Essential laminations and deformations of homotopy equivalences: from essential pullback to homeomorphism

Mark Brittenham

In this paper we develop techniques for using essential laminations to deform homotopy equivalences of 3-manifolds to homeomorphisms.

0. Introduction

One of the most important principles which guides our understanding of 3-manifolds is the belief that one can derive geometric-topological consequences about a 3-manifold from homotopy-theoretic data. The Poincaré conjecture is just one example of this. Another (more pertinent) example is the belief that if two (closed, say) 3-manifolds M and N are homotopy equivalent, then they are homeomorphic. This of course is false - lens spaces give counterexamples. Still, it is widely believed that if we restrict our attention to manifolds which have infinite fundamental group (to avoid lens spaces) and are irreducible (to avoid their connected sums), then this should be true: such 3-manifolds should be determined up to homeomorphism by their fundamental group alone (since under the above conditions they are both K(π,1)-spaces, hence determined up to homotopy equivalence by their fundamental groups).

This conjecture has been verified in many instances, usually by imposing some extra conditions on one of the 3-manifolds involved. Of these, perhaps the most important has been that one of the manifolds contains a 2-sided incompressible surface (i.e., the manifold is Haken). With this assumption the conjecture was proved by Waldhausen [Wa], who in fact proved something somewhat stronger:

Theorem [Wa]: If f:M→N is a homotopy equivalence between closed 3-manifolds, with M irreducible and N Haken, then f is homotopic to a homeomorphism.

Other instances in which the conjecture has been proved include when one of the manifolds is a Seifert-fibered space [Sc], when one contains a π_1-injective immersed surface of the right sort [H-S], and when both are hyperbolic [Mo].

In this paper we begin a program to generalize Waldhausen's result, by generalizing his assumption (in a different direction than [H-S]): we will assume instead that N contains an *essential lamination.* Essential laminations are a recently-defined [G-O] generalization of both the incompressible surface and the codimension-1 foliation without Reeb components. In essence, it is a codimension-1 foliation of a closed subset of N, which has π_1-injective leaves and irreducible complement. The presence of such objects has very nice consequences for their ambient 3-manifold ([G-O],[G-K]), similar to those known for incompressible surfaces; at the same time, essential laminations seem to exist in far greater abundance (see, e.g., [De], [Na], [Ro]) than incompressible surfaces do. In fact, it has been suggested that, except for a class of Seifert-fibered spaces which are known ([Br1], [Cl]) not to contain essential laminations, *every* irreducible 3-manifold with infinite fundamental group contains an essential lamination.

In proving this generalization, our intent is to follow the outline of Waldhausen's proof as closely as we can. Waldhausen's proof consists of choosing a (connected, 2-sided) incompressible surface S⊆N, making f transverse to S, and *surgering* its inverse image $f^{-1}(S)$ (by homotoping the map f) until it is incompressible. With further work (using, essentially, the h-cobordism theorem for 2-manifolds) one can then make $f^{-1}(S)$ connected, and the restriction of f, $f{:}f^{-1}(S)\rightarrow S$, a homeomorphism). One then splits M and N open along these surfaces, and continues with the induced map f. The argument is finished by an induction on a hierarchy of N - N can be successively split along incompressible surfaces into a collection of balls. The base case is Alexander's Theorem.

What we show in this paper is that we can achieve the 'last 2/3rds' of this outline:

Theorem: If f:M→N is a homotopy equivalence of (closed, orientable) 3-manifolds, M irreducible, $\mathcal{L}\subseteq$N essential and transversely orientable, and if f is transverse to $\mathcal{L}$ and $f^{-1}(\mathcal{L})$ is essential in M, then f is homotopic to a homeomorphism.

This reduces the task of proving this (conjectured) generalization to showing that one can always deform a homotopy equivalence to give an essential pullback. To date it is not known if this is always possible - this is an area of very active research at this time, and may be considered one of the most important unsolved questions in the theory of essential laminations. In a sequel to this paper [Br3], we report on some progress on the question of essential pullbacks.

It is worth noting that our approach to this problem will in the end appeal to Waldhausen's Theorem, so we will not be giving an alternate proof of that theorem. Because we do not seek to reprove the old result, we can therefore focus on those cases in which the theorem will say something new. In other words, we can assume at the outset that neither of our 3-manifolds M or N contain an incompressible surface (since Waldhausen's Theorem says that one is Haken iff the other one is). Remarkably, this assumption is exactly what is needed to find a way to use Waldhausen's Theorem. Essential laminations in non-Haken 3-manifolds inherit extra structure from the fact that there are no incompressible surfaces around, and this added structure is exactly what makes it possible to utilize Waldhausen's Theorem in an underhanded way.

1. Preliminaries

The reader is referred to [G-O] for definitions and basic results concerning essential laminations. As with a foliation, a lamination can be covered by a collection of coordinate charts for the 3-manifold, in which the leaves of the lamination appear as horizontal plaques. A lamination has a well-defined tangent space T$\mathcal{L}$, and we assume that these tangent planes vary continuously as one moves from point to point in N. We assume that M and N are both orientable and non-Haken, and that the essential lamination $\mathcal{L}\subseteq$N is *transversely oriented*, i.e., it is carried by a transversely orientable branched surface B (this is the natural generalization of the 2-sidedness of incompressible surfaces). Consequently, every leaf L of $\mathcal{L}$ is orientable (and non-compact, because N is non-Haken).

We say that a map is *transverse* to a lamination if for every $x \in f^{-1}(\mathcal{L})$,

$$f_*(T_x M) + T_{f(x)}\mathcal{L} = T_{f(x)}N.$$

One of the simplest way of achieving this is by starting with a branched surface which carries $\mathcal{L}$, and making f transverse to B; then by embedding $\mathcal{L}$ in a fibered neighborhood N(B) of B with very short fibers, f will then be transverse to $\mathcal{L}$ as well. Some of our constructions will force us to work with the more general point of view, however. (This fact will cause us some trouble, by requiring us to prove some rather technical facts which, from the branched surface point of view, are essentially trivial.)

We will make it a common practice in this paper to label with a prime '$'$' any object which is the inverse image under f of something in N which f is transverse to (or otherwise corresponds to it); thus $f^{-1}(\mathcal{L}) = \mathcal{L}'$. This rule is broken so infrequently that it should cause no problem.

Every essential lamination $\mathcal{L}$ contains a minimal sublamination $\mathcal{L}_0$; this follows easily from Zorn's Lemma applied to the collection of non-empty sublaminations. Since $f^{-1}(\mathcal{L}_0)$ is essential if $f^{-1}(\mathcal{L})$ is (a sublamination of an essential lamination is essential), we may as well assume that $\mathcal{L}$ is its own minimal sublamination. Consequently, every leaf of $\mathcal{L}$ is dense in $\mathcal{L}$.

However, we cannot yet know that the same is true of $f^{-1}(\mathcal{L}) = \mathcal{L}'$ (it is, but will take some work). But because M is non-Haken and $\mathcal{L}'$ is (assumed) essential, we already know that this is 'almost' true:

Proposition 1: An essential lamination $\mathcal{L}'$ in a non-Haken 3-manifold M contains a unique minimal sublamination $\mathcal{L}_0$; in other words, for every leaf L' of $\mathcal{L}'$, $\mathcal{L}_0 \subseteq \overline{L'}$.

A proof of this may be found in [Br2]; a different and much shorter proof of this statement (giving less information about the structure of the complement of $\mathcal{L}'$) can also be found in [Br3].

Finally, we will be making heavy use of the following result in the next section - it gives us a convenient way to 'map out inverse images'.

Theorem [EMT]: The collection $\mathcal{H}\subseteq\mathcal{L}$ of leaves having no holonomy is dense in $\mathcal{L}$.

Recall that a leaf $L\subseteq\mathcal{L}$ has no holonomy if for every (embedded) loop γ in L, the leaves of $\mathcal{L}$ intersect a small (normal) annulus over γ only in closed loops.

Of course, in the present situation, that only means that there is at least one such leaf, but that one will be enough. Note that $f^{-1}(\mathcal{H})=\mathcal{H}'$ consists of leaves with no holonomy in $f^{-1}(\mathcal{L})=\mathcal{L}'$, because transverse pictures are preserved under f (a loop in $\mathcal{H}'$ with holonomy around it would be carried under f to a loop in $\mathcal{H}$ with holonomy around it), and that $\mathcal{H}'$ is dense in $\mathcal{L}'$, because f maps short arcs α transverse to L' homeomorphically to short arcs $f(\alpha)$ transverse to $\mathcal{L}$; then since $f(\alpha)$ intersects $\mathcal{H}$ in a set dense in $\mathcal{L}$, α intersects $\mathcal{H}'$ in a set dense in $\mathcal{L}'$.

2. The inverse image of a leaf is connected

Because f is a homotopy equivalence of orientable, closed (because they are non-Haken) manifolds, f therefore has degree 1 or -1, and we can, by choosing orientations appropriately, assume that f has degree 1. We are going to show that the inverse image $f^{-1}(L)$ of a leaf L is a single leaf of $\mathcal{L}'$, by counting the degrees of the map f, thought of as a (smooth) map from a leaf of $\mathcal{L}'$ to a leaf of $\mathcal{L}$, and showing that these degrees are always 1.

It is not hard to see that the induced maps between the leaves of $\mathcal{L}'$ and $\mathcal{L}$, $f{:}f^{-1}(L)\to L$, where each has the leaf space topology (making each a smooth surface), are proper; the inverse images of compact sets are compact. This is because if $C\subseteq L$ is compact (hence compact in N, since the leaf space topology is finer than the subspace topology), then

$f^{-1}(C) \subseteq f^{-1}(L)$ is compact in M. If it is not compact in $f^{-1}(L)$, then there is a sequence $\{x_n\}_{n=1}^{\infty}$ which has no convergent subsequence in (the leaf space topology on) $f^{-1}(L)$. But since $f^{-1}(C)$ is compact in M, there is a subsequence converging to a point x of $f^{-1}(C)$ (in the topology on M). Looking in a coordinate chart about x, this convergence must be occurring in the transverse direction; the subsequence must be living in infinitely-many distinct plaques of $\mathcal{L}'$ (see Figure 1). But pushing this forward under f, we can then conclude (since f is transverse to $\mathcal{L}$) that $f(f^{-1}(C))=C$ intersects infinitely-many plaques of $\mathcal{L}$ in a coordinate neighborhood of f(x), a contradiction, since a choice of a point in each would give an infinite discrete set in C (in the leaf topology on L).

Figure 1

Because $\mathcal{L}$ (and therefore $\mathcal{L}'$) are transversely-oriented, we can, using the orientations of M and N, assign orientations to the leaves of $\mathcal{L}$ and $\mathcal{L}'$ (as, say, the second and third vectors of an orienting frame which starts with the normal orientation). These maps can therefore be assigned degrees, by the usual method of counting inverse images of regular values, giving each a sign according to the orientations described above.

Because $\mathcal{L}$ and $\mathcal{L}'$ are both essential, their leaves π_1-inject into N and M, respectively, and because f itself is an injection on π_1, it follows that for any leaf L of $\mathcal{L}$, and any leaf L$'$ of $\mathcal{L}'$ in $f^{-1}(L)$, the induced map $f:L' \rightarrow L$ is an injection on the level of π_1. Therefore, the next result becomes relevant:

Proposition 2: Let $f:S \rightarrow T$ be a π_1-injective, proper map between connected, orientable, non-compact open surfaces, $T \neq \mathbf{R^2}$. Then either

(a) (degree(f)=0) $S=\mathbf{R^2}$ or $\mathbf{S^1} \times \mathbf{R}$, and f can be properly homotoped outside of any compact subset of T, or

(b) (degree(f)$\neq$0) f is properly homotopic to a finite-sheeted covering map.

The proof [Br4] is somewhat tedious, though not hard, and is in fact much in the spirit of a published proof of the analogous 3-dimensional result; see [B-T].

Lemma 3: Let $\mathcal{L}_0$ be the minimal sublamination of $\mathcal{L}'$. Then f($\mathcal{L}_0$)=$\mathcal{L}$.

Proof: $\mathcal{L}_0$ contains a non-simply-connected leaf L$'$, for otherwise $\mathcal{L}_0$ is an essential lamination consisting of planes, and a result of Gabai [Ga] implies that M is the 3-torus, which is Haken. Let L be the leaf that L$'$ maps to; then f:L$'\to$L satifies the conditions of the proposition above. So either (in case (b)) f(L$'$)=L, so f($\mathcal{L}_0$)=f($\overline{\mathrm{L}'}$)$\supseteq$ $\overline{\mathrm{f}(\mathrm{L}')}$ =$\overline{\mathrm{L}}$=$\mathcal{L}$ (where the middle containment is because f is a closed map), or (in case (a)) L$'$=$\mathbf{S^1} \times \mathbf{R}$ and maps onto some end ϵ of L. But then f($\mathcal{L}_0$)=f($\overline{\mathrm{L}'}$)$\supseteq$ $\lim_\epsilon$(L)=$\mathcal{L}$ (because $\lim_\epsilon$(L) is a non-empty sublamination of $\mathcal{L}$). ∎

Using an argument similar to the one above, we can show that for every leaf L, f^{-1}(L) consists of only countably-many leaves. For if there were uncountably-many leaves, then for any finite cover of M by coordinate charts for $\mathcal{L}'$, some chart intersects uncountably-many of these leaves, and so contains uncountably-many plaques of f^{-1}(L). But it is a standard fact (by repeatedly cutting an interval in half) that for any uncountable set A in an interval I, there is a point x of the interval so that every neighborhood of x intersects A in an uncountable set. Therefore, there is a point y in a plaque of $\mathcal{L}'$ (corresponding to x) so that every neighborhood of x intersects uncountably-many plaques of f^{-1}(L). In particular, any (short) transverse arc through y intersects an uncountable number of such plaques. But the image under f of this arc then gives an arc transverse to $\mathcal{L}$ intersecting uncountably-many plaques of L, in some coordinate neighborhood about f(y), an impossibility, since L can intersect a chart in at most countably-many plaques (surfaces do not contain uncountable discrete sets, being second countable).

Therefore, we can apply Sard's Theorem to conclude that our map f:f^{-1}(L)$\to$L has regular values. In particular, since f is transverse to $\mathcal{L}$, any regular value for this map is also a regular value for f:M$\to$N. Therefore, there are regular values for f which live in leaves of $\mathcal{L}$. Given any such regular value x, we can find a small neighborhood $\mathcal{V}$ (in N) of points which are also regular values for f, and $\mathrm{f}^{-1}(\mathcal{V})=\mathcal{V}_1, \cup \cdots \cup \mathcal{V}_\mathrm{n}$, each mapping homeomorphically down onto $\mathcal{V}$ under f. In particular we can find a small arc I$\subseteq\mathcal{V}$ (which we can think of as a short vertical interval in some coordinate chart) transverse to $\mathcal{L}$(and therefore intersecting every leaf of $\mathcal{L}$, since every leaf is dense) consisting of regular values for f. This arc then has n disjoint inverse images $I_1 \cup \ldots \cup I_n$, and every $y \in I \cap \mathcal{L}$, $y \in L_1 \subseteq \mathcal{L}$, is a regular value for $f{:}f^{-1}(L_1) \to L_1$) with $f^{-1}(y)$ consisting of n points. We can therefore use the points of $I \cap \mathcal{L}$ to calculate the degrees of the maps induced from f on the leaves of $\mathcal{L}'$.

Lemma 4: $\mathcal{L}_0 = \mathcal{L}'$.

Proof: If not, then there is a leaf L' of $\mathcal{L}'$, not contained in $\mathcal{L}_0$, with $\mathcal{L}_0 \subseteq \overline{L'}$ and $\mathcal{L}_0 \neq \overline{L'}$. Consider a point $x \in I \cap \mathcal{L}$ and the points $\{x_1, \ldots, x_n\} = f^{-1}(x) \subseteq \mathcal{L}'$. Some collection $\{x_1, \ldots, x_k\}$ lies in $\mathcal{L}' \setminus \mathcal{L}_0$, and the rest $\{x_{k+1}, \ldots, x_n\}$ lie in $\mathcal{L}_0$. Because $\mathcal{L}_0$ is a closed set, for each $x_i \in \mathcal{L}' \setminus \mathcal{L}_0$ there is an open neighborhood $\mathcal{O}_i$ of x_i in $I_i \subseteq f^{-1}(I)$ which misses $\mathcal{L}_0$.

Consider $\mathcal{O} = f(\mathcal{O}_1) \cap \ldots \cap f(\mathcal{O}_k) \subseteq I$. Because f maps the arcs I_i homeomorphically to I, this is an open subset of I containing x. Consider $f^{-1}(\mathcal{O}) \subseteq f^{-1}(I)$; note that $f^{-1}(\mathcal{O}) \cap I_i \subseteq \mathcal{O}_i$ for $1 \leq i \leq k$.

Now look at $f^{-1}(\mathcal{O}) \cap I_{k+1} = \mathcal{O}'$. This is an open neighborhood of x_{k+1} in I_{k+1}. Because $\overline{L'}$ contains ($\mathcal{L}_0$ and hence) x_{k+1}, there are points of L' in I_{k+1} which pass arbitrarily close to x_{k+1} and hence are contained in $\mathcal{O}_i$. Choose one; call it x'_{k+1}. Then $x'_{k+1} \notin \mathcal{L}_0$, and $f^{-1}(f(x'_{k+1})) \cap I_i = f^{-1}(y') \cap I_i \in \mathcal{O}_i$ are not in $\mathcal{L}_0$ for $1 \leq i \leq k$. So we have increased the number of points in the inverse image of a point of $\mathcal{L} \cap I$ which are not in $\mathcal{L}_0$. Continuing, we can therefore find a point y of $\mathcal{L} \cap I$ with $|f^{-1}(y) \cap (\mathcal{L}' \setminus \mathcal{L}_0)| = n$, i.e., $f^{-1}(y) \subseteq \mathcal{L}' \setminus \mathcal{L}_0$. But this means that $y \notin f(\mathcal{L}_0)$, contradicting the previous lemma. So $\mathcal{L}_0 = \mathcal{L}'$. ∎

Lemma 5: All leaves of $f^{-1}(\mathcal{H})$ map to their corresponding leaves of $\mathcal{L}$ with the same degree.

Proof: Suppose not; let L_1 and L_2 be leaves of $f^{-1}(\mathcal{H})$ which map with different degrees. Because $\mathcal{L}' = \mathcal{L}_0$, we have $L_2 \subseteq \overline{L_1}$ and $L_1 \subseteq \overline{L_2}$. Because they have different degrees, one of them has non-zero degree, say L_1. f maps L_1onto some leaf L of $\mathcal{L}$.

Pick a point y in the arc I which lies in the leaf L, and consider all of the points $\{y_1, \ldots, y_k\}$, $k \geq 1$, of $f^{-1}(y)$ which lie in L_1. We can join them together by arcs, each point joined to y_1, say, to form a (singular) tree , =, $_y$ in L_1. Each arc descends under f to a loop in the leaf of $\mathcal{L}$ containing $f(L_1)$; because this leaf is in $\mathcal{H}$, the normal fence over this loop meets all of the nearby leaves of $\mathcal{L}$ in closed loops. In M, this means that the normal fence over each of our arcs meets all of the nearby leaves in arcs whose endpoints map to the same point under f, i.e., the lifts , $_z$ of , $_y$ along the normal fence (lifting y_1 to z_1 over z) have all of their endpoints lying in $f^{-1}(z)$.

But because L_2 passes arbitrarily close to L_1, It must intersect the normal fence over , $_y$ in some , $_z$ and so $\{z_1, \ldots, z_k\} \subseteq L_2 \cap f^{-1}(z)$. But because Σ(local degrees at z_i)=Σ(local degrees at y_i)=degree(L_1)$\neq$degree(L_2), there must be additional points of $f^{-1}(z)$ in L_2. But then we can apply the same argument in the other direction to show that there must be another point w of I for which L_1 contains even more points of $f^{-1}(w)$. But because $|f^{-1}(w)| = |f^{-1}(x)| = n$ for all w in I, continuing such arguments eventually forces us into a contradiction; we can't always find more points. So degree(L_1)=degree(L_2). ∎

Corollary 6: Every leaf of $f^{-1}(\mathcal{H})$ maps with degree one; in particular, for every leaf L of $\mathcal{H}$, $f^{-1}(L)$ is connected.

Proof: Every leaf of $f^{-1}(L)$, $L \in \mathcal{H}$, maps with the same degree, but the sum of these degrees is the degree of $f: M \to N$, which is one. Therefore, there can be only one leaf in each inverse image, and it maps with degree one. ∎

Proposition 7: For every leaf L of $\mathcal{L}$, $f^{-1}(L)$ is connected.

Proof: Let L' be a leaf of $f^{-1}(L)$; because $\overline{L'} = \mathcal{L}'$, L' limits on a leaf L_1 of $f^{-1}(\mathcal{H})$. Pick a point y of I contained in $f(L_1)$; because $f^{-1}(y) \subseteq L_1$, we can join all of the inverse images together in a tree , $_y \subseteq L_1$. But then L' meets the normal fence over , $_y$ arbitrarily close to , $_y$, so there is a point z in I near y for which , $_z \subseteq L'$, and so L' contains all of the inverse images of z. Therefore degree(L')=degree(f)=1; as before, this implies that $L' = f^{-1}(L)$. ∎

3. The homotopy equivalence of $M\backslash\gamma$

Choose a non-simply-connected leaf L of $\mathcal{L}$ (again, [Ga] says that either such a leaf exists or N is a 3-torus, hence Haken). $L'=f^{-1}(L)$ is connected, hence is mapped to L by f with degree one, so $f:L'\to L$ is properly homotopic to a 1-sheeted covering map, i.e., a homeomorphism. In particular (from the proof in [Br4]), if γ is an essential simple loop in L, there is a homotopy H of $f:L'\to L$, supported on a compact set $C\subseteq L'$, making $f^{-1}(\gamma)=\gamma'$ a connected, essential, simple loop mapping homeomorphically to γ under f.

For our next step we must first alter $\mathcal{L}$ (and $\mathcal{L}'$) by splitting $\mathcal{L}$ along L (see [G-O]), and filling in the resulting (product) I-bundle with (a Cantor's set-worth of) parallel copies of L. We also do the same thing with L' in $\mathcal{L}'$. Then by redefining f to be $f|_{L'}\times I$ on the I-bundle and the old f elsewhere, we get a new map (still called f), homotopic to the old one, transverse to a new essential lamination (still called $\mathcal{L}$), and with a new essential lamination (still called $\mathcal{L}'$) as a pullback. We label as L (and L') the leaves in the 'center' of the two I-bundles. These laminations no longer have every leaf dense (but we won't be needing that fact any more).

We can then deform f to a new map (still called f) with $f^{-1}(\gamma)=\gamma'$ by redefining f on $C\times I$ (in the I-bundle over L) to be two copies of H glued together along the face where the inverse image of γ is γ'. This map is still transverse to $\mathcal{L}$ and $\mathcal{L}'$ is its pullback.

By taking small tubular neighborhoods of γ and γ', we get induced maps $f:M\backslash\gamma'\to N\backslash\gamma$ and $f:M\backslash\mathring{N}(\gamma')\to N\backslash\mathring{N}(\gamma)$, with $f:\partial N(\gamma')\to\partial N(\gamma)$ a homeomorphism. Note that the inclusions $M\backslash\mathring{N}(\gamma')$ and $N\backslash\mathring{N}(\gamma)$ into $M\backslash\gamma'$ and $N\backslash\gamma$ are homotopy equivalences, and the map $f:M\backslash\mathring{N}(\gamma')\to N\backslash\mathring{N}(\gamma)$ has degree one, hence is a surjection on π_1. Also, $M\backslash\mathring{N}(\gamma')$ and $N\backslash\mathring{N}(\gamma)$ are irreducible, with incompressible boundary: for example, if $S\subseteq N\backslash\mathring{N}(\gamma)$ were a reducing sphere, then in N it bounds a ball; this ball would then have to (meet, hence) contain $N(\gamma)$, implying that γ was not essential (in N, hence) in its leaf. Also, if $\partial N(\gamma)$ were compressible in $N\backslash\mathring{N}(\gamma)$, then $N\backslash\mathring{N}(\gamma)$ would be a solid torus, implying that N was the union of two solid tori joined along their boundary. But being laminar, N has universal cover $\mathbf{R^3}$ [G-O], while the union of two solid tori does not.

So $M\setminus \mathring{N}(\gamma')$ and $N\setminus \mathring{N}(\gamma)$ are $K(\pi,1)$ spaces, and so the map $f{:}M\setminus \mathring{N}(\gamma')\rightarrow N\setminus \mathring{N}(\gamma)$ is a homotopy equivalence iff it is injective on π_1, or equivalently, if $f{:}M\setminus\gamma'\rightarrow N\setminus\gamma$ is injective on π_1.

Proposition 8: $f{:}M\setminus\gamma'\rightarrow N\setminus\gamma$ is injective on π_1.

Proof: Let α' be a loop in $M\setminus\gamma'$ with $f(\alpha')=\alpha$ null-homotopic in $N\setminus\gamma$, and let $H{:}D^2\rightarrow N\setminus\gamma$ be a null-homotopy. α is in particular then null-homotopic in N, so by the π_1-injectivity of $f{:}M\rightarrow N$, α' is null-homotopic in M, by some null-homotopy $H'{:}D^2\rightarrow M$. Because

$$f\circ H'|_{\partial D^2}=f\circ\alpha'=H|_{\partial D^2},$$

we can form a map $G=f\circ H'\cup H{:}S^2\rightarrow N$, a map equal to $f\circ H'$ on the upper hemisphere $\mathbf{S^2_+}$ and to H on the lower hemisphere $\mathbf{S^2_-}$ (see Figure 2). We can assume this map is transverse to $\mathcal{L}$ and γ, by making each piece so transverse, using a Morse theory argument. Cover our laminations $\mathcal{L}$ and $\mathcal{L}'$ by finitely-many coordinate charts; in each chart our lamination looks like a Cantor's set-worth of horizontal disks. Then we can put the maps H and H′ into Morse-normal form w.r.t the vertical direction of each chart, dragging the singularities up or down into the regions between leaves, if necessary.

Figure 2

Now consider $\lambda=G^{-1}(\mathcal{L})\subseteq S^2$; it is a 1-dimensional lamination in S^2 transverse to the equator $S^1\subseteq S^2$. Because $\mathcal{L}$ is essential and G is transverse to $\mathcal{L}$, λ has no monogons and no non-trivial holonomy around closed loops, and so it follows (see, e.g., [G-O, p.51]) that every leaf of λ is a closed loop. Also, $G^{-1}(\gamma)$ consists of a finite number of points, all in the upper hemisphere $\mathbf{S^2_+}$.

What we wish to do is to remove the points $G^{-1}(\gamma)= (f\circ H')^{-1}(\gamma)=(H')^{-1}(\gamma')$ from the upper hemisphere $\mathbf{S^2_+}$. If some point x is in a loop of $G^{-1}(\mathcal{L})$ which is entirely contained in the upper hemisphere (so it is in a loop of $(f\circ K)^{-1}(\mathcal{L})= K^{-1}(\mathcal{L}')$), then this loop bounds a disk Δ in the upper hemisphere, which represents a null-homotopy (under H$'$) of the loop into M. Because $\mathcal{L}'$ is essential, this loop is also null-homotopic in the leaf L$'$ of $\mathcal{L}'$ containing it. By redefining H$'$ on Δ to lie in L$'$, and pushing off L$'$ slightly, we can remove the loop from $\mathrm{G}^{-1}(\mathcal{L})$, in so doing removing points of $\mathrm{G}^{-1}(\gamma)$ from the upper hemisphere. In this way we can remove any points of $\mathrm{G}^{-1}(\gamma)$ which lie in such loops, by simply redefining the null-homotopy.

Our only problem is that some points could lie on loops which are not entirely contained in the upper hemisphere. What we need to do is to show how to 'pull' these loops into the upper hemisphere. This will involve deforming the loop α'; but we will show that this deformation need not cross γ'. So in the end we will have shown that α' is freely-homotopic, in the complement of γ', to a loop which is null-homotopic in the complement of γ'. Therefore, α' itself is null-homotopic in M$\setminus\gamma'$.

It is actually technically easier to pull *all* of $\mathrm{G}^{-1}(\mathcal{L})$ into $\mathbf{S^2_+}$; this is what we will seek to do. This will require three facts. First, the induced map $\mathrm{f}_*{:}\pi_1(\mathrm{L}\setminus\gamma) \to \pi_1(\mathrm{L})$ is injective: this follows because γ is essential in L. Second, the map $\mathrm{f}_*{:}\pi_1(\mathrm{M},\mathrm{L}'\setminus\gamma')\to\pi_1(\mathrm{N},\mathrm{L}\setminus\gamma)$ is injective. This follows from the commutative diagram

$$\begin{array}{ccccccc} \pi_1(\mathrm{L}'\setminus\gamma') & \to & \pi_1(M) & \to & \pi_1(M, L'\setminus\gamma') & \to & \pi_0(L'\setminus\gamma') \\ \downarrow & & \downarrow & & \downarrow & & \downarrow \\ \pi_1(L\setminus\gamma) & \to & \pi_1(N) & \to & \pi_1(N, L\setminus\gamma) & \to & \pi_0(L\setminus\gamma) \end{array}$$

where all of the maps are induced by f, together with (half of) the five lemma, since the first vertical arrow is surjective (it's degree one), the second is injective, and the fourth is injective for homological reasons (γ' separates iff γ does). Finally, the map $\mathrm{f}_*{:}\pi_1(\mathrm{L}',\mathrm{L}'\setminus\gamma')\to\pi_1(\mathrm{L},\mathrm{L}\setminus\gamma)$ is injective; this follows by a similar five lemma argument. (Of course, half of the objects in these diagrams are not groups! However, since by 'injective' we mean only that the only element that these maps send to the trivial element is the trivial element, and this is proved in the five lemma without appealing to any group structure, this lapse will cause no difficulty.)

In what follows we will always write as if the circles and arcs we are dealing with are contained in L and L'(so that we need to arrange things to avoid bumping into γ and γ'), even though this is not always true; this is assumed only for the convenience of the reader. If we are not working in these leaves, we run into fewer problems, and the reader can supply his or her own arguments.

By an argument like the one above we can remove any circles of $G^{-1}(\mathcal{L})=\lambda$ which are entirely contained in the lower hemisphere $\mathbf{S^2_-}$.

Now consider an *innermost arc* ω of an *outermost family* of parallel arcs of $\lambda\cap\mathbf{S^2_-}$, cutting off an outermost disk Δ from $\mathbf{S^2_-}$, and set $\alpha_0=\Delta\cap\partial\mathbf{S^2_-}$ (see Figure 3). What we wish to do is 'pull' this disk over from $\mathbf{S^2_-}$ to $\mathbf{S^2_+}$ (i.e., replace H restricted to this disk with the image under f of a map into $M\setminus\gamma'$). To be careful about this, we must first divide this family of parallel arcs into a finite number of (disjoint) families, on each of which we will then carry out this 'replacement'.

Figure 3

Choose a branched surface B carrying $\mathcal{L}$ and a fibered neighborhood N(B) with $\partial_h N(B)\subseteq\mathcal{L}$. By splitting B if necessary, we can assume that $N(B)|\mathcal{L}$, a collection of I-bundles over surfaces, has no component which is an I-bundle over a compact surface (by deleting such bundles and then collapsing the fibers). Because α and $\partial_h N(B)\cap\partial_v N(B)$ (a finite collection of loops) are both one dimensional, we can assume that they miss one another. Now look at the set $H^{-1}(N(B)\cap\alpha_0=A$. This is a closed subset of α_0, and (because H is transverse to $\mathcal{L}$) every point of $H^{-1}(\mathcal{L})\cap A$ is contained in a (maximal) interval in A which is not a point. Also, if an endpoint of one of these intervals maps to $\partial_h N(B)$ (hence to $\mathcal{L}$), then it is isolated on one side. Our first claim is that there are only finitely-many such intervals. For suppose there were infinitely-many, and suppose first that infinitely-many had endpoints mapping to $\partial_h N(B)$. This sequence of intervals must then have lengths tending to 0, and their endpoints contain a subsequence converging to a point x mapping to $\partial_h N(B)$ (since its inverse image is closed). But this point x is then isolated in A on one side, contained in a non-trivial interval of A on the other, and limited on by the endpoints of a sequence of intervals in A, which is clearly absurd. Therefore, infinitely-many of these intervals would have to have (both of their) endpoints in $\partial_v N(B)$. But since by hypothesis each also intersects $H^{-1}(\mathcal{L})$, we then get a sequence of points in $H^{-1}(\mathcal{L})$ converging to a point y mapping to $\mathcal{L}$. But since the lengths of our intervals tend to zero, their end points also converge to y, so (since the endpoints map to $\partial_v N(B)$) y also maps to $\partial_v N(B)$, so $f(y) \subseteq \mathcal{L} \cap \partial_v N(B) = \partial_h N(B) \cap \partial_v N(B)$, a contradiction. So there are only finitely-many such intervals.

N(B) can be foliated (by foliating the I-bundles between the leaves of $\mathcal{L}$) with $\mathcal{L}$ as a sublamination. By imagining that we locally crush leaves of this foliation to a point (in coordinate charts, say) we obtain local maps of I into an interval, which (since H is transverse to $\mathcal{L}$) is non-singular (i.e., locally injective) at the points of $I \cap H^{-1}(\mathcal{L}) = C$. Our next claim is that we can find a finite number of disjoint subintervals of I, whose union contains C, each of which is mapped injectively under one of our local maps. This is because since the local maps are non-singular on C, for each point z of C, there is a neighborhood of it on which the local map is injective. This gives an open cover of C, which has a finite subcover; by breaking ties (C is a Cantor set and the cover consists of intervals - pick points not in C in each overlap and shorten each interval using that point), we obtain our disjoint intervals.

This gives us in total a finite number of intervals covering one end of our family of arcs (by only taking those which intersect that end), for which the points of C map injectively into a local model of N(B) with its leaves crushed to points. These intervals partition our family of parallel arcs into a finite number disjoint, parallel subfamilies. We will now show how to pull the (outermost) subfamily into $\mathbf{S}^2_+$, leaving the rest of the subfamilies fixed; by induction we can then pull the entire family into $\mathbf{S}^2_+$.

Let ω_0 be the innermost arc of this outermost subfamily, cutting off the subdisk Δ_0 of Δ, and let $\alpha_0 = \Delta_0 \cap S^1$. Under H, Δ_0 represents a homotopy of $f \circ \alpha|_{\alpha_0} : (\alpha_0, \partial\alpha_0) \to (N, L \backslash \gamma)$ into $L \backslash \gamma$; by the injectivity above, in follows that there is a homotopy J' of

$$H'|_{\alpha_0} : (\alpha_0, \partial\alpha_0) \to (M, L' \backslash \gamma')$$

to an arc ω_0' in $L' \backslash \gamma'$. Because L' is two-sided, we can continue the homotopy and push this off of L' (in the direction that the arc α_0 was heading at its endpoints) to an arc β'. This gives us a map $J' : \Delta_0' \to M$, which is transverse to $\mathcal{L}'$ along its boundary; we can therefore, by a small deformation fixed on the boundary, make J' transverse to $\mathcal{L}'$ (and γ'); any loops in $(J')^{-1}(\mathcal{L}')$ can be removed as before (see Figure 4).

Figure 4

$f \circ \omega_0'$ and $H \circ \omega_0$ are both homotopic rel endpoints to $H \circ \alpha_0 = f \circ H' \circ \alpha_0'$ (under $f \circ J'$ and H), and both map into $L \backslash \gamma$, so together they form a loop in $L \backslash \gamma$ null-homotopic in N; by the injectivity of the composition $\pi_1(L \backslash \gamma) \hookrightarrow \pi_1(L) \hookrightarrow \pi_1(N)$, this loop is therefore null-homotopic in $L \backslash \gamma$, by a homotopy J^+. After pushing this homotopy off of L (as before), we can think of it as a homotopy in the complement of $\mathcal{L}$ between $f(\beta')$ and β. We can now replace $H|_\Delta$ with $(f \circ J') \cup J^+$ to get a picture like Figure 5; this, together with $f \circ H'$ and $H|_{\mathbf{S}^2_- \backslash \Delta}$ form a new map G transverse to $\mathcal{L}$. Notice that the lamination $(J')^{-1}(\mathcal{L}')$ consists of arcs joining the *same* pairs of points as $H^{-1}(\mathcal{L}) \cap \Delta_0$ did; for if any are knocked 'off-line' we would be able to find either a monogon side of our subfamily (the (image of) the disk the outermost such arc cuts off can be completed to a monogon, because $N(B)|\mathcal{L}$ has no compact I-bundles; see Figure 6), or non-trivial holonomy around a null-homotopic loop in $G^{-1}(\mathcal{L})$ (which is impossible for an essential lamination - it implies non-trivial holonomy around a loop null-homotopic in its leaf). If we redefine $K = K \cup J$, we then succeed in 'pulling' the outermost family of arcs into $\mathbf{S}^2_+$. J may meet γ'; however, it is still true that $G|_{\partial(\mathbf{S}^2_+ \cup \Delta)}$ is homotopic to $G|_{S^1} = \alpha$ in the complement of γ'.

Figure 5 Figure 6

To see this, consider the (finite number of) points of $(J')^{-1}(\gamma')$; they are contained in (finitely-many) arcs $\beta'_1, \ldots, \beta'_k$ of $J^{-1}(L')$. The endpoints of each of these arcs β'_i bound an arc β_i of $H^{-1}(\mathcal{L})\cap\Delta$; since both $f(\beta'_i)$ and β_i are homotopic rel endpoints to (a subarc of) α_0, they are therefore homotopic rel endpoints to one another in N. They are therefore homotopic rel endpoints to one another in L, by the π_1-invectivity of L. In other words, $f(\beta'_i)$ is homotopic rel endpoints, in L, to an arc missing γ (because β_i misses γ). By the third injectivity result above, this implies that β'_i is homotopic rel endpoints, in L', to a map missing γ'. If we imagine tacking these homotopies onto J' (in the normal direction; see Figure 7), they give us a prescription for *lifting* α'_0 in M *over* the points of $(J')^{-1}(\gamma')$ to ω'_0, i.e., α_0 is homotopic to ω'_0, rel endpoints, in the complement of γ'.

Figure 7

Continuing inductively, since there are only finitely many parallel families of arcs in $\lambda\cap\mathbf{S}^2_-$, and only finitely-many subfamilies in each parallel family, we can pull all of the arcs of $G^{-1}(\gamma)$ in the lower hemisphere into the upper hemisphere (at the expense of deforming α' in the complement of γ'); then by doing disk-replacements as above, we can remove all points of $(H')^{-1}(\gamma')$ from $\mathbf{S}^2_+$, achieving our desired null-homotopy of α'. ∎

This proposition completes the theorem. $f{:}M\setminus\mathring{N}(\gamma')\to N\setminus\mathring{N}(\gamma)$ is then a homotopy equivalence between 3-manifolds with incompressible boundary, which is a homeomorphism on the boundary; Waldhausen's Theorem [Wa] says that f is homotopic, rel boundary, to a homeomorphism. By gluing on the homeomorphism $f{:}N(\gamma')\to N(\gamma)$ by a constant homotopy, this gives a homotopy of $f{:}M\to N$ to a homeomorphism.

4. Concluding remarks

The question that we leave unanswered here - can we deform a homotopy equivalence to give an essential pullback - seems to be a rather elusive one. It is even unclear whether one should expect a finite number of 'surgeries' to suffice, as in Waldhausen's original proof, or whether something more like a 'convergence' result (like that of [M-S], for example) will be required. Evidence seems to point to the latter [Ga2], but it is possible that this could be changed by once again invoking the hypothesis that both manifolds be non-Haken [Br3].

One direction in which to improve the result presented here is to remove the hypothesis that $\mathcal{L}$ be transversely orientable. This hypothesis was used to allow us to make coherent choices of orientations for leaves, in order to carry out our degree calculations. Losing this hypothesis would only affect the proof of Lemma 5, however. For example, if we only assume that every leaf in the range is orientable (a fact which seems to be true for nearly all of the known examples), we can by the same method of proof recover the fact that all leaves in the inverse image of a leaf with no holonomy have the same degree, up to sign (and that degree must therefore be 1). For all leaves in the domain must then be orientable, since otherwise an orientation-reversing loop in a leaf L′ is a transverse orientation-reversing loop (since M is orientable), so maps (since f is transverse to $\mathcal{L}$) to a transverse orientation-reversing loop in a leaf L, hence (since N is orientable) an orientation-reversing loop.

In fact, if we assign a (temporary) orientation to our transverse arc I, in so doing assigning orientations to each of its pullbacks (since they map homeomorphically), we can then assign (temporary) orientations to the leaves of L′, by completing the transverse orientation to the orientation of M. The orientations so assigned at the inverse image of a point all agree with orientations on the leaves containing them; otherwise, as before, we could draw an arc between two points with opposite orientation, which therefore flips the transverse orientation, as well; the arc would then map to a transverse orientation-reversing loop in our leaf.

Then the local degree of f at a point agrees with the local degree of f restricted to our leaf L′; this gives us a degree of the map f:L′→L, well defined up to sign, which can be calculated by using the local degrees of f thought of as a map from M to N. But then if two leaves in $f^{-1}(\mathcal{H})$ have different degrees (up to sign), we can use the same see-saw method of proof in Lemma 5 to manufacture a contradiction.

This allows us to give a different proof of Lemma 5 (and Corollary 6) if we assume that our lamination $\mathcal{L}$ is (a sublamination of) a foliation with orientable leaves. For then by passing to a 2-fold covering in the range (and the corresponding 2-fold covering in the domain), we can make our split foliation (and hence its pullback) transversely orientable. Then since (by [Br2]) both laminations still have every leaf dense, we can then recover Lemma 5 and Corollary 6 for these new laminations.

But this in turn allows us to easily recover these results for our original lamination. For either a leaf in the range has a single inverse image under the orienting cover, or two inverse images. If it has one inverse image, then the cover of the pullback also consists of a single leaf, and so the pullback does. If the leaf in the range has two inverse images under the covering, then the covering of the pullback consists of two leaves, so the pullback consists of one or two leaves. But if there are two, they map to our original leaf with the same degree, up to sign, which implies that our map f has even degree, a contradiction. So all leaves still have inverse images consisting of a single leaf.

References

[Br1] Brittenham, *Essential laminations in Seifert-fibered spaces*, Topology **32** no. 1 (1993), 61-85

[Br2] Brittenham, *Essential laminations in non-Haken 3-manifolds*, to appear in Topology and its Applications

[Br3] Brittenham, *Essential laminations and deformations of homotopy equivalences, II: the structure of pullbacks*, in preparation

[Br4] Brittenham, *π_1-injective, proper maps of open surfaces*, preprint

[B-T] Brown and Tucker, *On proper homotopy theory for non-compact 3-manifolds*, Trans. A.M.S. **188** no. 2 (1974), 105-126

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

[De] Delman, *Essential laminations in knot complements*, Thesis, Cornell University, 1992

[EMT] Epstein, Millet, and Tischler, *Leaves without holonomy*, J. London Math. Soc. (2) **16** (1977), 548-552

[Ga] Gabai, *Foliations and 3-manifolds*, Proc. ICM Kyoto 1990, 609-619

[Ga2] Gabai, personal communication

[G-K] Gabai and Kazez, *3-manifolds with essential laminations are covered by solid tori*, preprint

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

[H-S] Hass and Scott, *Homotopy equivalence and homeomorphism of 3-manifolds*, preprint

[M-S] Morgan and Shalen, *Degenerations of hyperbolic structures,* II: *measured laminations, trees, and 3-manifolds*, preprint

[Mo] Mostow, *Quasi-conformal mappings in n-space and rigidity of hyperbolic space forms*, Publ. IHES, no. 34 (1968), 53-104

[Na] Naimi, *Constructing essential laminations in some 3-manifolds*, Thesis, Caltech, 1992

[Ro] Roberts, *Taut foliations in alternating knot exteriors*, Thesis, Cornell University, 1992

[Sc] Scott, *The are no fake Seifert-fibered spaces with infinite* π_1, Annals of Math. **117** no. 1 (1983) 35-70

[Wa] Waldhausen, *On irreducible 3-manifolds which are sufficiently large*, Annals of Math. **87** (1968), 56-88

Department of Mathematics

University of Texas at Austin

Austin, TX 78712

AMS Subject Classifications: 57M10, 57M99

Keywords: essential lamination, non-Haken, pullback

Research supported in part by NSF grants DMS-8610730 and DMS-9203435