Essential laminations and Haken normal form, II: regular cell decompositions

Mark Brittenham

This paper is a sequel to [B1]. In that paper we showed how, given a triangulation τ of a 3-manifold M, to use an essential lamination $\mathcal{L}$ in M to find a (usually different) essential lamination $\mathcal{L}_0$ which was in Haken normal form with respect to the triangulation τ. In this paper we show how to extend this result to more general cell decompositions.

A cell decomposition of a 3-manifold is called *regular* if every k-cell is a polyhedron, and every face of every k-cell is glued to a (k-1)-cell by a homeomorphism. Such a decomposition arises very naturally, for example, from a Heegard decomposition of M; the decomposition then has one 0-cell and one 3- cell. A lamination is in *normal form* w.r.t. a regular decomposition if it is transverse to the decomposition, and it meets the 3-cells B^3_k in disks, each meeting every 1-cell in the induced cell decomposition of ∂B^3_k at most once. The main result of this paper is:

Theorem: If M is a 3-manifold with a regular cell decomposition $\{B^3_k\}$, and M contains an essential lamination $\mathcal{L}_0$, then there is an essential lamination $\mathcal{L}$ in M which is in normal form w.r.t. the cell decomposition.

This result extends that of [B1], showing that normal essential laminations can be found for all reasonable decompositions of 3-manifolds. Essential laminations therefore behave in essentially the same way that incompressible surfaces do, in terms of Haken normal form.

One interesting consequence of this theorem is:

Theorem: If M=$H_1 \cup H_2$ is a Heegard decomposition of M and if M contains an essential lamination $\mathcal{L}_0$, then there is an essential lamination $\mathcal{L}$ in M such that:

(1) $\mathcal{L}$ is transverse to $\partial H_1 = \partial H_2$,

(2) $\mathcal{L} \cap H_1$ consists of compressing disks for ∂H_1, which form $\leq$genus(H_1) parallel families, and

(3) $\mathcal{L} \cap H_2$ is a lamination in H_2 with π_1-injective leaves.

The proof of the main theorem is an extension of the proof found in [B1]. We start by turning our cell decomposition into a triangulation, using a slight variation of the first barycentric subdivision. We then use [B1] to develop an (in general infinite-time) isotopy which finds an essential lamination $\mathcal{L}_1$ which is in normal form w.r.t. the triangulation. It will not in general be in normal form w.r.t. the cell decomposition, so we take steps to make it look 'more' like one in normal form. Then we turn the isotopy machine on again, to get another essential lamination $\mathcal{L}_2$ in normal form w.r.t. the triangulation. We continue this process (taking infinitely-many infinite isotopies), all the time watching the intersections of the laminations with the 1-cells of our decomposition. These turn out to 'almost' form a nested sequence, and by altering $\mathcal{L}_0$ slightly, can in fact be made into a nested sequence, whose intersection forms a set of stable points under these isotopies. Then, as in [B1], a new lamination can be seen to grow out of these points, which will turn out to be essential and in normal form w.r.t. the cell decomposition.

This paper uses to a large extent the same techniques that were developed in [B1], so a familiarity with that paper will be assumed. The reader is also referred to that paper and [G-O] for definitions and basic notions regarding essential laminations.

The author wishes to thank Kerry Jones for several rewarding conversations.

1. Proof of the theorem: the isotopies

Given a regular cell decomposition of M, we can subdivide the 3-cells in a natural way to get a triangulation τ of M whose 3-simplices are embedded, except possibly at their vertices, and so that the 1-skeleton of τ contains the 1-cells of M, with no subdivision - see Figure 1. To fix notation, we will call the i-skeleton of the original cell decomposition $\mathcal{K}^{(i)}$. The union of the 1- simplices of τ which are contained in the 2-skeleton of the original decomposition will be called $\tau_2^{(1)}=\tau^{(1)}\cap\mathcal{K}^{(2)}$; it therefore contains $\mathcal{K}^{(1)}$. The union of the 1-simplices of τ which meet the interior of the 3-cells of the original decomposition will be called $\tau_3^{(1)}$; see Figure 1.

Figure 1

We assume, as in [B1], that every leaf of $\mathcal{L}_0$ is non-compact, and further, by passing to a sublamination, that $\mathcal{L}_0$ is its own minimal sublamination, that is, every leaf of $\mathcal{L}_0$ is dense in $\mathcal{L}_0$. These assumptions are made largely for convenience, to streamline the argument; they do not weaken the results, because we are only interested in, and in general can only expect, an existence result, so our starting point is fairly arbitrary. Any compact leaf could be put into normal form using the usual techniques ([Ha], [Sch]).

Because the isotopies developed for triangulations in [B1] work perfectly well for a 'triangulation' with bad vertices (i.e., in which some of the 3-simplices have vertices identified), as the reader can readily check, we can run the isotopies described there on any essential lamination, to give an essential lamination in Haken normal form w.r.t. τ. In an abuse of notation we will call the end result of these 'isotopies' $\mathrm{I}_1(\mathcal{L}_0)=\mathcal{L}_1$, thinking in our minds that it is a lamination isotopic to $\mathcal{L}_0$ (instead of a sublamination of a splitting of the closure of the eventually stable portions of $\mathcal{L}_0$ under an infinite sequence of isotopies!).

This lamination, however, is almost certainly not in normal form w.r.t. the original cell decomposition. We next try to right all of the obvious flaws that this lamination has, as far as normal form for the decomposition is concerned.

The procedure of making a lamination meet an embedded 3-cell in normal disks is fairly straightforward, and entirely similar to that of a 3-simplex (see [B1] for a discussion of that case). The problem here is that our 3-cells in general are not embedded, which makes the notion of 'pushing your problems out' of the 3-cell intractible; things will generally get pushed back in, from an unexpected direction. However, we can start by identifying all of the things in a given 3-cell B_1 which we want to push, and which, pushed out of an embedded 3-cell, would give us a collection of normal disks, and push them, not worrying about what is getting pushed back in.

Figure 2

Number the 3-cells of M by $B_1,\ldots,B_r$. First notice that $\mathcal{L}_1$ meets each ∂B_k in a lamination consisting of circles, because ∂B_k is a 2- sphere; otherwise, $\mathcal{L}_1$ has a monogon or non-trivial holonomy around a null-homotopic loop. Start with B_1, and surger $\mathcal{L}_1$ along a 2-sphere in B_1 parallel to ∂B_1, and throw away any 2-sphere leaves created; as in [B1], this gives a lamination isotopic to $\mathcal{L}_1$, which meets B_1 in finitely-many parallel families of disks. Then we can locate a finite number of disjoint embedded '∂-compressing' disks $\Delta^2_1,\ldots,\Delta^2_n$ in B_1 (see Figure 2) for which $\mathcal{L}_1$ meets $B_1 \setminus \bigcup(\Delta^2_k \times I)$, in normal disks, where the '1-skeleton' of this ball is obtained from the 1-skeleton of B_1 by deleting the arcs $\Delta^2_k \cap \partial B_1 \subseteq \mathcal{K}^{(1)}$ and replacing them with $\Delta^2_k \cap \mathrm{int}(B_1)$ (see Figure 3). If we now,one at a time, isotope $\mathcal{L}_1$ using these disks as a guide, we will, as a result of this isotopy I^+_1 have removed points of intersection of $\mathcal{L}_1 \cap \partial\Delta^2_k \subseteq \mathcal{K}^{(1)}$ from $\tau^{(1)}_2$, at the expense of *adding* to the points of intersection of $\mathcal{L}_1$ with $\tau^{(1)}_3$, since both of these procedures could add points to the intersection of $\tau^{(1)}$ with the interior of the 3-cell. We now have a new lamination $\mathrm{I}_1^+(\mathcal{L}_1)=\mathcal{L}_1^+$ transverse to $\tau^{(1)}$, and we can use the machinery of [B1] to create an infinite isotopy I_2 to create a new essential lamination $\mathcal{L}_2$ which is once again in normal form w.r.t. τ. In this way we can create an *infinite collection of infinite isotopies* I_k, along with transition isotopies I_k^+, so that $\mathcal{L}_\mathrm{k}=\mathrm{I}_\mathrm{k}(\mathcal{L}_{\mathrm{k}-1}^+)$ is in normal form w.r.t. the triangulation τ, and $\mathcal{L}_\mathrm{k}^+=\mathrm{I}_\mathrm{k}^+(\mathcal{L}_\mathrm{k})$, where I_k^+ has attempted to make $\mathcal{L}_\mathrm{k}$ meet B_j in normal disks, where $\mathrm{k}\equiv\mathrm{j}(\mathrm{mod}\ \mathrm{r})$.

Figure 3

The isotopies I_k^+ really achieve very little in a practical sense, and it is likely the case that there will be no time during the sequence of isotopies in which $\mathcal{L}_\mathrm{k}$ will meet all of B_j in normal disks. However, we have won a moral victory - we know that everything that is bad about the intersections of our lamination $\mathcal{L}_\mathrm{k}$ with the 3-cells of M *eventually* moves. The other main benefit gained from this collection of isotopies is that we have arranged that $\mathcal{L}_\mathrm{k}\cap\tau^{(1)}$ is contained in $\mathcal{L}_{\mathrm{k}-1}^+\cap\tau^{(1)}$, modulo the splitting along finitely-many leaves which is required to obtain $\mathcal{L}_\mathrm{k}$ from $\overline{\mathrm{X}}_\mathrm{k}$; see [B1]. Also, $\mathcal{L}_\mathrm{k}^+\cap\tau_2^{(1)}\subseteq\mathcal{L}_\mathrm{k}\cap\tau_2^{(1)}$, although $\mathcal{L}_\mathrm{k}^+\cap\tau_3^{(1)}$ may be much larger than $\mathcal{L}_\mathrm{k}\cap\tau_3^{(1)}$, because the surgeries along the 2-spheres in the B_j will in general create many additional points. We will see in the next section, however, that by an a posteriori splitting of $\mathcal{L}_0$ along a *countable* number of leaves, we can arrange that $\mathcal{L}_\mathrm{k}\cap\tau^{(1)}\subseteq\mathcal{L}_{\mathrm{k}-1}^+\cap\tau^{(1)}$, on the nose. Then, out of the stable set - the intersection of these nested collections of points - stable normal disks for the cell decomposition begin to grow. Then we will once again be in the same situation encountered in [B1]; after altering the union of these stable disks, we are able to find a lamination in normal form w.r.t. the cell decomposition which, with some work, can be shown to be essential.

2. Making the isotopies accomodate one another

We have now defined our sequence of isotopies, which, up to splitting, only remove points of $\mathcal{L}_0 \cap \tau_2^{(1)}$, at the expense of possibly adding points to $\mathcal{L}_0 \cap \tau_3^{(1)}$. It is actually the case that there are points of $\mathcal{L}_0 \cap \tau_2^{(1)}$ which are never removed by these isotopies, although they are hard to see, because they are continually moved by the splittings. To enable ourselves to see them, let $\mathrm{p_k}$ be the projection which is the inverse of the splitting required to get (in the notation from [B1]) the lamination $\mathcal{L}_\mathrm{k}$ from the closure $\overline{\mathrm{X}}_\mathrm{k}$ of the stable disks of the k-th isotopy. Then $\mathrm{p_k}(\mathcal{L}_\mathrm{k} \cap \tau_2^{(1)}) = \overline{\mathrm{X}}_\mathrm{k} \cap \tau_2^{(1)} = \mathrm{X}_\mathrm{k} \cap \tau_2^{(1)} \subseteq \mathcal{L}_{\mathrm{k}-1} \cap \tau_2^{(1)}$. So setting $\pi_\mathrm{k} = \mathrm{p_k} \circ \cdots \circ \mathrm{p}_1$, we have $\pi_\mathrm{k}(\mathcal{L}_\mathrm{k} \cap \tau_2^{(1)}) \subseteq \pi_{\mathrm{k}-1}(\mathcal{L}_{\mathrm{k}-1} \cap \tau_2^{(1)}) \subseteq \mathcal{L}_0 \cap \tau_2^{(1)}$, and since each of the $\mathrm{p_k}$'s are closed maps, $\pi_\mathrm{k}(\mathcal{L}_\mathrm{k} \cap \tau_2^{(1)})$ is a closed set. So we have a nested sequence of closed sets in the compact set $\tau_2^{(1)}$; so either eventually they are all empty, or their intersection is non-empty.

But if one of these sets $\pi_\mathrm{k}(\mathcal{L}_\mathrm{k} \cap \tau_2^{(1)})$ is empty, then of course $\mathcal{L}_\mathrm{k} \cap \tau_2^{(1)}$ is empty. Then because $\mathcal{L}_\mathrm{k}$ is in normal form w.r.t. the triangulation τ, it cannot meet $\mathcal{K}^{(2)}$ either, because $\tau_2^{(1)} \subseteq \mathcal{K}^{(2)}$ cuts the 2-cells into 2-simplices, and any loop of $\mathcal{L}_\mathrm{k} \cap \mathcal{K}^{(2)}$ in a 2-simplex bounds a disk in both 3-simplices containing it, after the initial surgery- isotopy, hence is contained in a sphere leaf of $\mathcal{L}_\mathrm{k}$, a contradiction. So $\mathcal{L}_\mathrm{k}$ would be contained in the interior of the 3-cells, which is also a contradiction. The only alterative is that $\mathcal{L}_\mathrm{k}$ is in fact empty. But this is still another contradiction; the isotopies in [B1] always gave a non- empty lamination if they started with one, so by induction if $\mathcal{L}_0$ is non- empty, so are all of the $\mathcal{L}_\mathrm{k}$.

So now we know that something survives the isotopies, but we can't necessarily see it, because it might be forever moved by the splittings. We will now remedy this situation by an a posteriori splitting of $\mathcal{L}_0$.

This next step is motivated by the following observation. Suppose L is a leaf of X_k, the stable leaves of the isotopy I_k, which meets (and hence is contained in) one of the singular leaves of $\overline{X}_k$, but is isolated on one side, in the transverse direction. Then when the procedure of [B1] requires us to split along the leaf L and throw away anything which is isolated on both sides, this amounts to just taking the singular leaf of $\overline{X}_k$ and throwing away those portions of it which are isolated on both sides (see Figure 4). So, for example, if all of the leaves of $\mathcal{L}_{k-1}$ which give rise to singular leaves in $\overline{X}_k$ are isolated on one side, then we may take $\mathcal{L}_k \cap \tau_2^{(1)} \subseteq \mathcal{L}_{k-1} \cap \tau_2^{(1)}$. The intersections of the $\mathcal{L}_k$ with $\tau_2^{(1)}$ would be nested.

Figure 4

Consider the finite collection of all leaves of $\mathcal{L}_k$ which are not ordinary leaves of the isotopy I_k; that is, they are leaves of $\mathcal{L}_k$ which are not contained in X_k - they arise by splitting. Under π_k, these leaves are mapped into, but not necessarily onto, a finite number of leaves of $\mathcal{L}_0$. Taking the union over all k, we get a countable collection $\mathcal{S}'$ of leaves of $\mathcal{L}_0$. Now among these, choose those which are limited upon on both sides, and call this collection $\mathcal{S}$. Split $\mathcal{L}_0$ along these leaves; we still call this new lamination $\mathcal{L}_0$. This is in principle a splitting along a countable number of leaves; but it can be carried out in entirely the same spirit as the splitting of a lamination along a single leaf.

Splitting along a single leaf L amounts to replacing $\mathcal{L}_0$ in a branched surface neighborhood N(B) by $(\mathcal{L}_0 \setminus L) \cup \partial N(L)$ in $N = (N(B) \setminus L) \cup N(L)$ (see Figure 5). By making sure that the lengths of the I-fibers of N(L) tend to 0 fast enough as we tend to ∞ in L, we can insure that N is (I- fiber-preservingly) homeomorphic to N(B); this amounts to insuring that the sum of the lengths of the I-fibers in N(L) added to each I-fiber of N(B) is finite.

Figure 5

For a countable number of leaves L_i, we must then require that the maximum amounts of 'air' that each splitting adds to some I-fiber of N(B) together form a convergent series. The sum, by scaling the splittings, we can then make as small as we like. This gives us a lamination - the resulting set of 'leaves' forms a closed set, because it meets the I-fibers in closed sets - in (something homeomorphic to) N(B). By [G-O], this new lamination, which we will still call $\mathcal{L}_0$, is therefore also essential, since it, too, is carried with full support by the (we may assume) essential branched surface B.

Recall that a monogon number for a lamination $\mathcal{L}_0$ carried by a branched surface B w.r.t. an arc or loop γ, transverse to $\mathcal{L}_0$ and meeting N(B) in I-fibers, is a number ϵ so that any two points in $\mathcal{L}_0 \cap \gamma$ within ϵ of one another are in the same vertical fiber of N(B). Now if we assume that, by scaling the splittings, our collection of splittings moves points in N(B) by at most $\epsilon/3$, then $\epsilon/3$ is a monogon number for our new $\mathcal{L}_0$ w.r.t. B. This is because any two points within $\epsilon/3$ of one another, when we collapse the splittings to retrieve the old $\mathcal{L}_0$, are then within ϵ of one another, so are contained in the same fiber of N(B).

After carrying out this splitting, we can now run the exact same isotopies that we had previously built, which allowed us to identify the leaves of $\mathcal{S}$; but where previously we had an arc or disk in one of the L_i which needed pushing, now we find two parallel arcs or disks, both of which we push. But now any time one of the isotopies requires a splitting, the leaf we split along has already been split, so it is isolated on one side. Then the splitting is really just an erasure, so the intersections of the associated $\mathcal{L}_k$ with $\tau_2^{(1)}$ will be nested.

3. Stability

We are now in a position to establish the basic stability results analogous to [B1]. We have, beginning with an essential lamination $\mathcal{L}_0$, an infinite collection $\{\mathrm{I}_\mathrm{k}\}$ of infinite isotopies giving essential laminations $\mathcal{L}_\mathrm{k}$ in normal form w.r.t. the triangulation τ, and transition isotopies $\{\mathrm{I}^+_\mathrm{k}\}$ of $\mathcal{L}_\mathrm{k}$ which attempt to put the lamination $\mathcal{L}_\mathrm{k}$ into normal form w.r.t. the cell decomposition of which τ is a subdivision. These isotopies are compatible, in the sense that, for every k,

$$\mathrm{I}_{\mathrm{k}+1}(\mathcal{L}^+_\mathrm{k})\cap\tau_2^{(1)}=\mathcal{L}_{\mathrm{k}+1}\cap\tau_2^{(1)}\subseteq\mathrm{I}^+_\mathrm{k}(\mathcal{L}_\mathrm{k})\cap\tau_2^{(1)}=\mathcal{L}^+_\mathrm{k}\cap\tau_2^{(1)}\subseteq\mathrm{I}_\mathrm{k}(\mathcal{L}^+_{\mathrm{k}-1})\cap\tau_2^{(1)}=\mathcal{L}_\mathrm{k}\cap\tau_2^{(1)}.$$

In general, however, $\mathcal{L}_\mathrm{k}\cap\tau_3^{(1)}$ is not under such good control. We can assume, by passing to a sublamination at each stage, that every leaf of the $\mathcal{L}_\mathrm{k}$ is dense in $\mathcal{L}_\mathrm{k}$, and we can assume that $\mathcal{L}_\mathrm{k}$ doesn't contain, hence isn't, a compact leaf - otherwise, we could just quit and put it into normal form in the

As before, the set of points $\mathrm{P}=\bigcap_{k\geq1}(\mathcal{L}_\mathrm{k}\cap\tau_2^{(1)})$ must be non- empty, since each $\mathcal{L}_\mathrm{k}$ is non-empty, and it consists of (not necessarily all) points of $\mathcal{L}_0\cap\tau_2^{(1)}$ which are never moved under any of the isotopies I_k and I^+_k. We will now study the intersections of the $\mathcal{L}_\mathrm{k}$ with the 2-skeleton $\mathcal{K}^{(2)}$ of M, and see, as in [B1], that stable arcs in the 2-cells begin to grow out of each of these stable points. These arcs will together form the boundaries of disks in the 3-cells which are themselves eventually stable. The arguments are very similar to those given in [B1], except for the additional case when the point of P is in a splitting leaf for one of the isotopies.

Now consider a point x$\in$P and a 2-simplex Δ^2 of $\tau^{(2)}\cap\mathcal{K}^{(2)}$ containing x in its boundary. For each k, x is contained in an arc α_k of $\mathcal{L}_\mathrm{k}\cap\Delta^2$. One end of α_k is fixed (it is x), while the other changes, as α_k 'grows', by boundary-compressions, as in [B1]. We should note that the proof in [B1] of eventual stability of the other endpoint, under the isotopy

I_1, used only the existence of a monogon number for the lamination $\mathcal{L}_0$, and gave an upper bound N, depending only on this monogon number, for the number of ∂-compressions that the end of a half-anchored arc can undergo. But this upper bound can be chosen to be universal over the entire collection of isotopies, at least for ordinary leaves of the $\mathcal{L}_k$; because if an arc with one endpoint in P undergoes more than N ∂-compressions in the first k isotopies (or transition isotopies), then the exact same argument will allow us to find two points of the arc too close to one another, which give a monogon in the 2-simplex. We can then let this monogon flow back under all of the isotopies; this is really only a finite amount of flowing, because in each isotopy the arc is eventually stable. This will produce an arc in a leaf of $\mathcal{L}_0$ which, because the ends are too close together, gives a monogon for $\mathcal{L}_0$, a contradiction. If we are dealing with an arc in some split-and-paste leaf, we choose a nearby arc in an ordinary leaf to apply this to; our original arc might not be able to flow all of the way back. In particular, an arc, one of whose endpoints is a point of P which is contained in an ordinary leaf of all of the laminations $\mathcal{L}_k$, is eventually stable.

Unfortunately, we cannot guarantee that such points exist; it could be the case that every point of P is contained in a splitting leaf of some $\mathcal{L}_k$. Then the argument above might not work - the arc α_k may 'jump' to the other side of the 2-simplex (see Figure 6), instead of being ∂-compressed. However, the arc α_k can jump only in the direction in which the leaf containing it is *not isolated*, i.e., only towards the side in which it is limited upon by other leaves (see Figure 6). So if this arc in a splitting leaf is not eventually stable, it is either ∂-compressed infinitely-often, which could be detected as above, or it has to jump infinitely-often, always compressing less than N-times in between, in order not to be detected as the ordinary leaves were. But because it can jump in only one direction, it has to be ∂-compressed back at some stage. So if the arc α_k jumps N times, it is also

∂-compressed at least N times. But at whatever stage so many jumps have occurred (say $\mathcal{L}_k$), there is an arc in a nearby leaf, which is ordinary for all of the isotopies up to I_k - there are, after all, uncountably-many ordinary leaves in any neighborhood of α_k. This arc we can think of as having had its end near x anchored throughout these isotopies, and has been ∂-compressed at least N-times; if it is close enough to α_k , it has been ∂-compressed each time that α_k has been. But an ordinary arc cannot be ∂-compressed this often; so the arc α_k must in fact eventually stabilize, too.

Figure 6

Now the arguments from [B1] can be applied. These stable arcs start gluing end-to-end around the boundaries of the 3-cells, and so must eventually close up - otherwise the union of the arcs will be forced to wander from 2-cell to 2-cell and so eventually cross a 1-cell in $\mathcal{K}^{(1)}$ twice. This is because these arcs cannot wander in the interior of a 2-cell indefinitely. In fact, such a collection of arcs can consist of only n arcs, where n=the number of 2- simplices in the 2-cell we are wandering in, before it violates either the normality of $\mathcal{L}_k$ w.r.t τ, or will be pushed by one of the transitional isotopies I_k^+ (see Figure 7).

But then the transitional isotopy I_k^+, for some k, would try to push a stable arc, which is impossible. Therefore the arcs eventually grow together into stable loops in the boundaries of the 3-cells, which are in normal form, and eventually are bounded by disks in some $\mathcal{L}_k$. These disks are stable - their boundaries no longer move, and the next isotopy makes it parallel to ∂B^3, hence a union of stable normal disks. The union of these stable disks is an object we will call X.

Figure 7

4. The normal lamination

Now the argument finishes in much the same manner as in [B1]. X is a collection of 1-to-1 immersed surfaces with $X \cap \mathcal{K}^{(1)}$ closed in $\mathcal{K}^{(1)}$, and which meets each 3-cell B_k in normal disks. $\overline{X}$ is in general not a lamination, as a result of bad limiting behavior similar to that encountered in [B1]. This, however, occurs only finitely-often (see Figure 8); there can of course be more than two normal disks in each 3-cell added to X to give $\overline{X}$, but, as in [Br1], these disks occur only where different normal disk types meet. After a smoothing operation (as in [B1]) we can make the non-manifold points of $\overline{X}$ consist of loops in leaves of X; the same parity condition which allowed us to smooth away high-valency vertices of the graph , $=X \cap$(the added normal disks) will persist. We then split $\overline{X}$ along the finitely-many singular leaves of $\overline{X}$ to get a lamination $\mathcal{L}'$. Then, as in [B1], we have the following facts:

Figure 8

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

Proof: Let L be a leaf of X and γ a loop in X which is immersed, transverse to itself, and is null-homotopic in M. Then γ meets only finitely-many of the normal disks of L, so is contained a leaf L_k of $\mathcal{L}_k$ for some k. It is therefore null-homotopic in L_k, and in fact if we let A be the union of a neighborhood of γ in L_k together with all of the pieces of the complement of this neighborhood which are disks, then γ is null- homotopic in A, because then all of the components of ∂A π_1-inject into $L_k \setminus A$. $L_k \setminus \text{int}(N(\gamma))$ has some finite number $n=n(k)$ of disk components, and none of these disk components D, $\partial D=\beta$, can be later replaced by a non-disk one in some later $\mathcal{L}_j$. For if this were to occur it would do so by the splitting of some $\overline{X}_j$ to get $\mathcal{L}_j$. But because the boundary of the disk is stable - all of $N(\gamma)$ is - and its boundary has no holonomy, the disk can be lifted to nearby ordinary leaves for the isotopy I_j, and because these loops are eventually stable, they bound disks in the ordinary leaves, which, because they are compact, also eventually stabilize under I_j.

But because $\mathcal{L}_j$ is essential, β bounds a disk in its leaf; since by assumption it doesn't bound a disk on the side it used to, it now bounds one on the 'other' side. But then this disk can be lifted to nearby ordinary leaves, implying that the lifts of β bound disks on both sides. So nearby ordinary leaves are 2-spheres, contradicting the essentiality of $\mathcal{L}_j$, because these spheres consist of only finitely-many normal disks, so they stabilize in finite time, so are isotopic to leaves of $\mathcal{L}_{j-1}$.

Therefore this number n(k) can only increase with k, and is bounded above by the number of components of $\partial N(\gamma)$, so n(k) eventually stabilizes. This collection of disk components of the complement of $N(\gamma)$ then later stabilize also; they could a priori continue jumping around, if they were in a splitting leaf, but to do so infinitely often would require some arc with one anchored end to jump infinitely often, which is impossible.

As in [B1], every cut-and-paste leaf of $\mathcal{L}'$ is limited upon by ordinary leaves of $\mathcal{L}'$; otherwise $\mathcal{L}'$ would have only finitely many leaves, all of which were limit leaves, giving rise, transversely, to non-empty closed countable perfect sets, an impossibility. We also have:

Lemma: If γ is a loop in a leaf L of $\mathcal{L}'$ which bounds a disk D in $M|\mathcal{L}'$, then L has trivial holonomy around γ.

Proof: The proof is the same as in [B1], although the notation means something different - we isotope γ so that it meets $\mathcal{K}^{(1)}$ in a point of P and look in the normal fence over γ. Non-trivial holonomy would imply the existence of a (possibly different loop γ_0 in a leaf of $\mathcal{L}'$, with an infinite ray in an ordinary leaf spiralling down towards it. Eventually, though, two points of this ray must be within ϵ of one another along the the 1-cell of $\mathcal{K}^{(1)}$

meeting γ. The compact piece β of the arc between them meets only finitely-many normal disks and so is contained in a leaf of $\mathcal{L}_k$ for some k. If we graft the normal fence onto the disk D and make this disk D^+ transverse to $\mathcal{L}_k$ (rel β), then the arc β will be contained in a leaf ℓ_0 of $\mathcal{L}_k \cap D^+$; see Figure 9. Then, as in [B1], we can use this leaf to find a disk D_0 with $\partial D_0 = \alpha_0 \cup \beta_0$, where $\beta_0 \subseteq \ell_0$, α_0 has length less than ϵ, and $\partial\alpha_0 \subseteq \mathcal{K}^{(1)}$. But then letting a nearby arc, which is contained in a leaf which was never split under any of the previous isotopies, flow back along the isotopies I_j, $j \leq k$, gives an arc in $\mathcal{L}_0$ which violates the monogon number ϵ for $\mathcal{L}_0$. Therefore there can be no holonomy around γ.

Figure 9

These three facts are the ones which were used in [B1] to show that $\mathcal{L}'$ has a sublamination $\mathcal{L}$ which is essential. Using the same procedure on our lamination $\mathcal{L}'$, we can conclude that $\mathcal{L}'$ contains an essential sublamination $\mathcal{L}$. Because $\mathcal{L}'$ was a union of normal disks, this lamination $\mathcal{L}$ is in normal form w.r.t. the cell decomposition of M.

5. Heegard decompositions

Given a Heegard decomposition $M = H_1 \cup H_2$ of M, H_i=handlebodies of genus g, there is a standard way to obtain a regular cell decomposition, given as a handle decomposition, for M with H_1=a neighborhood of the 1-skeleton=the union of a 0-handle and a finite number of 1-handles, and H_2=the union of 2-handles and one 3-handle. Meridians for the 1-handles of H_2 are thought of as giving attaching maps for the 2-handles. Given an essential lamination $\mathcal{L}_0$, the above result allows us to find an essential lamination $\mathcal{L}$ which is in normal form w.r.t. this cell decomposition. What we will show is that interpreted as being a lamination transverse to the surface $F = \partial H_2 = \partial H_2 = \partial N(\mathcal{K}^{(1)})$, this lamination is in 'normal form' w.r.t. the Heegard decomposition, in the sense of the theorem stated in the introduction. The first two properties are immediate, provided we think of H_1 as being a very small regular neighborhood of $\mathcal{K}^{(1)}$, because $\mathcal{L}$ is transverse to $\mathcal{K}^{(1)}$. It is the third and last property which will need to appeal to the normality of $\mathcal{L}$ w.r.t. the 3- cells of M.

To show that the third property holds we will show that, for $\mathcal{L}_2{=}\mathcal{L}\cap H_2$, $M_0{=}H_2|\mathcal{L}_2$ has incompressible boundary. Because $\mathcal{L}_2$ must be end- incompressible - otherwise $\mathcal{L}$ would not be, either - and has no leaves which are spheres or compressible tori, the proof of [G-O, Theorem 1(a)] will show that $\mathcal{L}_2$ has π_1-injective leaves. So suppose D is a disk in $H_2|\mathcal{L}_2{=}M_0$, $\partial D{=}\gamma{\subseteq}L{\subseteq}\mathcal{L}_2$; we will show that it is isotopic in M_0, rel ∂D, to a disk in $\mathcal{L}_2$.

Make γ and D transverse to the 2-cells $\{D_i\}$ of M; then $D\cap D_i$ consists of a finite number of circles and arcs, for each i, and by disk-swapping we can remove all of the circles of intersection. Now choose an outermost arc α of $D\cap D_i$, cutting off a disk $\Delta{\subseteq}D$ which meets γ in an arc γ_i (see Figure 10).

Figure 10

Because Δ misses the 2-cells D_i except on its boundary, it is contained in one of the 3-cells B_j, and so γ_i is contined in a normal disk D' in B_j. γ_i cuts this disk into two disks, one of which, Δ', misses the 1-skeleton of M; this is because D' is normal and the endpoints of α are contained in the same face of B_j. But then Δ together with Δ' form a disk in D_j with boundary in a single face of D_j, and this cuts off a 3-ball from D_j which misses the 1-skeleton of M. This ball gives us a way to isotope D, keeping it's boundary in L, to remove the arc α from $D\cap D_i$. Continuing inductively, inducting on the number of arcs in the intersection of D with the D_i, we can find an isotopy of γ in $L\backslash\mathcal{K}^{(1)}$ to a loop bounding a disk D which misses the 2-cells. But then this loop is contained in a normal disk which is contained in L, and D is isotopic to the disk that γ cuts off in this normal disk. Letting things flow back along the isotopy, we see that our original loop γ bounds a disk in L, which, because the isotopy never forced γ to cross the 1-skeleton, misses $\mathcal{K}^{(1)}$, i.e., is contained in $\mathcal{L}_2$. Therefore $H_2|\mathcal{L}_2$ has incompressible boundary.

We should note that the lamination $\mathcal{L}_2 \subseteq H_2$ <u>cannot</u> be essential; a lamination in a handlebody, other than a collection of compressing disks, can never be ∂-incompressible. For an essential lamination in a handlebody could be successively pushed off of a set of core disks for the 1-handles of H_2, by the standard disk-swapping techniques, and so would live in a 3-ball.

6. Concluding remarks

This paper marks a first step toward finding an algorithm to determine if an irreducible 3-manifold M contains an essential lamination. This paper reduces the problem to determining if, w.r.t. some given cell decomposition of M, M contains a normal essential lamination, i.e., one carried with full support by one of a finite number (see [F-O]) of 'normal' branched surfaces.

To find such an algorithm, it seems that two pieces remain. The first is to generate a new finite list of branched surfaces, which are essential, and, assuming there <u>is</u> an essential lamination, at least one of which carries a lamination with full support. (This sentence is slightly strange, since, technically, one of the conditions of essentiality is that a branched surface carry a lamination with full support - we therefore mean essential *minus* this condition.) Techniques exist, largely using Haken's normal surface theory [Ha], to determine (more or less) algorithmically if a 'normal' branched surface is essential (in this sense), so this step appears to be within the reach of the present technology.

The second step is to find an algorithm to determine when a branched surface carries a lamination with full support. This appears to be far more elusive. Sufficient conditions exist (e.g., [B2], [C1]), as well as necessary ones (e.g., [G-O]); the difficulty, of course, lies in achieving both at the same time! To date, work of Joe Christy [C2] seems to give the best candidate.

With these pieces in place, we could construct our algorithm. If an essential lamination existed, then one of our finite number of 'normal' essential branched surfaces could be shown to carry a lamination with full support. If however, none of them did, then M would contain no essential laminations.

This paper also gives further evidence that the notion of an infinite isotopy will be an increasingly useful tool in controlling essential laminations in a wider variety of contexts. For example, with similar techniques it might be possible to show that a knot κ in a 3-manifold M can be pulled taut w.r.t. some essential lamination $\mathcal{L}$; that is, that for some lamination $\mathcal{L}$, $\mathcal{L}\setminus\text{int}(N(\kappa))$ is essential in $M\setminus\text{int}(N(\kappa))$. This would be a large step towards answering the question of what manifolds obtained by Dehn surgery on a knot contain essential laminations. It is interesting to note that the corresponding 'normal form' problem for incompressible surfaces is, like Haken norm form, both short and easy.

Bibliography

[B1] Brittenham, *Essential laminations and Haken normal form*,

[B2] Brittenham, *A sufficient condition for a branched surface to fully carry a lamination*, in preparation

[C1] Christy, *Branched surfaces and attractors*, I:*dynamic branched surfaces*, Trans. AMS (to appear)

[C2] Christy, personal communication

[F-O] Floyd and Oertel, *Incompressible surfaces via branched surfaces*, Topology **23** (1984), 117-125

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

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

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