i$ which is made up of rectangles in each of the faces of Δ^3_i; see Figure 1a.

By boundary-surgeries, realized as an isotopy of $\mathcal{L}$, we can now, working family by family, arrange that each loop of $\mathcal{L}\cap\partial\Delta^3_i$ meets each 1-simplex of $\partial\Delta^3_i$ at most once (see Figure 1b). Finally, by isotopy one pushes any disk D of $\mathcal{L}\cap\Delta^3_i$ which now entirely misses the 1-skeleton of Δ^3_i out of Δ^3_i.

Figure 1

Using this isotopy I, we have arranged that $I(\mathcal{L})\cap\Delta^3_i$ consists of a collection of normal compact disks. Also, by the structure of the isotopy (see Figure 1b), we have the following very useful property for I: $I(\mathcal{L})\cap\tau^{(1)}\subseteq\mathcal{L}\cap\tau^{(1)}$, *and the isotopy* I *is fixed on all points of* $I(\mathcal{L})\cap\tau^{(1)}$.

The idea now is to string isotopies like the one above together to give an 'isotopy' for $\mathcal{L}$, which attempts to put $\mathcal{L}$ into normal form.

A triangulation τ of M has a finite number of 3-simplices which we number $\Delta^3_1,\dots,\Delta^3_n$. By the above, we can perform an isotopy I_1 of $\mathcal{L}$ so that $I_1(\mathcal{L})\cap\Delta^3_1$ consists of normal disks. Then we build an isotopy $I_{2,1}$ of $I_1(\mathcal{L})$ so that $I_{2,1}(I_1(\mathcal{L}))\cap\Delta^3_2\equiv I_2(\mathcal{L})\cap\Delta^3_2$ consists of normal disks. Notice, however, that since $I_{2,1}$ may require boundary-surgeries, which push problems out of Δ^3_2, and hence possibly into Δ^3_1, we can no longer insure that $I_2(\mathcal{L})\cap\Delta^3_1$ consists of normal disks. Continuing inductively, working cyclically through the 3-simplices Δ^3_i, we can build an isotopy $I_{r,r-1}$ of $I_{r-1}(\mathcal{L})$ so that $I_{r,r-1}(I_{r-1}(\mathcal{L}))\cap\Delta^3_i\equiv I_r(\mathcal{L})\cap\Delta^3_i$ is a collection of normal disks, for $r\equiv i(\mathrm{mod}\ n)$. Notice that this collection of isotopies I_r satisfies $I_r(\mathcal{L})\cap\tau^{(1)}\subseteq I_{r-1}(\mathcal{L})\cap\tau^{(1)}$, and the isotopy I_r is fixed on all points of $I_r(\mathcal{L})\cap\tau^{(1)}$. We will adopt the notation that, for $r\geq s$, $I_{r,s}=I_r\circ I_s^{-1}=I_{r,r-1}\circ I_{r-1,r-2}\circ\ldots\circ I_{s+1,s}$, so $I_{r,s}\circ I_s=I_r$.

If for some s, $I_s(\mathcal{L})$ is in normal form w.r.t. τ, then for all $r\geq s$, $I_{r,s}=\mathrm{Id}$, since the isotopies move only portions of $\mathcal{L}$ which are not in normal form. We can then set $\mathcal{L}_0=I_s(\mathcal{L})$, giving the lamination required for the theorem. In what follows now we will therefore assume that for no r is $I_r(\mathcal{L})$ in normal form; note that this is equivalent to $I_{r+n}(\mathcal{L})\cap\tau^{(1)}\neq I_r(\mathcal{L})\cap\tau^{(1)}$, for all r. Now we will observe how these isotopies effect $\mathcal{L}$, as r continues to get larger and larger. What we will find is that pieces of the $I_r(\mathcal{L})$ begin to stabilize, becoming fixed under all further isotopies. These pieces will form the 'core' of the essential lamination $\mathcal{L}_0$ of the theorem.

2. Stability.

Given the sequence of isotopies described in the previous section, consider the sets $I_r(\mathcal{L})\cap\tau^{(1)}=C_r$. Each is a closed set, and is non-empty; for if $C_r=\emptyset$, then it is easy to see that $I_r(\mathcal{L})\cap\partial\Delta^3_i$ is a collection of loops all in the interiors of the faces of $\partial\Delta^3_i$, for all i. But then $I_{r+n}(\mathcal{L})\cap\partial\Delta^3_i$ would be empty, for all i, implying that the essential lamination $I_{r+n}(\mathcal{L})$ is contained in a ball, which is impossible (see §5 below).

Since $C_{r+1}\subseteq C_r$. it follows that $C=\cap C_r$ is a non-empty closed set; it is the intersection of a nested sequence of non-empty closed sets in the compact set $\tau^{(1)}$. C is by construction the set of points of $I_0(\mathcal{L})\cap\tau^{(1)}=\mathcal{L}\cap\tau^{(1)}$ which are fixed by all of the isotopies I_r; it is the set of *stable points* of the isotopy I_r.

The points of C represent the 'seeds' of the lamination $\mathcal{L}_0$; the next lemma begins to show how $\mathcal{L}_0$ will grow out of these points.

Lemma: *Given* $x\in C$, *and a 2-simplex* Δ^2 *of* τ *with* $x\in\partial\Delta^2$, *then for some* s *the arc* α_s *of* $I_s(\mathcal{L})\cap\Delta^2$ *containing* x *is stable, i.e.,* α_s *is fixed under all isotopies* $I_{r,s}$, $r\geq s$.

Proof: It suffices to show that for some s, the arc α_s has its other endpoint also in C; the arc must then be stable, because by construction the only way an arc in a loop of some $I_r(\mathcal{L})\cap\partial\Delta_i^3$ can move is if one of its endpoints is removed from $I_r(\mathcal{L})\cap\tau^{(1)}$. So we will assume that for no s is the other endpoint x_s of the arc α_s contained in C, i.e., for some $r\geq s$, $x_s\notin I_r(\mathcal{L})\cap\tau^{(1)}$, and show how this leads to a contradiction.

The first thing we must understand is how the arc α_s evolves under further isotopies. Because one endpoint, x, is fixed, α_s can change only by 'splicing'; a small neighborhood of x_s in α_s is removed, and an arc in a small neighborhood of $\partial\Delta^2$ is added to it, joining α_s to another similarly shortened arc. This could happen a finite number of times. The union of α_s and these arcs is α_r (see Figure 2). This is described as a 'conservative isotopy' in [1].

Figure 2

Now our hypothesis implies that splicing must be occurring infinitely often, which in turn implies that the number of points in $\alpha_r\cap\gamma$ must be getting arbitrarily high, where γ is the 'neighbor loop' parallel to $\partial\Delta^2$ which cuts off an annulus where all of this splicing is occurring (see Figure 2). It therefore follows that, for some r, two points of $\alpha_r\cap\gamma$ must lie within ϵ of one another along γ, where ϵ is a monogon number for $\mathcal{L}$ w.r.t. γ. But then there are also points a, b of $\alpha_r\cap\gamma$ consecutive along γ with this property. Then the short arc of γ between a and b together with the arc of α_r between them form a simple loop in Δ^2, bounding a disk D in Δ^2. Reversing the isotopies carried out so far, and applying them to the arc of α_r together with the disk D, exhibits a homotopy of a vertical arc in N(B), rel its boundary, into a leaf of $\mathcal{L}$. This, however, contradicts [7, Theorem 1(d)], which says that such homotopies are impossible. ∎

Lemma: *Given* $x\in C$ *and* Δ_i^3 *with* $x\in\partial\Delta_i^3$*, then there exists an* s *such that* x *is contained in a stable normal disk* $\Delta_{x,i}$ *of* $I_s(\mathcal{L})\cap\Delta_i^3$*, i.e.,* $\Delta_{x,i}$ *is fixed under all further isotopies* $I_{r,s}$*,* $r\geq s$.

Proof: It suffices to show that, for some s, x is contained in a disk $\Delta_{x,i}$ of $I_s(\mathcal{L})\cap\Delta_i^3$ with $\Delta_{x,i}\cap\tau^{(1)}\subseteq C$, since then $\partial\Delta_{x,i}$ is stable. Then since a disk can be moved by the isotopies only if its boundary moves, $\Delta_{x,i}$ must be stable.

But this is straightforward, given the previous lemma. The points x lies in two faces of $\partial\Delta^3_i$; pick one, Δ^2. For some s, the arc α_s of $I_s(\mathcal{L})\cap\Delta^2$ containing x is stable; its other endpoint $x_s$is in C. Set $x_s=y_1$; it is contained in another 2-simplex Δ^2_1 of $\partial\Delta^3_i$, and we can apply the previous lemma to y_1 in this other 2-simplex to get $y_2\in C$ joined to y_1 by an arc of some $I_r(\mathcal{L})\cap\Delta^2_1$. The point y_2 is now contained in still another 2-simplex of $\partial\Delta^3_i$, and this process can be continued.

But at some point in every n isotopies the union of these arcs, because they are stable, is contained in the boundary of a normal disk D of some $I_r(\mathcal{L})\cap\Delta^3_i$. This disk D meets $\tau^{(1)}$ in at most 4 points, and therefore meets C in at most 4 points. Consequently, the chain above must eventually close, hence $y_i=x$, for some $i\leq 4$. So eventually x is contained in a stable normal loop of some $I_r(\mathcal{L})\cap\partial\Delta^3_i$, which in turn bounds a stable normal disk in $I_{r+n}(\mathcal{L})\cap\Delta^3_i$. ∎

3. The intermediate lamination $\mathcal{L}'$.

We now have that for each $x\in C$, and each i such that $x\in\partial\Delta^3_i$, there is an s for which x is contained in a stable normal disk $\Delta_{x,i}$ of $I_s(\mathcal{L})\cap\Delta^3_i$. Now consider the set $\bigcup\Delta_{x,i}=X$, where the union is taken over all $x\in C$ and all i as above. X is a union of disks which are normal w.r.t. τ, and $X\cap\tau^{(1)}=C$.

X can also be thought of as a disjoint union of 1-to-1 immersed surfaces in M; 2-disk neighborhoods for points in X can easily be seen by examining the cases $x\in\mathrm{int}(\Delta^3_i)$, $x\in\partial\Delta^3_i\setminus\tau^{(1)}$, and $x\in\tau^{(1)}$. X even has a local product structure, because it falls into blocks of parallel normal disks. Moreover:

Lemma: *Every leaf of* X *is* π_1*-injective in* M.

Proof: Let L be an leaf of X, and γ a loop in L null-homotopic in M. By general position we may assume that γ is self-transverse and immersed in L, and misses the 1-skeleton of τ. γ is compact, so it meets only finitely many of the normal disks D_i comprising X. By choosing an s large enough so that the disks that γ meets are stable normal disks for I_s, we can assume that $\gamma\subseteq I_s(\mathcal{L})$. Consequently, $\gamma\subseteq L'\subseteq I_s(\mathcal{L})$ for some leaf L' of $I_s(\mathcal{L})$. L' is π_1-injective in M, so there is a null-homotopy $F:D^2\to L'$ with $F|_{\partial D^2}=\gamma$.

Consider $\gamma\subseteq L'$. A small regular neighborhood of γ, $N(\gamma)\subseteq L'$, separates L' into a finite number of connected components; $L'\setminus\mathrm{int}(N(\gamma))=L'_1\cup\ldots\cup L'_r$. Some collection $L'_1,\ldots,L'_s$ of them are compact, and by setting $S=L'_1\cup\ldots\cup$

$L'_s \cup N(\gamma)$, we have that S is compact, $\gamma \subseteq S$, and $L' \setminus \mathrm{int}(S)$ has no compact components. It then follows from standard facts (since ∂S is π_1-injective in $L' \setminus \mathrm{int}(S)$) that $S \subseteq L'$ is π_1-injective in L', and so we can assume that the image of the null-homotopy F above in fact lies in S.

But now for large enough $r \geq s$, $I_{r,s}(S) \subseteq \mathcal{L}'$: S is compact, and its boundary, which is contained in $\partial N(\gamma)$, is stable. So the ordinary induction for Haken normal form of compact surfaces can be applied to S to show that eventually the isotopies for $\mathcal{L}$ must stabilize on S. It then follows that $I_{r,s} \circ F$ maps D^2 into X, i.e., into L, giving a null-homotopy for γ. ∎

In general, however, X is not a lamination - it fails to be a closed set. However, it does meet the 1-skeleton of each 3-simplex in a closed set. So it is a relatively simple task to categorize the kinds of bad limiting behavior that can be taking place in X; they occur only at interfaces between normal disk types, and only in one of 3 ways; see Figure 3.

Figure 3

One can rule out bad behavior between two parallel normal disks; in this case the portion of $\partial \Delta^3_i$ between the two disks is an annulus, made up of 3 or 4 rectangles, and an analysis like that in [1,Lemma 5.3] shows that once two parallel stable disks are within ϵ of one another along $\tau^{(1)}$, none of the disks in between can be moved by a boundary-surgery. The first such surgering disk, together with a tail over to the point along $\tau^{(1)}$ where the two leaves were closest, would give a disk violating [7,Theorem 1(d)], as before. Therefore any leaves between the two disks which move disappear completely under a surgery, since they are in 2-spheres which are thrown out. The disks which remain in X between the two parallel disks are therefore in a nested intersection of the disks which remain after each surgery, and hence form a closed set.

So bad limiting behavior can occur only when it involves disks of different normal types. Keeping in mind that $X \cap \tau^{(1)}$ is already a closed set, inspection leads to the catalogue of limiting behavior in Figure 3.

It then follows that $\overline{X}$ = the closure of X is the union of X and a finite number $\Delta^2_1, \Delta^2_2, \ldots, \Delta^2_k$ of normal disks. Let $\Gamma = (\bigcup \Delta^2_i) \cap X \subseteq \tau^{(2)}$; this is the set of points where $\overline{X}$ fails to be a lamination (see Figure 3). Γ is a finite graph, possibly containing some isolated vertices. What we wish to do now is to alter $\overline{X}$ to make it a lamination. The graph Γ acts very much like the branch locus of a branched surface, and we wish to do a splitting of $\overline{X}$ along $\Gamma$to obtain a lamination.

In general, a branched surface cannot be split open to become a lamination. But because the branch curves already have leaves of X limiting vertically down upon them, this allows us to see that a crucial parity condition is satisfied, and will allow us to build our lamination.

Figure 4

First we perform a preliminary 'double split' along the leaves L of X which intersect Γ: replace each leaf L with $L \cup \partial N(L)$ (see Figure 4). Replacing L with $\partial N(L)$ is what is ordinarily known as a splitting along L, so we are splitting X along L, but also keeping L (for reasons which will become apparent later). The graph Γ is then contained in $\bigcup N(L)$; in particular, each component of Γ is isolated in $\overline{X}$, in the transverse direction, on one side.

Figure 5

The next step is to alter Γ to a collection of disjoint circles. Call an edge of Γ a 'double cusp edge' if in both of the 3-simplices containing it, the edge is contained in one of our finite number of normal disks in $\overline{X} \backslash X$. Such edges can be split open; see Figure 5. In particular, since every hanging edge of Γ (one which contains a vertex with valency one) is of this type, such edges can be removed. After these splittings Γ may contain some isolated vertices; we then split along them, as well. After these preliminaries, Γ will be a graph with every vertex having valency at least two.

Every vertex v of Γ lies in a 2-disk neighborhood D in a leaf L of X, and in this disk we can assign normal orientations to the edges of Γ at v, pointing out of the 3-simplex that the 2-simplex that creates the edge is in (Figure 6). By the preliminary splitting, this is well-defined. Because each of the disks we have added to X to get $\overline{X}$ is limited upon by parallel normal disks- that is the only reason they were added - on the side away from L, we can conclude that the normal orientations of the edges must alternate as we travel around v in the disk D (see Figure 6). For in travelling around v in L, when we cross an edge of Γ, we pass (as we follow the orientation) from a part of L which is *not* limited upon by leaves of X to a part which *is*. Therefore, the next edge of Γ that we encounter must have the opposite orientation, in order for us to pass from points having leaves limiting on them to points that do not, as we cross the edge. In particular, there must be an even number of edges. So we have a situation as shown in Figure 7, and then by either flattening $\overline{\mathrm{X}}$ out at v, or splitting at v (it doesn't matter which) we can remove v from Γ , in so doing splicing adjacent edges together. Induction on the number of vertices in Γ of valency $\geq$3 finishes the argument.

Figure 6 Figure 7

The object that we now have is what is known as a *pre-lamination*; it is a branched 2-manifold all of whose branch curves are circles. Such an object can be turned into a lamination by splitting open the singular leaves. In fact, our pre-lamination has a very simple form - in consists of a collection of surfaces, with a finite number of compact surfaces (made up of our added 2-simplices) glued to them along disjoint circles. In particular, these added surfaces are never glued to one another. This pre-lamination can therefore be turned into a lamination simply by deleting those portions of our leaves ∂N(L) which are isolated on both sides. More precisely, consider the union of the ∂N(L) cut open along the loops Γ, and let Y be the collection of components of $\bigcup(\partial \mathrm{N(L)})\backslash\Gamma$ which are isolated on both sides in $\overline{\mathrm{X}}$; see Figure 8. $\mathcal{L}'$=$\overline{\mathrm{X}}\backslash$Y is then a lamination; it is a closed set because we have only removed pieces of $\overline{\mathrm{X}}$ which nothing limits on transversely, and it has the local product structure because the only places where $\overline{\mathrm{X}}$ failed to be a lamination were along Γ, but by construction the piece of ∂N(L) to one side of the circle was isolated on both sides, so has been removed.

Figure 8

$\mathcal{L}'$ is therefore a lamination, which by construction is in Haken normal form w.r.t. the triangulation τ, since it is a union of normal disks. $\mathcal{L}'$ need not be an essential lamination; however, we will see in the next sections that $\mathcal{L}'$ must in fact contain an essential *sublamination.*

5. Finding $\mathcal{L}_0$.

For convenience we will distinguish between leaves of $\mathcal{L}'$ which contain components of Γ from those which do not, calling the former split-and-paste leaves and the latter ordinary leaves. Notice that by the lemma above, all ordinary leaves are π_1-injective in M, since each is isotopic to a leaf of X.

Lemma: *Every split-and-paste leaf* L *of* $\mathcal{L}'$ *is limited on by ordinary leaves.*

Proof: If not, then for any x∈L there is an arc A meeting $\mathcal{L}'$ transversely and containing x in its interior, which meets only split-and-paste leaves of $\mathcal{L}'$.

But since every split-and-paste leaf is a limited upon by other leaves, by construction, these must all be split-and-paste leaves, as well. So the set $\mathcal{L}'\cap$A is a perfect set. Because there must be a sequence of points limiting on x, $\mathcal{L}'\cap$A is also infinite.

An infinite, perfect set in an interval is uncountable; but each leaf of $\mathcal{L}'$ can meet A in only a countable number of points, since in the leaf these points represent a discrete set in an open surface. It follows that there must be an uncountable (in particular, infinite) number of split-and-paste leaves. But since there are only finitely many split-and-paste leaves (since Γ had only finitely many components), we arrive at a contradiction. ∎

Consider the (possibly empty) collection $\mathcal{T}$ of leaves of $\mathcal{L}'$ which are compressible tori. Since all ordinary leaves are π_1-injective, hence incompressible, the leaves of $\mathcal{T}$ are all split-and-paste leaves. There are, consequently, only finitely many of them; $\mathcal{T}=\mathrm{T}_1 \cup \ldots \cup \mathrm{T}_\mathrm{k}$. Because M is irreducible, such a torus T_i either bounds a solid torus or is contained in a 3-ball; in either case, T_i separates M. Because T_i is a split-and-paste leaf, it is isolated on one side.

Notice also that by the split-and-paste construction, for every such torus, there is an ordinary leaf on the isolated side (the 'L' of 'L∪N(L)') which can be joined to the torus by an arc meeting no other leaves of $\mathcal{L}'$.

Now split M open along T_1 to get $M_1=M|T_1$. One of the two components of M_1, call it N_1, has boundary $=T_1$ which is isolated in $\mathcal{L}'\cap N_1$, i.e., $\mathcal{L}'\cap\mathrm{int}(N_1)$ is a lamination in N_1. Furthermore, $\mathcal{T}_1=\mathcal{T}\cap\mathrm{int}(N_1)$ has fewer components than $\mathcal{T}$. If we now continue, choosing a component of $\mathcal{T}_1$ and splitting, by induction we will eventually find a component N of some $M|\bigcup(T_{i_j})$ which has boundary $\bigcup(T_{i_j})$ isolated in $\mathcal{L}'\cap N$, and with $\mathcal{T}\cap\mathrm{int}(N)=\emptyset$. Consequently, $\mathcal{L}_0=\mathcal{L}'\cap\mathrm{int}(N)\subseteq N\subseteq M$ is a lamination in N; since N is a closed subset of M, $\mathcal{L}_0$ is also a lamination in M. By construction, $\mathcal{L}_0$ contains no compressible tori. It is also non-empty, because it contains the ordinary leaves described above on the isolated sides of each of the T_{i_j}.

Proposition: (a) $\mathcal{L}_0$ *contains no spheres,*
(b) $\partial(M|\mathcal{L}_0)$ *is incompressible in* $M|\mathcal{L}_0$*,*
(c) *Every leaf of* $\mathcal{L}_0$ *is end-incompressible, and*
(d) $M\setminus\mathcal{L}_0$ *is irreducible.*

Proof (a): If $\mathcal{L}_0$ contains a sphere leaf L', then it contains an sphere leaf L which is an ordinary leaf. For if L' were a split-and-paste leaf, then there are ordinary leaves limiting on it; but since L' is simply-connected, it lifts to the nearby leaves in its normal fence, so nearby leaves must cover L' , and hence be spheres themselves. But any compact ordinary leaf L of $\mathcal{L}_0$ is made up of only finitely many normal disks, so $L\subseteq I_r(\mathcal{L})$ for some r, i.e., $\mathcal{L}$ contains a leaf isotopic to L. But for $L=S^2$, this contradicts the essentiality of $\mathcal{L}$; $\mathcal{L}$ contains no spheres.

(b): Let γ be a simple loop in $\partial(M|\mathcal{L}_0)$, bounding a disk D in $M|\mathcal{L}_0$. If γ is contained in an ordinary leaf of $\mathcal{L}_0$, then this result follows immediately from the π_1-injectivity of that leaf. So we may suppose that γ is contained in a split-and-paste leaf L of $\mathcal{L}_0$.

Claim: L has trivial holonomy around γ; the normal fence A of γ meets all nearby leaves in closed loops.

Proof of claim: After an isotopy in L, we may assume that γ meets the 1-skeleton $\tau^{(1)}$ of τ. Now if there is non-trivial holonomy around γ, then there is a loop of $A\cap\mathcal{L}_0$, possibly equal to γ, with a non-compact arc α of $A\cap\mathcal{L}_0$ spiralling down towards it. This loop γ is also contained in a split-and-paste leaf L', because a null-homotopic loop in an ordinary leaf bounds a disk in that leaf, so has no holonomy around it. Since there are ordinary leaves limiting on L', there must be points in $A \cap \mathcal{L}_0$, contained in ordinary leaves, arbitrarily close to a point of α; such points must also be in half-infinite arcs limiting on γ. So we may assume that α above is in an ordinary leaf.

Figure 9

But such an arc, in spiralling down on γ, must eventually pass within ϵ of itself along a 1-simplex σ^1 of the 1-skeleton $\tau^{(1)}$, where ϵ is a monogon number for $\mathcal{L}$. The arc β in γ between two such points meets only finitely many normal disks of $\mathcal{L}_0$, and so is contained in some $I_r(\mathcal{L})$. Make the disk D^+, consisting of the compressing disk D and the positive half of the normal fence A, transverse to $I_r(\mathcal{L})$ by a small isotopy, which we may assume is fixed on the compact end of the arc α containing β. Then $D^+ \cap I_r(\mathcal{L})$ must consist of circles and arcs, because $I_r(\mathcal{L})$ is essential, and β is contained in one such arc, δ. If we look at $\delta \cap \eta$, where η is the arc of σ^1 between the endpoints of β, and choose two points of the intersection which are consecutive on η, the subarcs of β and η which they split off form a loop, bounding a disk in D^+; see Figure 9. Once again, however, this disk violates [7,Theorem 1(d)], a contradiction. This establishes the claim.

Since L must be limited on by ordinary leaves, it follows that near γ there must be loops in the normal fence, contained in ordinary leaves, which bound disks in the leaves of $\mathcal{L}_0$ containing them. Reeb stability implies that the set of loops of $\mathcal{L}_0 \cap A$ in the normal fence A which bound disks in their leaves is open in $\mathcal{L}_0 \cap A$, so the set Λ of loops which do not is closed. Such loops are contained in the split-and-paste leaves, so they meet a transverse arc in a closed, countable set. Therefore there is one such loop γ_0 which is isolated in Λ. But then γ_0 together with the loops, in ordinary leaves, in a short normal fence around it represent an embedded vanishing cycle for $\mathcal{L}_0$. Therefore, by arguments like those in [7,Lemma 2.8], the leaf of $\mathcal{L}_0$ containing γ_0 is a torus T bounding a solid torus. But this contradicts the construction of $\mathcal{L}_0$; it contains no compressible tori.

(c): Let ϵ be a monogon number for a branched surface B carrying $\mathcal{L}$. Since $\mathcal{L}'$ can have only finitely many compact split-and-paste leaves, we can find a branched surface B' carrying $\mathcal{L}'$ such that $\partial_h N(B') \subseteq L'$ and $N(B')|L'$ has no compact components.

If D_0 is a end-compressing disk for $\mathcal{L}'$, then since the union of fibers of $N(B')|\mathcal{L}'$ of length not less than $\epsilon/3$ is compact, and $M \setminus int(N(B'))$ is compact, eventually the tail of D_0 contains a fiber f of $N(B')|\mathcal{L}'$ of length less than $\epsilon/3$. Without loss of generality, we may assume that $f \subseteq \tau^{(1)}$ (by dragging D_0 there; every component of $N(B')|\mathcal{L}'$ meets $\tau^{(1)}$). If the leaf L of $\mathcal{L}'$ containing ∂D_0 is a split-and-paste leaf, then there are ordinary leaves limiting on it, and so there is an arc α in an ordinary leaf L' in the normal fence over the arc of ∂D_0 cut off by ∂f, within $\epsilon/3$ of ∂D_0. In particular, it is within $\epsilon/3$ of ∂f. Then $\alpha \cup$(the extension of f) bounds a disk D in M with $\alpha \subseteq L'$, and $f \subseteq \tau^{(1)}$ of length $< \epsilon$, and since α meets only finitely-many normal disks of L', we may assume that $\alpha \subseteq I_r(\mathcal{L})$ for some r. But this situation again violates [7,Theorem 1(d)], so $\mathcal{L}'$ is end-incompressible.

(d): Suppose $M \setminus \mathcal{L}_0$ is not irreducible, so there is a 2-sphere $S^2 \subseteq M \setminus \mathcal{L}_0$ not bounding a ball in $M \setminus \mathcal{L}_0$. Because M is irreducible, S^2 does bound a ball B^3 in M (on one side only; otherwise, $M = S^3$, a contradiction). It is easy to see that there must be innermost such 2-spheres, w.r.t inclusion of bounding 3-balls; all but finitely-many of the components of $M \setminus \mathcal{L}_0$ are contained in N(B), and so inherit an I-bundle structure from the fibering of N(B). They are therefore irreducible, and so contain no reducing spheres, since a reducing 2-sphere could be made transverse to the I-fibering. But then falling down the fibers we would find an S^2- or RP^2-leaf in the boundary of $M \setminus \mathcal{L}_0$ (i.e., in $\mathcal{L}_0$), a contradiction; nearby ordinary leaves would also have to be S^2's or RP^2's, which (being compact) would be leaves of the original $\mathcal{L}$. The remaining components of $M \setminus \mathcal{L}_0$ can each have only finitely-many 'reducing' 2-spheres which are nested in M, since for the 2-spheres to be non-nested in their component there must be leaves of $\mathcal{L}_0$ in between, hence one of the finitely-many boundary leaves of $\mathcal{L}_0$ (corresponding to $\partial_h N(B)$) in between, as well.

We therefore may assume, by passing to an innermost such S^2, that every reducing sphere in $B^3 \setminus \mathcal{L}_0$ is parallel to ∂B^3, i.e. (setting $\mathcal{L}_1 = \mathcal{L}_0 \cap B^3$ and capping B^3 off with a ball to get S^3), $S^3 \setminus \mathcal{L}_1$ is irreducible. $\mathcal{L}_1 \subseteq S^3$ also satisfies (1)-(3); any compressing or end-compressing disk could be pushed off the capping 3-ball (it's just a big point), so could be thought of as living in (B^3, hence) M.

But $\mathcal{L}_1$ cannot be essential, since 3-spheres don't contain essential laminations [7], so $\mathcal{L}_1$ must contain a compressible torus. But, back in M, this torus leaf of $\mathcal{L}_1$ is contained in a ball, so is compressible in M. But this contradicts the construction of $\mathcal{L}_0$. ∎

$\mathcal{L}_0$ is also non-empty and contains no tori bounding solid tori, by construction. It is therefore (by definition [7]) an essential lamination. Because it is a sublamination of $\mathcal{L}'$, it is also in Haken normal form w.r.t. τ, and gives us the lamination required for the theorem. The corollary follows immediately from the construction of [5, Section 3].

6. Concluding remarks.

Exactly what 3-manifolds contain essential laminations is one of the more important questions left unanswered by the theory which has been developed in recent years. It has been shown in [1] and [4] that some of the irreducible Seifert-fibered spaces with infinite fundamental group do not contain essential laminations, but to date no non-Seifert-fibered example is known. It has in fact been suggested by Gabai that (assuming the Geometrization Conjecture) no such examples exist: every hyperbolic 3-manifold will contain an essential lamination.

The result presented here gives a new tool for searching for essential laminations (and for manifolds which do not contain them), by allowing us to restrict our attention to laminations which have more structure than what a garden variety lamination could be expected to have. The key result in [1], for example, is a structure theorem for essential laminations in Seifert-fibered spaces; the additional restriction this places on the 'shape' of the lamination led, using work of Eisenbud, Hirsch, and Neumann on foliations of Seifert-fibered spaces, to the non-existence result mentioned above. Hopefully the result presented here will be of similar help in more general manifolds .

References.

[1] Brittenham, M.: *Essential Laminations in Seifert-fibered Spaces*, Topology **32** no. 1 (1993) 61-85

[2] Brittenham, M.: *Essential laminations and Haken normal form, II: regular cell decompositions*, preprint

[3] Brittenham, M.: *Essential laminations and Haken normal form, III: laminations with no holonomy*, preprint

[4] Claus, W.: *Essential laminations in closed Seifert-fibered spaces*, Thesis, University of Texas at Austin, 1991

[5] Floyd, W., Oertel, U.: *Incompressible surfaces via branched surfaces*, Topology **23** no. 1 (1984) 117-125

[6] Gabai, D.: *Foliations Transverse to Laminations in 3-manifolds*, in preparation

[7] Gabai, D., Oertel, U.: *Essential Laminations in 3-manifolds*, Annals of Math. **130** (1989), 41-73

[8] Haken, W.: *Theorie der Normalflaschen*, Acta. Math. **105** (1961), 245-375

[9] Hatcher, A.: *Measured Laminations in 3-manifolds*, preprint

[10] Jaco,W., Oertel,U.: *An Algorithm to Determine if a 3-manifold is a Haken Manifold*, Topology **23** (1984), 195-209

[11] Schubert, H.: *Bestimmung der Primfaktorzerlegung von Verkettungen*, Math. Zeit. **76** (1961), 116-148