Atoroidal surface bundles over surfaces

Brian H. Bowditch
Mathematics Institute, University of Warwick,
Coventry CV4 7AL, Great Britain.

[First draft, June 2007. Revised: October 2009]

0. Introduction.

The main aim of this paper is to prove a finiteness result for atoroidal surface bundles over surfaces. It can be viewed from a number of different perspectives, and one can give several essentially equivalent statements. This, and related questions, have been considered by a number of authors. See [R] for a recent survey.

First, we express it in group theoretical terms. By a *surface group* we mean the fundamental group of a closed orientable surface other than the 2-sphere. By a *surface-by-surface group*, we mean a group G, with a normal subgroup, N, such that both N and G/N are surface groups. We refer to these as the *fibre* and *base* groups respectively. One can show that there are only finitely many choices of fibre for a given such group [J,Hi]. However, it will be convenient for us to regard the fibre subgroup as constituting part of the structure of a surface-by-surface group, and so we can bypass this particular issue. We refer to the genera of the respective surfaces as the *fibre genus* and *base genus*. We say that such group is *atoroidal* if it contains no free abelian subgroup of rank 2. We show:

Theorem 0 : *There are only finitely many isomorphism classes of atoroidal surface-by-surface groups with any given base and fibre genera.*

Here an isomorphism is assumed to respect the fibre subgroups. (Though by the earlier observation, this does not alter the strength of the statement.) Our argument does not give a computable bound on this number.

Daniel Groves informs me that he has an independent argument for Theorem 0, as well as the more general statement for one-ended groups, which we discuss later.

It is important to note that, in fact, no atoroidal surface-by-surface groups are known to exist at all. This is a well known question, which is the main theme of the survey [R]. It is easily seen, however, that atoroidal surface-by-surface groups cannot exist if either the fibre or base has genus 1. In the former case, it is sufficient to note that $SL(2, \mathbf{Z})$ has no surface subgroup. In the latter case, we would get a rank-2 free abelian subgroup of the mapping class group. However, the centraliser of any pseudoanosov element is virtually cyclic, and so any free abelian subgroup must contain a non-pseudoanosov element, which contradicts the atoroidal assumption. Therefore, in this paper, we shall be assuming that both fibre and base genera are at least 2. We remark that one can eliminate the orientability assumption for the base surface. This is a consequence of a more general result as discussed in [Bow5] (see also [DaF]). It seems likely that one can also remove the orientability assumption for the fibre, though since many of the available results for the curve complex etc. refer only to orientable surface, we shall not consider that case here.

Recall that the mapping class group of an orientable surface can be identified with the outer automorphism group of the fundamental group. (We are allowing orientation reversing mapping classes.) Any surface-by-surface group gives us a monodromy homomorphism from the base group, Γ, to the mapping class group, $\mathcal{M}$, of the fibre surface. Conversely, one can reconstruct a unique surface-by-surface group from such a homomorphism (see [Hi], or the discussion in Section 1). In the atoroidal case this homomorphism is injective, and its image is "purely pseudoanosov", that is, every non-trivial element is pseudoanosov. Thus Theorem 0 is equivalent to saying that there are only finitely many conjugacy classes of purely pseudoanosov surface subgroups of $\mathcal{M}$ of given genus.

The existence of surface subgroups is already non-trivial (see [GoH]). Examples of surface subgroups that are "almost" purely pseudoanosov are constructed in [LR]. From this point of view, the theory is parallel to the theory of kleinian groups and that of Gromov hyperbolic groups [Gr]. There are analogues of Theorem 0 in the context of hyperbolic and relatively hyperbolic groups, see [De,Da]. Some general aspects of the geometry of subgroups of $\mathcal{M}$ are discussed in [Mo].

In this language, we can generalise Theorem 0 to the case where Γ is a finitely presented one-ended group. The statement again says that there are only finitely many copies of Γ in $\mathcal{M}$ up to conjugacy. This statement is proven in [Bow5] using Theorem 8.1 of this paper. In fact, Dahmani and Fujiwara have independently obtained this result from Theorem 8.1 by different methods. As mentioned earlier, Daniel Groves informs me that he has a different approach to this.

Theorem 0 can also be viewed from the perspective of 4-manifolds. Detailed background material can be found in [Hi]. Here we are interested in smooth surface bundles over surfaces. Their fundamental groups are surface-by-surface groups. Indeed every surface-by-surface group arises in this way. Theorem 0 then states that there are only finitely many such bundles up to homotopy equivalence respecting the fibre class. In fact such a homotopy equivalence implies the existence of a bundle isomorphism. In particular, two equivalent bundles are diffeomorphic. We remark that any closed smooth 4-manifold with surface-by-surface fundamental group is topologically s-cobordant to such a surface bundle. See [Hi] for a discussion of these matters.

There are many related questions concerning geometric, symplectic and complex structures on surface bundles, see [Hi,R] and further references therein. For example it is not known if any hyperbolic 4-manifold can be a surface bundle. Indeed it appears to be unknown if every hyperbolic 4-manifold is finitely covered by a surface bundle (cf. the virtual fibre conjecture for hyperbolic 3-manifolds). Certainly, there can only be finitely many surface bundles admitting a hyperbolic structure for fixed genera, since the genera determine the Euler characteristic, hence the volume, and there are only finitely many hyperbolic 4-manifolds of any given volume [W]. (This can be generalised with little change to 4-manifolds of uniformly pinched negative curvature.) It is also known that a surface bundle cannot admit any complex hyperbolic structure [K]. We also note an analogous statement to Theorem 0 from complex analysis, which is discussed in [Mc]. This deals with holomorphic fibrations, where "atoroidal" is replaced by a notion of "truly varying", and the base is assumed to have fixed conformal structure.

Our point of view in this paper ties in with the large scale geometry of Teichmüller space. This general area has seen a great deal of activity recently, notably linked to the curve complex [Ha], and its application to the solution of Thurston's Ending Lamination Conjecture for hyperbolic 3-manifolds [MaM1,MaM2,Mi2,BrCM]. We make much use of these ideas. We base our arguments on a simplified presentation of some of that work to be found in [Bow4].

Although most of the argument is combinatorial, we need to make an appeal to hyperbolic geometry in Section 6. This is needed to find a suitable basepoint for our construction, that is, a set of curves filling the fibre surface with respect to which the monodomies are controlled. This introduces a non-constructive element into the proceedings. It would be nice to find a means of bypassing this, but it is complicated by the fact that, a-priori, a combinatorial model for a hyperbolic 3-manifold is not canonical, and so cannot be assumed invariant under a cyclic monodromy action.

I am grateful to Thomas Delzant for some early discussions about this, and related, questions while I was visiting Strasbourg. The first draft of this paper was written at the University of Southampton. I thank the referee for some very detailed comments and suggestions.

1. Fibrations.

We discuss the background to the main theorem, and its relation to 3-manifolds. We shall describe this mainly from the point of view of surface bundles over surfaces.

Going down a dimension, the case of surface fibrations over a circle are very much better understood. Let Σ be a closed surface, and let $\mathcal{M}$ be the mapping class group of Σ. Thus $\mathcal{M}$ is the group of self homeomorphisms of Σ (or equivalently, self homotopy equivalences) up to homotopy. Let $X(\Sigma)$ be the set of essential simple closed curves in Σ defined up to homotopy. Note that $\mathcal{M}$ acts on $X(\Sigma)$ with finite quotient.

Suppose that $\Sigma \longrightarrow N \longrightarrow S^1$ is a surface bundle. There is a natural cyclic cover $M \cong \Sigma \times \mathbf{R}$ so that $N = M/\mathbf{Z}$. If we choose a preferred homotopy class of homotopy equivalence $\Sigma \longrightarrow M$, then we get a homomorphism from $\mathbf{Z}$ to $\mathcal{M}$. The image of 1 gives the "monodromy" $\phi \in \mathcal{M}$ of N. We write $N = M_\phi$ — the mapping torus of ϕ. Thus M_ϕ is determined by ϕ up to homeomorphism commuting with projection to S^1. If we allow ourselves to vary the preferred class $\Sigma \longrightarrow M$, then it only depends on the conjugacy class of ϕ in $\mathcal{M}$.

If there is some $n \in \mathbf{Z}\setminus\{0\}$ and some $\alpha \in X(\Sigma)$ with $\phi^n\alpha = \alpha$, then we get an essential torus in M_ϕ, i.e. a π_1-injective map of $S^1 \times S^1$ into M_ϕ, and hence a $\mathbf{Z} \oplus \mathbf{Z}$ subgroup of $\pi_1(M_\phi)$. (We can assume that ϕ^n respects an orientation of α.) The converse to this is also true, and is not hard to see by "cut and paste" arguments, but it also follows from Theorem 1.1 below.

Key to understanding this is the Nielsen-Thurston classification of surface automorphism. We say that $\phi \in \mathcal{M}$ is *pseudoanosov* if given any $\alpha \in X(\Sigma)$ and $n \in \mathbf{Z}$ with $\phi^n\alpha = \alpha$, then $n = 0$.

The fibred case of Thurston's hyperbolisation theorem [T1,O1] states:

Theorem 1.1 : *If $\phi \in \mathcal{M}$ is pseudoanosov, then M_ϕ admits a (unique) hyperbolic structure.*

(This, of course, implies that M_ϕ is atoroidal.)

In this case, the cyclic cover, $M \cong \Sigma \times \mathbf{R}$ is a "doubly degenerate" kleinian group. Much is now understood about the geometry of such groups in the light of the work on the Ending Lamination Conjecture. This is discussed further in Section 5.

We now move on to consider surface bundles over surfaces. Suppose that S is a closed surface of genus at least 2, and that $\Sigma \longrightarrow \Pi \longrightarrow S$ is a Σ-bundle. Let $\Sigma \longrightarrow \tilde{\Pi} \longrightarrow \tilde{S}$ be the cover corresponding to the universal cover, $\tilde{S}$, of S. Thus the covering group, $\Gamma \cong \pi_1(S)$ acts on $\tilde{\Pi}$ with $\Pi = \tilde{\Pi}/\Gamma$.

Note that any essential curve γ in S (not necessarily simple) determines a 3-manifold, M_γ, fibring over the circle, together with preferred π_1-injective homotopy class of maps $M_\gamma \longrightarrow \Pi$. Clearly any essential torus in M_γ gives us an essential torus in Π.

In fact, any essential torus, $h : \Delta \longrightarrow \Pi$ arises in this way (up to homotopy). To see this, note that S is atoroidal, so we can homotope the projection of $h(\Delta)$ to S into a closed curve, γ in S. Since Σ is atoroidal, it follows that γ must be essential in S (otherwise we could homotope $h(\Delta)$ into a fibre of Π). We can now homotope h so that it factors through the map $M_\gamma \longrightarrow \Pi$ of the surface bundle M_γ. In particular, the monodromy of M_γ is not pseudoanosov.

Suppose now that we choose a preferred class of homotopy equivalence $\Sigma \longrightarrow \tilde{\Pi}$. We get a homomorphism $\phi : \Gamma \longrightarrow \mathcal{M}$, and hence an action of Γ on $X(\Sigma)$.

From the above discussion, we see:

Lemma 1.2 : *Π is atoroidal if and only if whenever $g \in \Gamma$ and $\alpha \in X(\Sigma)$ satisfy $\phi(g)\alpha = \alpha$, then g is the identity.*

In particular, the homomorphism ϕ is injective. This hypothesis on the monodromy homomorphism, or its image, is commonly termed *purely pseudoanosov*.

Suppose that $\Sigma \longrightarrow \Pi \longrightarrow S$ and $\Sigma' \longrightarrow \Pi' \longrightarrow S'$ are smooth surface bundles. We say that they are *bundle equivalent* if there is a diffeomorphism, $f : \Pi \longrightarrow \Pi'$ and $g : S \longrightarrow S'$ commuting with projections $\Pi \longrightarrow S$ and $\Pi' \longrightarrow S'$. If $S = S'$, we say they are strongly equivalent if g can be taken to be the identity. The notion of equivalence translates to isomorphism of fundamental groups respecting the fibre group (see [EE,Hi]). (Indeed, this is clear if we replace "diffeomorphism" by "homotopy equivalence", in the above. One needs more work to promote homotopy equivalence to diffeomorphism, as noted below.)

Clearly strongly equivalent bundles have monodromy homomorphisms that are conjugate in $\mathcal{M}$. (The conjugacy is determined by the choices of homotopy equivalence $\Sigma \longrightarrow \tilde{\Pi}$ and $\Sigma \longrightarrow \tilde{\Pi}'$.) Conversely, the following follows from [EE].

Theorem 1.3 : *Suppose that Π and Π' are Σ-bundles over S and that the monodromy homomorphisms of Γ to $\mathcal{M}$ are conjugate in $\mathcal{M}$. Then Π and Π' are strongly equivalent.*

We remark that if we are only interested in equivalence up to homotopy, this is much easier. By choosing the homotopy equivalances $\Sigma \longrightarrow \tilde{\Pi}$ and $\Sigma \longrightarrow \tilde{\Pi}'$ appropriately, we can assume these homomorphisms to be equal. One can now construct an equivalence on the complement of a fibre $\pi_S^{-1}x$, so that the map of base surfaces $g : S \setminus \{x\}$ is homotopic to the identity. To extend over the fibre up to homotopy, we use Lemma 1.4 below.

Lemma 1.4 : *Any hoomotopy equivalence $f : \Sigma \times \partial D \longrightarrow \Sigma \times \partial D$ extends to a homotopy equivalence $f : \Sigma \times D \longrightarrow \Sigma \times D$. Here D denotes the closed disc.*

Proof : It's not hard to see that $f(\{y\} \times \partial D)$ must be homotopically trivial in $\Sigma \times D$ for any $y \in \Sigma$. We thus extend f over $\{y\} \times D$, and then to all of $\Sigma \times D$. ♢

Taking D to be a small regular neighbourhood of $x \in S$, we can now extend our equivalence over the fibre.

To get a smooth equivalence one needs more. In [EE] is shown that the space of diffeomorphisms homotopic to the identity is contractible. From that, one can modify the above argument to give a diffeomorphism.

To deal with general equivalence of fibrations, we need to consider automorphisms of S as well as Σ. We note the well known fact that the mapping class group of S can be identified with the outer automorphism group of $\Gamma \cong \pi_1(S)$. From this, and Lemma 1.2, it is easily seen that two bundles are equivalent if and only if the monodromies $\Gamma \longrightarrow \mathcal{M}$ are equal up to automorphism of S and conjugacy in $\mathcal{M}$.

To apply this we make the following definition:

Definition : A *standard* set of generators of Γ is an ordered set, $g_1, g_2, \ldots, g_n$, where n is twice the genus of S and $[g_1, g_2][g_3, g_4] \cdots [g_{n-1}, g_n] = 1$.

There is exactly one standard set of generators up to automorphism of Γ.

We deduce:

Lemma 1.5 : *Suppose that Π and Π' are Σ-bundles over S and S' respectively, with monodromy homomorphisms $\phi : \pi_1(S) \longrightarrow \mathcal{M}$ and $\phi' : \pi_1(S') \longrightarrow \mathcal{M}$. Then Π and Π' are equivalent if and only if there are standard generating sets, $g_1, \ldots, g_n$ for $\pi_1(S)$ and $g'_1, \ldots, g'_n$ for $\pi_1(S')$ such that $\phi(g_i) = \phi'(g'_i)$ for all i.* ♢

The following is now a simple consequence:

Lemma 1.6 : *Suppose that $\mathcal{P}$ is a set of Σ-bundles over a closed surface of given genus. Suppose that there is a finite subset, $\mathcal{M}_0 \subseteq \mathcal{M}$ such that for any $\Pi \in \mathcal{P}$, there is a standard generating set, $g_1, \ldots, g_n$, of the fundamental group of the base surface, such that, for all i, $\phi(g_i) \in \mathcal{M}_0$, where ϕ is the monodromy homomorphism associated to Π. Then $\mathcal{P}$ meets only finitely many equivalence classes.* ♢

The aim of this paper is therefore to find such a set $\mathcal{M}_0$, when $\mathcal{P}$ is the set of all atoroidal fibrations.

We remark that, in fact, the assumption that the generating set is standard can be dropped from the hypotheses of Lemma 1.6. It is sufficient to show there only finitely many conjugacy classes of purely loxodromic subgroup. Once this is established, one could arbitrarily choose a standard generating set for each. Their monodromies therefore all lie in another finite set. However our methods give a standard generating set anyway, and perhaps the statement is more natural in this form.

2. Outline of the proof.

We begin by considering a more general set up.

Let Γ, H, be finitely generated groups, let $\mathcal{H}$ be a set of monomorphisms from Γ to H, and let A be a finite generating set of Γ. If we can find some finite subset $F \subseteq H$, such that $\phi(A) \subseteq F$ for all $\phi \in \mathcal{H}$, then clearly $\mathcal{H}$ must be finite. This is equivalent to putting a bound on the word length of $\phi(g)$ for all $\phi \in \mathcal{H}$ and all $g \in A$, where word length is measured in terms of a fixed finite generating set of H. Put another way, we are bounding the maximal displacement of the identity in the Cayley graph $\mathcal{G}(H)$ of H, by elements of this form. Typically, we will only want to control ϕ up to conjugacy in H (that is, up to postcomposition by an inner automorphism of H). So for any given ϕ, it is enough to find some vertex of $\mathcal{G}(H)$ whose maximal dispacement by $\phi(A)$ is bounded.

We can also allow ourselves to precompose by an element of $\mathrm{Aut}(\Gamma)$, which corresponds to changing the generating set. Let $\mathcal{A}$ be an $\mathrm{Aut}(\Gamma)$-invariant set of finite generating sets of Γ. For the purposes of showing that $\mathcal{H}$ is finite up to such pre and post composition, given $\phi \in \mathcal{H}$, it suffices to find some $A \in \mathcal{A}$ and some vertex $x \in V(\mathcal{G}(H))$ whose maximal displacement by $\phi(A)$ is bounded. This in turn shows that there are only finitely many possiblities for $\phi(\Gamma) \leq H$ up to conjugacy (cf. Lemma 1.6). In this setting, we do not need any control on $\mathcal{A}$ — the set of all finite generating sets of Γ will do. Retrospectively, we could reduce it to a set that is finite modulo $\mathrm{Aut}(\Gamma)$. In practice, however, we will need to find such a cofinite $\mathcal{A}$ a-priori. One way of expressing this, for general finitely presented groups, is in terms of bounding the "complexity" of a generating set, as in Section 8. However, in this section, and most of the paper, we will restrict our attention to when Γ is a surface group. We also remark that we can more conveniently work with $\mathrm{Out}(\Gamma)$ instead of $\mathrm{Aut}(\Gamma)$, since any inner automorphism of Γ can be pushed over to H.

To be more specific, let S be a closed orientable surface, and let $\Gamma = \pi_1(S)$. Let χ be a one-vertex triangulation, i.e. consisting of a single vertex at some $x_0 \in S$, together with a set, $E(\chi)$, of loops based at x_0. (See Section 7 for more discussion.) Each (oriented) $\epsilon \in E(\chi)$ determines an element $g(\epsilon) \in \Gamma$, and the set $A(\chi) = \{g(\epsilon) \mid \epsilon \in E(\chi)\}$ generates Γ. There are only finitely many possibilities for χ up to self-homeomorphism of S, hence $\mathcal{A}$ is finite modulo $\mathrm{Aut}(\Gamma)$, or modulo $\mathrm{Out}(\Gamma)$ neglecting basepoints.

We are interested in the case where H is the mapping class group of a surface. But to begin, we illustrate the idea with reference to a different but simpler case, namely the fundamental group of closed hyperbolic manifold. One can then progressively adapt the argument to hyperbolic groups, and finally via the curve complex, to the mapping class group.

Suppose that N is a closed hyperbolic manifold (of any dimension) and that $H = \pi_1(N)$. Here a monomorphism, ϕ, from Γ to H determines a homotopy class of maps from S to N. We use free homotopy since we are working modulo conjugation in H. We begin with an arbitrary one-vertex triangulation, χ_0. We can realise the homotopy class so that each triangle of χ_0 is sent to a geodesic triangle in N. (This is a well-known construction due to Thurston and Bonahon [T1,Bon].) The pull-back path-metric, ρ, on S, is hyperbolic, with cone singularity at the vertex x_0. By choosing an arbitrary edge of χ_0 and sending it to a closed geodesic in N, we can arrange that the cone angle there is at least 2π. The injectivity radius of (S, ρ) is bounded below by that of N. Now by Gauss-Bonnet its area and hence its diameter of S is bounded above. From this, it is not hard to construct another one-vertex triangulation, χ, all of whose edges have bounded ρ-length. Their images in N also have bounded length, and so the word length of each element of $\phi(A(\chi))$ in H is bounded. Thus, there are only finitely many isomorphic copies of Γ in in H.

The same holds true if H is any given Gromov hyperbolic group, though one needs to phrase the argument differently. (See for example, the more general statement in [De].) One can also use a combinatorial argument involving "carrying graphs" as in [Ba]. Thus, if $\mathcal{G}(H)$ is the Cayley graph of H, and ϕ is any monomorphism from Γ to H, we can find a one-vertex triangulation, χ, and a vertex of $\mathcal{G}$, moved a bounded distance by $\phi(A(\chi))$. (In Section 1, this is phrased in terms of "standard" generating sets, but it boils down to the same thing.) The key fact, that makes this work is that the stable length, $||h||$, is bounded below by a positive constant for each infinite order $h \in H$. (See Section 3 or definitions.) This plays the role of the lower bound on injectivity radius featuring in the earlier argument.

In fact, the result of [Ba] is much more general. Suppose $\Gamma = \pi_1(S)$ acts on any hyperbolic graph (not necessarily locally finite). We write $\mathrm{inj}(\mathcal{G}, \Gamma) = \inf\{||g|| \mid g \in \Gamma \setminus \{1\}\}$. The following formulation is given as Theorem 3.3 in the paper.

Theorem 2.1 : [Ba] : *If* $\mathrm{inj}(\mathcal{G}, \Gamma) = \eta > 0$*, then there is a one-vertex triangulation* χ *and some point* $x \in \mathcal{G}$ *displaced a bounded distance by each element of* $A(\chi)$*, where the bound depends only on* $\mathrm{genus}(S)$ *and* η*.*

Now let Σ be another closed orientable surface, and let $\mathcal{G} = \mathcal{G}(\Sigma)$ be the curve graph of Σ (see Section 3). This admits an action of the mapping class group, $\mathcal{M}$. The following (given as Theorems 3.1 and 3.2 here) was proven in [MaM1]:

Theorem 2.2 : [MaM1] : *For any given* Σ*, the graph* $\mathcal{G}$ *is hyperbolic, and* $||h||$ *is bounded below for all pseudoanosov* $h \in \mathcal{M}$*.*

Thus, if $\phi : \Gamma \longrightarrow \mathcal{M}$ is a purely pseudoanosov homomorphism, we get an action of Γ on $\mathcal{G}$ with $\mathrm{inj}(\mathcal{G}, \Gamma)$ bounded below. Thus, there is some χ and a curve $\alpha \in V(\mathcal{G})$, so that $d(\alpha, \theta\alpha)$ is bounded above for each monodromy $\theta \in \phi(A(\chi))$. Note that the same is true of $d(a, \theta a)$ where a is any set of curves of bounded diameter and a bounded distance from α in $\mathcal{G}$. We thus have quite bit of flexibility in the choice of "basepoint" a, and we will need to narrow this down.

To reduce to a finite set of monodromies, we need to bound the geometric intersection number, $\iota(a, \theta a)$, for each $\theta \in \phi(A(\chi))$, where a is a marking of Σ, as defined in Section 4. (The essential point about a marking is that it has bounded self intersection and cuts Σ into discs.) By Lemma 4.4, this bounds the word length of each such θ in $\mathcal{M}$. We have already bounded distances in $\mathcal{G}$, so we need to compare this with intersection number. This brings us to "subsurface projections", as defined in [MaM2] and discussed in Section 4.

Let Φ be a subsurface of Σ (possibly an annulus). Suppose that the curves, $\alpha, \beta \in V(\mathcal{G}(\Sigma))$ cross, or lie in, Φ. One can define the "projected distance", $p(\Phi; \alpha, \beta)$, between α and β in Φ. This satisfies a coarse triangle inquality, namely $p(\Phi; \alpha, \gamma) \leq p(\Phi; \alpha, \beta) + p(\Phi; \beta, \gamma) + 3$ for all $\alpha, \beta, \gamma \in X(\Sigma)$. Note $p(\Sigma; \alpha, \beta) = d(\alpha, \beta)$ is just distance in $\mathcal{G}(\Sigma)$. The following follows from [MaM2] and is discussed in Section 4 (see Proposition 4.3) :

Theorem 2.3 : [MaM2] *If $\alpha, \beta \in X(\Sigma)$ with $p(\Phi; \alpha, \beta)$ at most some constant q for all subsurfaces Φ, then $\iota(\alpha, \beta)$ is bounded above by some function of $\kappa(\Sigma)$ and q.*

In other words, if $d(\alpha, \beta)$ is bounded, but $\iota(\alpha, \beta)$ is large, then there must be a proper subsurface, Φ, with $p(\Phi; \alpha, \beta)$ large. There may be many such subsurfaces, but if we only consider those Φ of maximal topological complexity such that $p(\Phi; \alpha, \beta)$ is large, then we are reduced to a bounded number of possibilities (see Proposition 4.6). Of course, we will need to properly quantify this statement — there may be different degrees of largeness and boundedness involved. Note that in the above, we could replace α, β by any sets, $a, b \subseteq X(\Sigma)$ of bounded self-intersection, in particular to multicurves or markings. Then $p(\Phi; a, b)$ is well defined up to an additive constant and we still have a coarse triangle inequality.

Let us now return to the above set-up where we have $d(a, \theta a)$ bounded for all $\theta \in \phi(A(\chi))$. If each $\iota(a, \theta a)$ is also bounded, we are done. If not, then there is some $\epsilon \in E(\chi)$ and subsurface, Φ, with $p(\Phi; a, \theta a)$ large, where $\theta = \phi(g(\epsilon))$. We say that ϵ is "bent" along Φ. We choose ϵ and Φ so that Φ has maximal possible complexity. Thus, each edge of χ can be bent along only boundedly many subsurfaces of this complexity. (Properly formulated, we will need to use a higher degree of bendiness for lower complexity surfaces, which will result in a more robust notion. The details are discussed in Section 7.)

Consider the lift $\tilde{\chi}$ of χ to the universal cover $\tilde{S}$, of S. We can equally well refer to the bending of edges in $\tilde{\chi}$. We know that some edge, ϵ, of $\tilde{\chi}$ is bent along Φ. (This entails a certain normalisation — we choose an arbitrary vertex $\tilde{x}_0$ of $\tilde{\chi}$, and equivariantly associate to this vertex the marking a, and associate to other vertices $\mathcal{M}$-images of a equivariantly.) Let $\epsilon', \epsilon'' \in E(\tilde{\chi})$ be the other edges of a triangle of $\tilde{\chi}$ incident on ϵ. Applying the coarse triangle inequality, we see that at least one of ϵ' or ϵ'', say ϵ', is bent along Φ (though perhaps not quite as bent). We now proceed to the triangle on the other side of ϵ' and continue. In this way, we construct a path, ζ, in $\tilde{S}$, such that each edge it crosses is bent along Φ. We will take it to be sufficiently long depending on $\kappa(\Sigma)$ as required below. We now go back down to S. The projected path, $\hat{\zeta}$ must cross some edge, e, of χ many times. Each time it crosses, it means that e is bent along some $\mathcal{M}$-image of Φ. (Between two such crossings there are segments of e and of $\hat{\zeta}$ which together give rise to a monodromy in $\mathcal{M}$.) But all these subsurfaces have the same complexity, and are hence bounded in number. This means we must return to the same subsurface twice, and so we can close up our path to give us a closed path, π, in S (not necessarily simple). (Note that we needed the bound on the number of possibilities for the subsurface since we may lose a certain amount of bendiness each time we cross a triangle, so our initial supply of bendiness may eventually be exhausted.) If we follow a boundary curve of Φ in Σ around π in S (possibly several times) we get a periodic curve (or a torus in the associated surface bundle). If π is essential in S, then our homomorphism, ϕ, cannot be purely loxodromic, giving a contradiction. In Section 7, we will present this argument more formally in terms of dual graphs to χ, rather than paths. (We will argue by assuming the bundle to be atoroidal, and inductively place bounds on projected distances as a function of complexity.)

We need to explain why π can be assumed non-trivial. This will entail a careful choice of basepoint, a. This will be either a complete multicurve (pants decomposition) of or marking (a complete multicurve together with marking curves), see the definitions in Section 4. (At present, the basepoint is only determined up to bounded distance in $\mathcal{G}$.) If we were careless, we might expect π just to run around the vertex, x_0, of χ. Returning to the analogy with the case of maps from S into a hyperbolic manifold N, there we chose an arbitrary edge, $\epsilon_0 \in E(\tilde{\chi})$, and mapped it to the closed geodesic in N in the given free homotopy class. In the new situation the "homotopy class" corresponds to a monodromy, $\theta_0 = \phi(g(\epsilon_0)) \in \mathcal{M}$, giving rise to a 3-manifold $M_0 = M_{\theta_0}$ fibring over the circle. As noted in Theorem 1.1, this admits a hyperbolic structure (the analogue of the closed geodesic in N) and we shall use this structure to choose a. We will require (among other things) that a should have bounded length when realised in M_0. It follows that a is a bounded distance from an "axis" of θ_0 acting on $\mathcal{G}$ (see Lemma 3.5), and so, after applying a suitable power of θ_0, and using the hyperbolicity of $\mathcal{G}$, we can assume that $d(\alpha, a)$ is bounded, where $\alpha \in V(\mathcal{G}(\Sigma))$ was our original attempt at a basepoint (as given by [Ba]). Thus, a will serve just as well. In other words, $d(a, \theta a)$ is bounded for all $\theta \in A(\chi)$. (We can now forget about the original α.)

Here is a more precise recipe to find a marking a. Suppose for the moment that we can find a uniformly lipschitz map, $f : (\Sigma, \sigma) \longrightarrow M_0$, in the fibre class, with $f(\Sigma)$ not meeting any Margulis tube in M_0, and where σ is a hyperbolic structure on Σ. (Indeed any metric of curvature at most -1 will do.) The injectivity radius of (Σ, σ) is bounded below, and so its diameter is bounded above. Thus, it is easy to construct a marking, a, in Σ, of bounded length. Its image in M_0 is also of bounded length.

To apply this, we need to relate projected distances to the geometry of M_0, or more precisely, to the infinite cyclic cover, $M \cong \Sigma \times \mathbf{R}$, of M_0. Note that f lifts to a bi-infinite sequence of maps $f_i : \Sigma \longrightarrow M$, so that $f_i \equiv f_0 \circ \theta_0^i$. We regard f_0 as determining a preferred class of homotopy equivalence from Σ to M. Suppose $i < j$, and $p(\Phi; \theta_0^i a, \theta_0^j a)$ is large. If Φ is an annulus, then there is a Margulis tube in M, in the homotopy class of Φ, which is homologically trapped between $f_i(\Sigma)$ and $f_j(\Sigma)$ (i.e. any homotopy between them maps with non-zero degree to this tube). If Φ is more complex, then there is a "band" in the homotopy class trapped between $f_i(\Sigma)$ and $f_j(\Sigma)$. Here, a band is a topological copy of $\Phi \times [0, 1]$, with $\partial\Phi \times [0, 1]$ lying in boundaries of Margulis tubes, and which is "long" in the direction of the second co-ordinate. Note that there can only be one tube or band in M in any given homotopy class. These notions arise out of the approach to the Ending Lamination Conjecture in [Mi2,BrCM] and in [Bow4], and are discussed in Section 6, see in particular, Proposition 6.10.

Now suppose that our path, π, is trivial in S. It lifts to a closed path, $\tilde{\pi}$, in $\tilde{S}$. Now there must be some component, $\tilde{\epsilon}_0$, of the lift of ϵ_0 in $\tilde{S}$, such that $\tilde{\pi}$ crosses $\tilde{\epsilon}_0$ in at least two distinct edges (cf. Lemma 7.1). By construction, each of these edges is bent along Φ. From the definition of "bent" this means that there are some $i < j$, so that $p(\Phi; \theta_0^i a, \theta_0^{i+1} a)$ and $p(\Phi; \theta_0^j a, \theta_0^{j+1} a)$ are both large. (As stated, this entails a certain choice of normalisation so that $\tilde{x}_0 \in \tilde{\epsilon}_0$, though the argument would work just as well with any $\mathcal{M}$-image of Φ.) This in turn implies that there are two tubes or two bands both homotopic to Φ, and respectively trapped between $f_i(\Sigma)$ and $f_{i+1}(\Sigma)$ and between $f_j(\Sigma)$ and $f_{j+1}(\Sigma)$. Thus, $f_j(\Sigma)$ must homologically separate them. In particular, they must be different, contradicting the fact that there is only one tube or band in a given homotopy class. There is a subtle issue here about the "order" in which the $f_i(\Sigma)$ occur in M. To make sense of this, we will use suitable homotopies to embedded surfaces in M (cf. Lemma 6.9).

To recap, once we can find such a map, $f : (\Sigma, \sigma) \longrightarrow M_0$, and make an appropriate choice of basepoint, the above argument shows that the lift, $\tilde{\epsilon}_0$, can contain at most one edge that is bent along any given subsurface. This shows that there can be no closed path in $\tilde{S}$ passing only through bent edges. The path, π, that we construct is therefore essential, and so gives a contradiction to the purely pseudoanosov hypothesis (cf. Lemma 7.4).

We now return to the issue of the existence of a uniformly lipschitz map $f : (\Sigma, \sigma) \longrightarrow M_0$. For this, we will use minimal surfaces in M_0 [MeSY,SY,FHS] (see Lemma 6.12). If the injectivity radius of M_0 is uniformly bounded below (i.e. there are no Margulis tubes) we can construct f by taking a minimal surface in the fibre class. In this case, f, will be an embedding. If each Margulis tube has bounded volume (which is equivalent to there being no bands), then we can push f off the Margulis tubes, while maintaining control of the lipschitz constant. (Here f remains an embedding.) If there are bands in M_0, we can push f off all the bands, (making use of Lemma 6.6) but it might be impossible to push it off all the Margulis tubes. We therefore have to live with the fact that there may be very short curves in (Σ, σ), so we can no longer find a marking of bounded length. We can, however, find a complete multicurve (i.e. a pants decomposition) of bounded length, and this will serve as our basepoint in this case. Note that the existence of a band means that ϵ_0 is bent along a non-annular surface, and so the maximal complexity surface, Φ, cannot be an annulus. In this case, a complete multicurve is sufficient to make sense of the above homological separation argument, again giving a contradiction.

The requirements of the choice of basepoint needed for the above argument, can be formulated as a purely combinatorial statement, without explicit reference to 3-manifolds, namely Lemma 6.1. The purpose of Section 6, and the introduction of hyperbolic 3-manifolds is just to prove this statement.

In principle, one might hope for a purely combinatorial proof based, for example, on the combinatorial models of Masur, Minsky et al., as described in Section 5. The difficulty arises from the fact that the model is not equivariant, and does not descend to M_0 —

a feature that was essential in making sense of the ordering and separation properties of surfaces. In a hypebolic 3-manifold we have minimal surfaces at our disposal to use as a starting point.

To make all this work, we need to take care of the quantification, which will involve us in a series of constants bounding lengths, distances, injectivity radii, lipschitz constants etc. Ultimately, these will all depend only on genus(S) and genus(Σ). The logic is laid out in Section 7, where it is presented as an induction on the complexity of subsurfaces.

3. The curve graph.

In this section, for the purposes of giving inductive arguments later, we allow Σ to be a compact surface with $p \geq 0$ boundary components. We define the *complexity* of Σ as $\kappa(\Sigma) = 3\,\mathrm{genus}(\Sigma) + p - 3$. We write $X(\Sigma)$ for the set of homotopy classes of non-trivial non-peripheral simple closed curves in Σ. Given $\alpha, \beta \in X(\Sigma)$, write $\iota(\alpha, \beta)$ for their geometric intersection number (i.e. the minimal number of intersection points).

A certain number of "exceptional" cases call for special attention. Note that if $\kappa(\Sigma) = -1$ (the annulus) or $\kappa(\Sigma) = 0$ (the three-holed sphere, denoted "3HS"), then $X(\Sigma) = \emptyset$. If $\kappa(\Sigma) = 1$, then Σ is either a four-holed sphere, "4HS", or a one-holed torus, "1HT". (We won't need to deal with the disc, sphere, or closed torus here.)

Suppose that $\kappa(\Sigma) \geq 2$ ("the non-exceptional" case). We define the *curve graph*, $\mathcal{G} = \mathcal{G}(\Sigma)$ by taking the vertex set, $V(\mathcal{G}) = X(\Sigma)$ and deeming $\alpha, \beta \in V(\mathcal{G})$ to be adjacent (or equal) if $\iota(\alpha, \beta) = 0$. (This is the 1-skeleton of the curve complex introduced in [Ha].) We write $d = d_{\mathcal{G}(\Sigma)}$ for the combinatorial metric on $\mathcal{G}(\Sigma)$. One can show that $d(\alpha, \beta)$ is bounded above in terms of $\iota(\alpha, \beta)$. This is implicit in work of Lickorish, and one can find an explicit bound for example, in [MaM1]. Clearly, $\mathcal{M}$ acts on $\mathcal{G}$ with finite quotient. The following is a key result:

Theorem 3.1 : [MaM1] $\mathcal{G}(\Sigma)$ *is Gromov hyperbolic.* ♢

The following notion makes sense for any action of a group, Γ, on a graph $\mathcal{G}$. Given $g \in \Gamma$, we write $||g||$ for its stable length. That is $||g|| = \lim_{n \to \infty} \frac{1}{n} d(x, g^n x)$ for some, hence any, $x \in \mathcal{G}$. We write $\mathrm{inj}(\mathcal{G}, \Gamma) = \inf\{||g||\}$ as g ranges over all infinite order elements of Γ. (In the torsion-free case, this is analagous to the injectivity radius of a negatively curved manifold.)

Returning to the case of $\mathcal{M}$ acting on the curve graph $\mathcal{G}(\Sigma)$, we have:

Theorem 3.2 : [MaM1] *There is some $\eta > 0$, depending only on* genus(Σ) *such that if $g \in \mathcal{M}$ is pseudoanosov, then $||g|| \geq \eta$.*

(This is also a consequence of the acylindricity of the action of $\mathcal{M}$ on $\mathcal{G}(\Sigma)$. In fact, stable lengths are uniformly rational [Bow3].)

The following is a general statement:

Theorem 3.3 : *Suppose that S is a closed surface and that $\Gamma = \pi_1(S)$ acts on a k-hyperbolic graph, $\mathcal{G}$, with $\mathrm{inj}(\mathcal{G},\Gamma) \geq \eta > 0$. Then there is a standard generating set, $g_1, \ldots, g_n$, for Γ, and some $x \in V(\mathcal{G})$, such that for all i, $d(x, g_i x) \leq s$, where s depends only on $\kappa(\Sigma)$, k and η.*

There are different methods available for proving this result. In particular, it was observed by Jason Manning that it follows from an elegant argument given in the appendix to the thesis of Josh Barnard. This has since been written up in more detail as [Ba]. (One can also generalise to non-orientable surfaces, as discussed in [Bow5].)

The idea is to consider a "carrying graph", $\Omega \subseteq S$, such that $S \setminus \Omega$ is a disc. Let $\tilde{\Omega}$ be its lift to the universal cover $\tilde{S}$. Given a Γ-equivariant map from the vertex set of $\tilde{\Omega}$ to $\mathcal{G}$, we can define a "length" of Ω, by taking a sum of distances between f-images of pairs vertices, corresponding to an orbit transversal of edges. We choose a pair (Ω, f) whose length is (nearly) minimal. One then shows that its length is bounded above in terms of genus(Σ), k and $\mathrm{inj}(\mathcal{G},\Gamma)$ — if there were a very long edge then a shortcutting operation would reduce the length. For this to work in general, one needs to know that two bi-infinite paths in $\mathcal{G}$ which are translated a bounded distance by some element, $g \in \Gamma$, are a uniformly bounded distance apart. This is where one uses the lower bound on $||g||$. We also need to make the additional observation that the generating set arising in this argument can be assumed to be standard. Note that the graph Ω remains embedded under the shortcutting operations, and so there are only finitely many combinatorial possibilities, up to the action of $\mathcal{M}$. We can therefore find a standard generating set represented by paths in Ω that cross each edge of Ω a bounded number of times.

There is only one standard generating set up to the action of $\mathcal{M}$. (In fact, it would be sufficient to know that there are only finitely many — see also the discussion of "complexity" of generating sets in Section 8.)

To apply Theorem 3.3, suppose that Σ and S are closed surfaces and that $\Sigma \longrightarrow \Pi \longrightarrow S$ is an atoroidal bundle. Then the monodromy homomorphism from Γ to $\mathcal{M}$ is injective, and so $\mathrm{inj}(\mathcal{G}(\Sigma),\Gamma) \geq \eta$, where η is the constant of Lemma 3.2. Thus:

Proposition 3.4 : *There is some $s = s(\kappa(\Sigma), \kappa(S))$ such that if $\Sigma \longrightarrow \Pi \longrightarrow S$ is atoroidal, then there is a standard set of generators, $g_1, \ldots, g_n$ for $\Gamma = \pi_1(S)$, and some $\alpha \in V(\mathcal{G})$ such that for all i, $d(\alpha, g_i\alpha) \leq s$.*

We note that one can generalise Proposition 3.4 to the case of a purely pseudoanosov homomorphism of a finitely presented group, Γ to $\mathcal{M}$. Here a "standard" generating set is replaced by one of bounded "complexity" as defined in Section 8. Independent proofs of this are given in [DeF] and [Bow5], both using results from [Bow3]. Either of these results is sufficient to get from Theorem 8.1 to the statement that there are only finitely many purely pseudoanosov isomorphic copies of Γ in $\mathcal{M}$ up to conjugacy. As mentioned in the introduction, an independent proof of this statement has been announced by Groves.

This is a step towards verifying the hypotheses of Lemma 1.6, though we need to do quite a bit more, since there are infinitely many elements of $\mathcal{M}$ moving any given curve a bounded distance.

We will also need to consider the curve graph of an exceptional surface where $\kappa(\Sigma) = 1$. Here it is defined differently. Again, $V(\mathcal{G}) = X(\Sigma)$, but α, β are deemed adjacent if $\iota(\alpha, \beta)$ is minimal (i.e. 1 for a 1HT and 2 for a 4HS). Thus, $\mathcal{G}(\Sigma)$ is a Farey graph.

The curve graph is central to the proof of the Ending Lamination Conjecture in [Mi,BrCM] (see also [Bow4]). Some of this will be discussed in Section 6. For the moment, we just note the following consequence of that work.

Suppose that M is a doubly degenerate 3-manifold homeomorphic to $\Sigma \times \mathbf{R}$ and without cusps (for example the cyclic cover of a bundle over a circle). We can realise any $\gamma \in X(\Sigma)$ as a closed geodesic in M, and write $l_M(\gamma)$ for its length.

It is a consequence of work in [Mi] or in [Bow4] that there is a bi-infinite geodesic, ζ, in $\mathcal{G}(\Sigma)$ such that if $\gamma \in X(\Sigma)$, then the distance in $\mathcal{G}(\Sigma)$ from γ to ζ is bounded above by some function of $l_M(\gamma)$. The geodesic ζ, is well defined up to bounded distance, depending only on $\kappa(\Sigma)$. We apply this as follows.

Suppose that ϕ is pseudoanosov.

Definition : An *axis* of ϕ in $\mathcal{G}(\Sigma)$, is a bi-inifinite geodesic, ζ, such that $\phi\zeta$ is a bounded distance from ζ.

Any two such axes (in particular ζ are $\phi\zeta$) are a uniformly bounded distance apart, depending only on $\kappa(\Sigma)$. The existence of an axis for any peudonanosov follows from [MaM1,MaM2].

Suppose that M is the cyclic cover of M_ϕ for some pseudoanosov $\phi \in \mathcal{M}$.

Lemma 3.5 : *Suppose that $\phi \in \mathcal{M}$ and that M_ϕ and M are as above. Suppose $\gamma \in X(\Sigma)$ with $l_M(\gamma) \leq l$ for some $l \geq 0$, then γ lies within a distance q of an axis of ϕ on $\mathcal{G}(\Sigma)$, where q depends only on $\kappa(\Sigma)$ and l.* $\diamondsuit$

4. Subsurface projections.

Subsurface projections were defined in [MaM2]. They play a crucial role in the approach of [Mi,BrCM] to the Ending Lamination Conjecture. Their use in this regard was circumvented in [Bow4], but they can also be tied in with the presentation there, as we discuss in Section 5.

Let Σ be a compact connected surface. For the moment, we can allow Σ to have boundary. Let $\Phi \subseteq \Sigma$ be a closed connected subsurface. We shall always assume its inclusion to be π_1-injective, i.e. $\Sigma \setminus \Phi$ has no disc or boundary-parallel annular components (though we need to allow for other annular components). If the inclusion is strict, then $\kappa(\Phi) < \kappa(\Sigma)$. We write $\mathcal{F} = \mathcal{F}(\Sigma)$ for the set of such Φ up to homotopy which satisfy $\kappa(\Phi) \geq 1$. We can identify the set of annuli up to homotopy with $X(\Sigma)$.

By an *arc* in Σ, we mean an embedded arc, α, with $\partial\alpha = \alpha \cap \partial\Sigma$. We consider this defined up to homotopy, sliding $\partial\alpha$ in $\partial\Sigma$. We also insist that it is homotopically non-trivial.

Let $X_A(\Sigma)$ be the set of arcs and curves in Σ, and let $\mathcal{G}_A$ be the arc-and-curve graph, that is, $V(\mathcal{G}_A) = X_A(\Sigma)$, and two elements of $X_A(\Sigma)$ are deemed adjacent or equal if they can be homotoped to be disjoint. It's not hard to see that the inclusion $X(\Sigma) \hookrightarrow X_A(\Sigma)$ induces a quasi-isometry from $\mathcal{G}(\Sigma)$ to $\mathcal{G}_A(\Sigma)$. (This holds even if $\kappa(\Sigma) = 1$.) One way to go back from an arc α in Σ is to take a boundary component of a regular neighbourhood of α union the boundary component, or components, it meets. (This is well defined up to distance 1.)

Suppose now that $\Phi \in \mathcal{F}$. We write $X(\Sigma) = X_I(\Sigma,\Phi) \sqcup X_D(\Sigma,\Phi) \sqcup X_T(\Sigma,\Phi)$, where $\alpha \in X(\Sigma)$ lies in $X_I(\Sigma,\Phi)$ if it is a curve in $X(\Phi)$, it lies in $X_D(\Sigma,\Phi)$ if it can be homotoped to be disjoint from Φ, and it lies in $X_T(\Sigma,\Phi)$ if it crosses some boundary curve of Φ. We write $X(\Sigma,\Phi) = X_I(\Sigma,\Phi) \cup X_T(\Sigma,\Phi)$. We define the *subsurface projection*, $\theta_\Phi : X(\Sigma,\Phi) \longrightarrow X(\Phi)$ as follows. Suppose that $\alpha \in X(\Sigma,\Phi)$. If $\alpha \in X_I(\Sigma,\Phi)$, set $\theta_\Phi\alpha = \alpha$. If $\alpha \in X_T(\Sigma,\Phi)$, let γ be any component of $\phi \cap \alpha$, and let $\theta_\Phi\alpha$ be the boundary component of a small neighbourhood of γ union the incident components of $\partial\Phi$. From the above discussion, this is well defined in $\mathcal{G}(\Phi)$ up to bounded distance (independent of $\kappa(\Sigma)$). Moreover, if $\alpha, \beta \in X(\Sigma,\Phi)$, then $\iota(\theta_\Phi\alpha, \theta_\Phi\beta)$ is bounded above in terms of $\iota(\alpha,\beta)$.

Given a subset $b \subseteq X$ we will write $\iota(b) = \max\{\iota(\alpha,\beta) \mid \alpha,\beta \in b\}$.

Definition : A *multicurve*, $a \subseteq X(\Sigma)$, is a (finite) subset of pairwise adjacent curves (i.e. $\iota(a) = 0$).

Such a set can be realised so that they are all simultaneously disjoint. We shall frequently think of a multicurve as a 1-manifold in Σ, defined up to homotopy, with each component essential. In this paper, we allow a multicurve to be empty.

Definition : We say that a multicurve, a, is *complete* if it is maximal, i.e. each component of $\Sigma \setminus a$ is a 3HS.

Suppose that $a \subseteq X(\Sigma)$ is a complete multicurve, and that $\alpha \in a$. Let $\kappa(\Phi)$ be the component of $\Sigma \setminus (a \setminus \alpha)$ containing α. Thus $\kappa(\Phi) = 1$. A *marking curve* of α is a curve β in Φ, adjacent to α in the Farey graph $\mathcal{G}(\Phi)$. A *complete marking* of a is a subset, $b \subseteq X(\Sigma)$, consisting of a, together with a marking curve for each element of a. Note that $\iota(b) \leq 2$, and that b fills Σ (i.e. each component of $\Sigma \setminus b$ is a disc.

Definition : A *marking* of Σ is the complete marking of some multicurve.

One can recover the multicurve uniquely from the marking. Note that if a is a marking, then $\iota(a) \leq 2$.

Any mapping class that preserves a complete multicurve must be a product of Dehn twists about these curves. Thus, any mapping class that preserves a complete marking is trivial.

If $a \cap X(\Sigma,\Phi) \neq \emptyset$, then we set $\theta_\Phi(a) = \theta_\Phi(\alpha)$ for some $\alpha \in a \cap X(\Sigma,\Phi)$. This is well defined up to bounded intersection, hence bounded distance in $X(\Sigma)$, depending only on $\kappa(\Sigma)$ and $\iota(a)$. (In practice we will only use this when $\iota(a) \leq 2$.)

To be more formal, we could define $\theta_\Phi \alpha$ or $\theta_\Phi a$ to be the set of all possible choices, thereby giving a canonical subset of bounded diameter. We make choices just for convenience of notation and exposition. We can assume these choices to be equivariant where convenient (going back to the more formal definition if necessary).

Definition : Given $\alpha, \beta \in X(\Sigma, \Phi)$, we will write $p(\Phi; \alpha, \beta) = d_{\mathcal{G}(\Phi)}(\theta_\Phi \alpha, \theta_\Phi \beta)$ for the *projected distance* between α and β in Φ. We similarly define $p(\Phi; a, b)$ if $a \cap X(\Sigma, \Phi) \neq \emptyset$ and $b \cap X(\Sigma, \Phi) \neq \emptyset$.

The above discussion applies only when $\kappa(\Phi) \geq 1$. We now move on to consider the case of annuli, or equivalently, elements of $X(\Sigma)$.

In what follows, it will be convenient to write $x \sim_k y$ to mean that $|x - y| \leq k$, for $x, y, k \in \mathbf{R}$.

We write τ_α for the positive Dehn twist about α. (This depends only on the orientation of the underlying surface.)

Suppose that R is an annulus, with core curve δ, and boundary components, δ_1 and δ_2. Suppose that ϵ, ζ are arcs connecting δ_1 to δ_2. There is a signed intersection number, $\hat{\iota}(\epsilon, \zeta)$, that counts the number of times ϵ crosses ζ in the positive sense. We will write $\iota(\epsilon, \zeta) = |\hat{\iota}(\epsilon, \zeta)|$ for the usual intersection number. (That is, the minimal number of intersections between representatives of ϵ and ζ in their respective homotopy classes relative to their endpoints.) If ϵ', ζ' are other such arcs, then it's not hard to see that $|\hat{\iota}(\epsilon, \zeta) - \hat{\iota}(\epsilon', \zeta')| \leq \iota(\epsilon, \epsilon') + \iota(\zeta, \zeta') + 2$. In particular, if ϵ, ϵ' are disjoint, and ζ, ζ' are disjoint, then $\hat{\iota}(\epsilon, \zeta) \sim_2 \hat{\iota}(\epsilon', \zeta')$. We also note that $\hat{\iota}(\tau_\delta^m \epsilon, \tau_\delta^n \zeta) = \hat{\iota}(\epsilon, \zeta) + n - m$. This, of course, depends on certain conventions of orientation. In particular, we need to assume that ϵ and ζ cross R in the same direction, but it doesn't matter which.

Suppose now that $\kappa(\Sigma) \geq 1$ and that $\delta \in X(\Sigma)$. Let $X(\Sigma, \delta) = \{\alpha \in X(\Sigma) \mid \iota(\delta, \alpha) \neq 0\}$. The annular cover of Σ corresponding to δ naturally compactifies to a closed annulus, $R(\delta) = R(\Sigma, \delta)$, with core curve δ.

If $\alpha \in X(\Sigma) \setminus \{\delta\}$, then α lifts to an arc whose closure in $R(\delta)$ is a closed arc, $\tilde{\alpha}$. This is well defined up to homotopy relative to its endpoints. If $\alpha \in X(\Sigma, \delta)$, we can choose this lift so that it crosses δ. Any two such choices will be disjoint.

Suppose now that $\alpha, \beta \in X(\Sigma, \delta)$. Then, up to $\sim_2$, $\hat{\iota}(\tilde{\alpha}, \tilde{\beta})$ is independent of the choice of lifts.

Definition : We set $\vec{p}(\delta; \alpha, \beta) = \hat{\iota}(\tilde{\alpha}, \tilde{\beta}) \in \mathbf{Z}$, for the *signed projection distance* and write $p(\delta; \alpha, \beta) = |p(\delta; \alpha, \beta)|$ for the unsigned projection distance.

We see easily that $\vec{p}(\delta; \alpha, \alpha) = 0$, and that for all $\alpha, \beta, \gamma \in X(\Sigma, \delta)$, $|\vec{p}(\delta; \alpha, \beta) + \vec{p}(\delta; \beta, \gamma) + \vec{p}(\delta; \gamma, \alpha)| \leq 3$. In other words, $\vec{p}(\delta; ., .)$ is a coarse 2-cycle. It follows that $p(\delta; ., .)$ satisfies a coarse triangle inequality: $p(\delta; \alpha, \beta) \leq p(\delta; \alpha, \gamma) + p(\delta; \gamma, \beta) + 3$.

Lemma 4.1 :
(1) $\vec{p}(\delta; \tau_\delta^m \alpha, \tau_\delta^n \beta) \sim_4 \vec{p}(\delta; \alpha, \beta) + n - m$.
(2) If $\epsilon \in X(\Sigma)$ *is disjoint from* δ*, then* $\vec{p}(\delta; \tau_\epsilon^m \alpha, \tau_\epsilon^n \beta) \sim_4 \vec{p}(\delta; \alpha, \beta)$.

Proof : (Some similar arguments can be found in Section 6 of [Mi1].) We begin with the second statement.

(2) Note that ϵ lifts to a set of disjoint arcs in $R(\delta)$. Each such arc bounds a disc in $R(\delta)$ disjoint from δ. The set of outermost such discs are disjoint. Now either α does not meet ϵ and is unaffected by Dehn twist, and $\tilde{\alpha}$ does not meet any of these outermost discs. Or else, α crosses ϵ and $\tilde{\alpha}$ enters a outermost disc on either side of δ. Once it enters, it never leaves. Now twisting along ϵ will have the effect of sliding the intersection of α with the boundaries these discs.

A similar discussion applies to β. But now $\tilde{\alpha}$ and $\tilde{\beta}$ can intersect at most once inside any of these outermost discs. Thus twisting along ϵ can change the total number of intersections of α and β by at most 4. (It can create or destroy at most two intersections in each such disc.) The statement (2) now follows.

(1) Now δ lifts to the core curve, δ, together with a disjoint collection of arcs. Thus a similar argument applies. We also need the observation about Dehn twists along δ in $R(\delta)$. ♢

We note, for future reference, that the following is implicit in the above proof. Suppose that $\delta \in X(\Phi)$ where $\Phi \in \mathcal{F}$. If $\alpha, \beta \in X(\Sigma, \delta)$, then (up to an additive constant) $\vec{p}(\delta; \alpha, \beta)$ depends only on how α and β intersect Φ (and hence, up to distance 2) only on their projections to $X(\Phi)$).

The following "projection bounds" theorem is central. It is proven in [MaM2], except that we have modified the statement by replacing "geodesic" with "quasigeodesic":

Theorem 4.2 : *Suppose that $\Phi \in X \cup \mathcal{F}$, and that $(\alpha_i)_i$ is a quasigeodesic in $X(\Sigma)$. Suppose that $\alpha_i \in X(\Sigma, \Phi)$ for all i. Then for all j, k, $p(\Phi; \alpha_j, \alpha_k) \leq v$, where v depends only on $\kappa(\Sigma)$ and the constants of quasigeodesicity.*

Proof : We can assume that α_j and α_k are the terminal points of $(\alpha_i)_i$. The statement, where $(\alpha_i)_i$ is assumed to be geodesic is proven in [MaM2]. For the above statement, note that, by hyperbolicity of $\mathcal{G}(\Sigma)$, $(\alpha_i)_i$ remains a bounded distance from a geodesic, π, in $\mathcal{G}(\Sigma)$ from α_j to α_k. Moreover, $X(\Sigma) \setminus X(\Sigma, \Phi)$ has diameter 1 in $\mathcal{G}(\Sigma)$. From this, we claim that α_i and α_j can be connected by a path consisting of at most five segments, alternately geodesic or bounded length, all lying in $X(\Sigma, \Phi)$. If $\pi \subseteq X(\Sigma, \Phi)$, then we can just use a segment of π, suppose that there is some $\beta \in \pi \setminus X(\Sigma, \Phi)$. If α_j and α_k are both a long way from β, then we follow π from α_k until we get close to to β, then cross to $(\alpha_i)_i$ by a path of bounded length not meeting β, and follow it, a bounded distance, until we are the other side of β and cross back to π, and follow it to α_k. If one or both of α_j or α_k are close to β, then we follow $(\alpha_i)_i$ near these points. This justifies the claim, and the result now follows from the statement for geodesics, together with the observation that $p(\Phi; \beta, \gamma)$ is bounded whenever $\beta, \gamma \in X(\Sigma, \Phi)$ are adjacent. ♢

The following is a simple consequence of Theorem 6.12 of [MaM2]. (It is a also a consequence of the results of Section 5 here.)

Proposition 4.3 : *Suppose $a, b \subseteq X(\Sigma)$ are complete markings. Suppose that, for some q, we have $p(\Phi; a, b) \leq q$ for all $\Phi \in X \cup \mathcal{F}$. Then $\iota(a \cup b)$ is bounded above in terms of $\kappa(\Sigma)$ and q.*

Proof : In [MaM2] the authors define a "distance" between two markings. They show that there is a fixed constant $q_0 = q_0(\kappa(\Sigma))$ such that if $q \geq q_0$ then the distance between any two markings a and b is bounded above by $C_0 \sum_{\Phi \in \mathcal{F}(a,b,q)} p(\Phi; a, b) + C_1$. Here, $\mathcal{F}(a, b, q) = \{\Phi \in X \cup \mathcal{F} \mid p(\Phi; a, b) > q\}$ and C_0, C_1 are constants depending only on $\kappa(\Sigma)$ and q. They also show that the intersection number, $\iota(a, b)$, is bounded above in terms of the distance. Under our hypotheses, we will have $\mathcal{F}(a, b, q) = \emptyset$, and so we get bound just in terms of C_1, and so in terms of $\kappa(\Sigma)$ and q as required. $\diamondsuit$

To apply this, we note the following simple observation:

Lemma 4.4 : *Suppose $a \subseteq X(\Sigma)$ is a complete marking of Σ. Then for any $r \geq 0$, the set of $\phi \in \mathcal{M}$ with $\iota(a \cup \phi a) \leq r$ is finite.* $\diamondsuit$

Proof : Since a fills Σ, the bound on $\iota(a \cup \phi a)$ bounds the number of possibilities for each component of ϕa. As observed earlier, any mapping class preserving each component of a complete marking is trivial. $\diamondsuit$

Corollary 4.5 : *Suppose that $a \subseteq X(\Sigma)$ is a complete marking. Given any $q \in \mathbf{N}$, the set $\{\phi \in \mathcal{M} \mid (\forall \Phi \in X \cup \mathcal{F})(p(\Phi; a, \phi a) \leq q)\}$ is finite.* $\diamondsuit$

The following will be proven in Section 5. (One could probably also use the technology of [MaM2].) Here a, b will be taken to be either complete multicurves or complete markings. (In practice, the essential point is that $\iota(a)$ and $\iota(b)$ are bounded.)

Proposition 4.6 : *Given $\kappa = \kappa(\Sigma)$, there is some $v_0(\kappa) \geq 0$ such that for all $v \geq v_0(\kappa)$, there is some $n \geq 0$ such that the following holds. Suppose that $a, b \subseteq X(\Sigma)$ are complete multicurves or complete markings. Suppose that $k \in \{-1, 1, 2, 3, \ldots, \kappa\}$, and that for all $F \in \mathcal{F}$ with $\kappa(F) > k$, we have $p(F; a, b) \leq v$. Then there are at most n subsurfaces $\Phi \in X \cup \mathcal{F}$ with $\kappa(\Phi) = k$ and with $p(\Phi; a, b) \geq v_0$.*

The interpretation where $k = -1$ is that we are assuming $p(F; a, b) \leq v$ for all $F \in \mathcal{F}$, and conclude that there are at most n curves $\delta \in X(\Sigma)$ with $p(\delta; a, b) \geq v_0$.

5. Annulus systems.

In this section, we tie in the discussion of subsurface projections with "annulus systems" as described in [Bow4]. These are a less sophisticated version of "hierarchies" as defined in [MaM2], but will be sufficient for our purposes.

For convenience, we will assume that Σ is a closed surface of given genus (at least 2). We also deal here in finite annulus systems, though the essential points can be readily adapted from the locally finite case in [Bow4]. We begin by recalling some definitions, and refer to the above for elaboration.

Let $I = [\partial_- I, \partial_+ I] \subseteq \mathbf{R}$ be a non-trivial compact interval, and let $\Psi = \Sigma \times I$. The first coordinate will be termed "horizontal" and the second "vertical". We write $\partial_\pm \Psi = \Sigma \times \partial_\pm I$, and let $\partial_H \Psi = \Sigma \times \partial I = \partial_- \Psi \sqcup \partial_+ \Psi$. We write $\pi_\Sigma : \Psi \longrightarrow \Sigma$ and $\pi_V : \Psi \longrightarrow I$ for the projections.

It will be convenient to consider all closed curves as realised as geodesics with respect to some fixed hyperbolic structure on Σ. This is really for ease of exposition — all the constructions are essentially combinatorial.

Let $\mathcal{S}$ be the set of homotopy classes of 3HS's in Σ. Given $\Phi \in \mathcal{S} \cup \mathcal{F}$, we can realise Φ so that all its boundary components are geodesic. (This may entail identifying certain pairs of boundary components.) By a *strip*, B, in Ψ, we mean a subset of the form $B = \Psi \times J$, where $\Phi \in \mathcal{S} \cup \mathcal{F}$ (realised as above) and $J = [\partial_- J, \partial_+ J] \subseteq I$ is a non-trivial closed interval. We write $\partial_\pm B = \Phi \times \partial_\pm J$, $\partial_H B = \Phi \times \partial J = \partial_- B \sqcup \partial_+ B$ and $\partial_V B = \partial \Phi \times J$. Thus, $\partial B = \partial_H B \cup \partial_V B$. We write $\kappa(B) = \kappa(\Phi)$ for its complexity.

A *vertical annulus* in Ψ is a subset of the form $\Omega = \gamma \times J$, where $\gamma \in X(\Sigma)$, and J is a non-trivial compact interval. We write $\partial_\pm \Omega = \alpha \times \partial_\pm J$ and $\partial_H \Omega = \alpha \times \partial J$. Note that if B is a strip, the $\partial_V B$ is a union of vertical annuli.

Definition : An *annulus system*, $\mathcal{W}$, in Ω is a finite set of disjoint vertical annuli.

We write $W = \bigcup \mathcal{W}$. We often refer to W itself as an "annulus system". We normally assume (without specifying) that W is *vertically generic*, that is, if $\Omega, \Omega' \in \mathcal{W}$ are distinct, then $\pi_V(\partial_H \Omega) \cap \pi_V(\partial_H \Omega') \subseteq \partial I$. This can always be arranged after a small adjustment of W.

We write $X_\pm(W) = \pi_\Sigma(W \cap \partial_\pm \Psi)$. We write $X(W) = \{\pi_\Sigma \Omega \mid \Omega \in \mathcal{W}\}$.

Definition : A *band* in Ω (with respect to W) is a strip, $B \subseteq \Psi$ such that $\partial_V B \subseteq W$, and such that no annulus of W crosses B (i.e. intersects both $\partial_+ B$ and $\partial_- B$).
A band is *full* if $\pi_\Sigma B = \Sigma$. Otherwise it is *proper*.

Note that we can view $W_B = W \cap B$ as an annulus system in B (though we need to allow for degenerate annuli in $\partial_H B$). This forms the basis of certain inductive arguments.

Given $u \in I$, write $\Sigma_u = \Sigma \times \{u\}$, and let $\omega(u) = \pi_\Sigma(W \cap \Sigma_u)$. This is a (possibly empty) multicurve in Σ. By definition, $\omega(\partial_\pm I) = X_\pm(W)$.

Definition : We say that W is *complete* if the multicurves $X_\pm(W)$ are complete, and if, for all $u \in I$, the multicurve is either complete, or complete minus at most one curve (i.e. we are allowing for at most one 1HT or 4HS component).

Let us assume that W is complete. Given a band, B, we write $X_\pm(B) = \pi_\Sigma(\partial_\pm B \setminus \partial_V B \cap W)$. If $X_\pm(B) \neq \emptyset$, we write $H_0(B) = d_{\mathcal{G}(\pi_\Sigma(B))}(X_-(B), X_+(B))$.

If $\kappa(B) \geq 2$, we define a *ladder* in B as a sequence, $(\Omega_i)_{i=0}^n$, of annuli in $\mathcal{W}$ such that $\partial_V B = \bigcup_{i=0}^n \pi_V \Omega_i$ and with $\partial_V \Omega_i \cap \partial_V \Omega_j \neq \emptyset$ if and only if $|i-j| \leq 1$. We write $H(B)$ for the minimum length of a ladder in B. Note that $H(B) \geq H_0(B)$. By assumption, $H(B) \geq 1$. We refer to $H(B)$ as the *height* of B.

Definition : We say that B is t-taut if $H(B) \leq H_0(B) + t$.

If $\kappa(B) = 1$, we say that B is 0-*taut* if the sequence of annuli, $(\Omega_i)_i$, in B gives rise to a geodesic $(\pi_\Sigma \Omega_i)_i$ in the Farey graph, $\mathcal{G}(\pi_\Sigma B)$.

Definition : We say that a complete annulus system, W, is *recursively t-taut* if every band, B with $\kappa(B) \geq 2$ is t-taut, and every complexity-1 band is 0-taut.

The following is proven in [Bow4]

Proposition 5.1 : *Suppose that $a, b \subseteq X(\Sigma)$ are complete multicurves. Then there is a complete annulus system with $X_-(W) = a$, $X_+(W) = b$ and satisfying:*
(W1) If $\Omega, \Omega' \in \mathcal{W}$, with $\pi_\Sigma \Omega = \pi_\Sigma \Omega'$, then $\Omega = \Omega'$,
(W2) W is recursively t-taut, where $t = t(\kappa(\Sigma))$ depends only on $\kappa(\Sigma)$; and
(W3) The set $X(W)$ arises inductively by taking tight geodesics in subsurfaces of Σ. $\diamondsuit$

We won't make (W3) precise here. Its only relevant consequence is the "a-priori bounds" result in hyperbolic 3-manifolds which will allow us to construct certain lipschitz maps — see Proposition 6.1. It will not be needed in this section.

We say that two bands are *parallel* if they have the same base surface. If $B \subseteq A$ are parallel bands with $B \cap \partial_H A = \emptyset$, then the closures of $A \setminus B$, if non-empty, are parallel bands, which we refer to as *collars* of B. We say that B is *h-collared* if it admits collars (on both sides) of height at least h. We observe that two non-parallel 1-collared bands of the same complexity must be disjoint.

Suppose that $H(\Psi) \geq 1$ (i.e. no element of $\mathcal{W}$ crosses Ψ). Let $\mathcal{W}_\partial = \{\Omega \in \mathcal{W} \mid \Omega \cap \partial_H \Psi \neq \emptyset\}$. We form a 3-manifold, Λ, by opening out each $\Omega \in \mathcal{W} \setminus \mathcal{W}_\partial$ to a torus, $\Delta(\Omega)$, and each $\Omega \in \mathcal{W}_\partial$ to an annulus, also denoted $\Delta(\Omega)$. (Thus Λ can be thought of as the metric completion of $\Psi \setminus W$.) There is a natural map $\pi_\Psi : \Lambda \longrightarrow \Psi$ which folds every $\Delta(\Omega)$ back down to Ω, and is the identity elsewhere. Note that, if $\Omega \in \mathcal{W} \setminus \mathcal{W}_\partial$, then the torus $\Delta(\Omega)$ comes with a natural (horizontal) homotopy class of *longitude*, λ, and a natural class of *meridian*, μ, where $\pi_\Psi(\mu)$ is homotopically trivial.

We form another space, $P(W)$, by gluing in a solid torus, $T(\Omega)$, to each $\Delta(\Omega)$. If $\Omega \notin \mathcal{W}_\partial$, then we take μ to be a meridian curve of $T(\Omega)$. Otherwise, we glue such a torus along an annulus of degree 1 in ∂T. This $P(W) \cong \Psi$. We can extend the map π_Ψ to a map $\pi_\Psi : P(W) \longrightarrow \Psi$, so that $T(\Omega) = \pi_\Psi^{-1} \Omega$ for all $\Omega \in \mathcal{W}$. We can assume that the preimages of $\partial_\pm I$ are the surface boundary components of $P(W)$, which we denote $\partial_\pm P(W)$. We write $\partial P(W) = \partial_- P(W) \cup \partial_+ P(W)$. By a *band* in $P(W)$ we will mean the preimage of a band in Ψ.

Now suppose that W is a complete annulus system. A *brick* in Ψ is a maximal band, B, with the property that $W \cap B \subseteq \partial B$. A brick has either complexity 0 (3HS) or complexity 1 (4HS or 1HT). In the latter case, $B \cap \partial_H \Psi = \emptyset$, and each of $X_-(B)$ and $X_+(B)$ is a single curve. If these two curves are adjacent in $\mathcal{G}(\pi_\Sigma B)$, we say that B is *efficient*. (This happens, for example, if every complexity 1 band is 0-taut.) In this situation, there are combinatorially only three sorts of bricks, one for each base surface. Let $\mathcal{B}(W)$ be the set of all bricks. These tile Ψ — they cover, and their interiors are disjoint. Note that $\mathcal{B}(W)$ lifts to a decomposition of Λ into "bricks" — also denoted $\mathcal{B}(W)$. This, in turn, gives rise to a tiling of $P(W)$, where each tile is either a brick or a tube, $T(\Omega)$. We can put a fixed riemannian metric on Λ by choosing a riemannian metric for each of the three brick types. The exact choice doesn't matter much, since any two such constructions will yield bilipschitz equivalent metrics on Λ. However, it is convenient to assume that each vertical boundary component of each brick is isometric to $(\mathbf{R}/\mathbf{Z}) \times [0,1]$, that each 3HS brick is isometric to a hyperbolic 3HS times $[0,1]$, and that the combinatorial symmetries of all bricks are induced by isometries.

Note that for each $\Omega \in \mathcal{W} \setminus \mathcal{W}_\partial$, $\Delta(\Omega)$ has the structure of a euclidean torus — it is the union of some number, $L(\Omega)$, annular intersections with bricks. We refer to $L(\Omega)$ as the *length* of Ω (or of $\Delta(\Omega)$). It is the imaginary part of the complex modulus of $\Delta(\Omega)$.

There is also a "twisting" associated to Ω which might be described as follows. There is a curve, μ_0, of length at most $L(\Omega) + \frac{1}{2}$ in $\Delta(\Omega)$ which together with λ generates the first homology of $\Delta(\Omega)$. We refer to this as the *vertical meridian*. It can be thought of as the real part of the complex modulus of Ω (up to adding $\pm\frac{1}{2}$). We can assume that $\mu = \mu_0 + h\lambda$ for some $h \in \mathbf{Z}$. We write $H(\Omega) = h$. This is well defined up to adding ± 1. The vertical meridian can be described purely combinatorially. (There is a further decomposition of Λ into truncated octahedra and truncated simplices, as described in [Bow4]. The key point is that the number of tiles that μ_0 passes through is bounded in terms of $L(\Omega)$. The exact nature of this construction is not important here.)

Another way to define twisting is as follows. Let $B_+(\Omega)$ be the complexity-1 brick with $\partial_+\Omega \subseteq \partial_- B \setminus \partial_V B$. Thus, $X_-(B_+) = \pi_\Sigma \Omega = \alpha$, say. Let $\beta_+(\Omega) = X_+(B_+)$. Since we are assuming B_+ to be efficient, α and $\beta_+(\Omega)$ are adjacent in $\mathcal{G}(\pi_\Sigma B)$. In particular, $\iota(\alpha, \beta_+(\Omega)) \in \{1, 2\}$. We similarly define $\beta_-(\Omega)$. We set $H_0(\Omega) = \vec{p}(\alpha; \beta_-(\Omega), \beta_+(\Omega))$. Assuming we have defined our orientation conventions appropriately, we have:

Lemma 5.2 : *For all $\Omega \in \mathcal{W}$, we have $|H(\Omega) - H_0(\Omega)| \leq s_0$, where s_0 depends only on $L(\Omega)$ and $\kappa(\Sigma)$.*

Proof : The first step is to reduce to the case where $H(\Omega) = 0$. Given any $n \in \mathbf{Z}$, we can construct another annulus system, W', by taking $\omega'(u) = \omega(u)$ for $u \leq \pi_V(\partial_+\Omega)$, and taking $w'(u) = \tau_\delta$ for $u \geq \pi_V(\partial_+\Omega)$ where $\delta = \pi_\Sigma(\Omega)$. (Recall that $\omega(u)$ is the multicurve at level u. We may need to straighten things out above Ω, so that the curves become geodesic again.) This effectively corresponds to performing a Dehn surgery on $P(W)$. Now $\beta_-(\Omega)$ remains unchanged, while $\beta_+(\Omega)$ changes to $\tau_\delta^n \beta_+(\Omega)$. (Note that this operation respects the brick decompositions.) Thus, both $H(\Omega)$ and $H_0(\Omega)$ get increased by n, and so $H(\Omega) - H_0(\Omega)$ remains unchanged. Setting $n = -H(\Omega)$, we can therefore suppose that $H(\Omega) = 0$, as claimed.

The aim is now to bound $|H_0(\Omega)|$ in terms of $L(\Omega)$. Now since, by definition, $\mu = \mu_0 + H(\Omega)\lambda$, we have $\mu = \mu_0$. Thus, there are only boundedly many combinatorial possibilities for the arrangement of bricks meeting Ω together with the meridian curves. We claim that this is the only information needed to compute $H_0(\Omega)$, up to an additive constant.

More precisely, we construct a sequence $\beta_-(\Omega) = \beta_0, \beta_1, \ldots, \beta_{n+1} = \beta_+(\Omega)$ in $X(\Sigma)$, each crossing δ, where n is bounded, and where $\iota(\beta_i, \beta_{i+1})$ is bounded for all i. Thus $p(\delta; \beta_i, \beta_{i+1})$ is bounded, and the result then follows using the coarse triangle inequality. To this end, we choose a maximal sequence $u_1 = \pi_V(\partial_-\Omega) < u_2 < \cdots < u_n = \pi_V(\partial_+\Omega)$, such that $\omega_i = \omega(u_i)$ is a complete multicurve with $\omega_i \neq \omega_{i+1}$ for all i. Note that $\delta \in \omega_i$, and that n is bounded above (by $L(\Omega)$). Let F_i be the component of $\Sigma \setminus (\omega_i \setminus \delta)$ containing δ. This is either a 1HT or a 4HS. Recall that we have a map $\pi_\Psi : \Lambda(W) \longrightarrow \Psi$ which maps the torus $\Delta(\Omega)$ to Ω. The preimige of F_i in $\Lambda(W)$ consists of either a single 3HS with two boundary components identified under π_Ψ, or else two 3HS's each with a boundary component of that get dentified. We write $\delta_i^0, \delta_i^1 \subseteq \Delta(\Omega)$ for these boundary components. In $\Lambda(W)$ these 3HS's each have the structure of a hyperbolic 3HS with geodesic unit length boundary components. However the original identifications of δ_i^0 and δ_i^1 need not respect this metric. To rectify this, we realise the meridian $\mu \subseteq \Delta(\Omega)$ as a euclidean geodesic. This gives basepoints $p_i^0 = \mu \cap \delta_i^0$ and $p_i^0 = \mu \cap \delta_i^1$. We can now identify δ_i^0 and δ_i^1 isometrically so as to identify p_i^0 with p_i^1. (This implicitly takes account of the homotopy class of the embedding of F_i in Σ.) We now choose $\beta_i \in X(F_i) \subseteq X(\Sigma)$ so has to have bounded length in this structure. Now since all 4HS and 1HT bricks have bounded geometry, we see that $\iota(\beta_i, \beta_{i+1})$ is bounded above for all i. It follows that $p(\delta; \beta_i, \beta_{i+1})$ is bounded. $\diamondsuit$

Lemma 5.3 : *Suppose that $W \subseteq \Psi$ is a recursively t-taut annulus system. Write $a = X_-(W)$ and $b = X_+(W)$. There is some v, depending only on $\kappa(\Sigma)$ and t with the following properties.*

(1) If $\gamma \in X(\Sigma)$ and $p(\gamma; a, b) \geq v$, then $\gamma \in X(W)$.

(2) If $\Phi \in \mathcal{F}(\Sigma)$ and $p(\Phi; a, b) \geq v$, the Φ is the base surface of some band in Ψ.

Proof : (1) We shall prove the contrapositive by induction on the complexity of the base. So suppose $\gamma \notin X(W)$. We want to bound $p(\gamma; a, b)$.

By tautness, there is a ladder, $(\Omega_i)_{i=0}^n$, in Ψ with $n \leq H(\Psi) + t = d_{\mathcal{G}(\Sigma)}(a, b) + t$. Let $\beta_i = \pi_\Sigma \Omega_i$. Thus $(\beta_i)_i$ is a uniform quasigeodesic (with multiplicative constant 1). Choose $u_i \in \partial_V \Omega_{i-1} \cap \partial_V \Omega_i$, setting $u_0 = \partial_- I$ and $u_{n+1} = \partial_+ I$. Set $\omega_i = \omega(u_i)$. Then $\omega_0 = a$ and $\omega_{n+1} = b$. We can assume that ω_i is complete for all i. Thus, γ intersects each ω_i. Let O_i be the full band, $\Sigma \times [u_i, u_{i+1}]$. Thus $X_-(O_i) = \omega_i$ and $X_+(O_i) = \omega_{i+1}$.

Let $N = \{i \mid \beta_i \cap \gamma = \emptyset\} \subseteq \{0, 1, \ldots, n+1\}$. Note that this has bounded cardinality (since (β_i) is quasigeodesic in $\mathcal{G}(\Sigma)$).

Suppose that $\beta_i \cap \gamma = \emptyset$. Then γ lies in some component, F, of $\Sigma \setminus \beta_i$. Let $B = F \times [u_i, u_{i+1}]$. Thus, $X_-(B) \subseteq \omega_i$ and $X_+(B) \subseteq \omega_{i+1}$. Moreover, $\gamma \notin X(W_B) \subseteq X(W)$. Thus, by induction on κ, $p(\gamma; X_-(W_B), X_+(W_B))$ is bounded above in terms of $\kappa - 1$ and t, and so therefore is $p(\gamma; \omega_i, \omega_{i+1})$.

Suppose that $\{j, j+1, \ldots, k\}$ is a maximal set of consecutive indices in $\{0, \ldots, n+1\} \setminus N$. In other words, γ crosses β_i for all $j \leq i \leq k$, and so $\beta_i \in X(\Sigma, \gamma)$. But $\beta_j, \beta_{j+1}, \ldots, \beta_k$

is a quasigeodesic, and so by Theorem 4.2, $p(\gamma; \beta_j, \beta_k)$ is bounded above in terms of $\kappa(\Sigma)$ and t. Since $\beta_j \subseteq \omega_j$ and $\beta_k \subseteq \omega_{k+1}$, we see that $p(\gamma; \omega_j, \omega_{k+1})$ is also bounded.

But we have seen that $p(\gamma; ., .)$ satisfies a triangle inequality up to the additive constant 3. It therefore follows that $p(\gamma; a, b) = p(\gamma; \omega_0, \omega_{n+1})$ is bounded in terms of $\kappa(\Sigma)$ and t as claimed.

(2) This is essentially the same argument, with Φ replacing γ. $\diamondsuit$

We can now deduce:

Lemma 5.4 : *Suppose that $W \subseteq \Omega$ is a recursively t-taut annulus system satisfying (W1) (i.e. no base curve occurs more than once). Let $a = X_-(W)$ and $b = X_+(W)$. There is a constant $v = v(\kappa(\Sigma), t)$ with the following properties.*

(1) Suppose $\Omega \in \mathcal{W}$ and set $\gamma = \pi_\Sigma \Omega$. Then $p(\gamma; a, \beta_-(\Omega)) \leq v$ and $p(\gamma; b, \beta_+(\Omega)) \leq v$.

(2) Suppose B is a maximal band, with base surface, Φ. Then $p(\Phi, a, X_-(B)) \leq v$ and $p(\Phi; b, X_+(B)) \leq v$.

Proof : (1) Let $B_- \in \mathcal{B}(W)$ be the complexity-1 brick immediately below Ω, so that $\beta_-(\Omega) = X_-(B)$. If $\partial_- B_- \subseteq \partial_- \Psi$, then $\beta_-(\Omega) \in a$, and so $p(\gamma; a, \beta_-(\Omega)) = 0$. Otherwise, let O be the full band with $\partial_- O = \partial_- \Psi$ and $\partial_+ O \supseteq \partial_- B_-$. We can assume that $X_+(O)$ is complete. Note that $X_-(O) = a$ and $X_+(O) \supseteq \beta_-$. Now, by (W1), $\gamma \notin X(W_O)$. Thus, by Lemma 5.3(1) applied to O, we get that $p(\gamma; X_-(O), X_+(O))$ is bounded, and so $p(\gamma; a, \beta_-(\Omega))$ is bounded as claimed.

A similar argument applies to $\beta_+(\Omega)$.

(2) Let O be the full band with $\partial_- O = \partial_- \Psi$ and with $\partial_+ O$ slightly below $\partial_- B$, so that $X_-(B) \subseteq X_+(O)$. We can assume that $X_+(O)$ is complete. By the maximality of B, and using (W1), we see that Φ cannot be the base surface of any band in O. By Lemma 5.3(2), applied to O, we see that $p(\Phi; X_-(O), X_+(O))$ and hence $p(\Phi; a, X_-(B))$ is bounded.

Similarly $p(\Phi; b, X_+(B))$ is bounded. $\diamondsuit$

Corollary 5.5 : *With the same hypotheses as Lemma 5.4, there is some $v' = v'(\kappa(\Sigma), t)$ with the following properties.*

(1) Suppose $\Omega \in \mathcal{W}$. Set $\gamma = \pi_\Sigma \Omega$. Then $|H(\Omega) - \vec{p}(\gamma; a, b)| \leq v'$.

(2) Suppose that B is a band in Ω. Set $\Phi = \pi_\Sigma B$. Then $H(B) \leq p(\Phi; a, b) + v'$.

Moreover, if B is a maximal band with this base surface, then $|H(B) - p(\Phi; a, b)| \leq v'$.

Proof : (1) By Lemma 5.2, $|h(\Omega) - \vec{p}(\gamma; \beta_-(\Omega), \beta_+(\Omega))|$ is bounded. By Lemma 5.4(1), $p(\gamma; a, \partial_-(\Omega))$ and $p(\gamma; b, \partial_+(\Omega))$ are bounded. We now use the fact that $\vec{p}(\gamma; ., .)$ is a coarse 2-cycle.

(2) Suppose first, that B is maximal. By tautness, $|H(B) - p(\Phi; X_-(B), X_+(B))| \leq t$. (since $p(\Phi; X_-(B), X_+(B)) = d_{\mathcal{G}(\Phi)}(X_-(B), X_+(B)) = H_0(B)$). By Lemma 5.4(2), we see that $p(\Phi; a, X_-(\Phi))$ and $p(\Phi; b, X_-(\Phi))$ are bounded. It follows that $|H(B) - p(\Phi; a, b)|$ is bounded.

For the general case, note that if $B \subseteq C$ are parallel bands, then $H(B) \leq H(C) + s$, where s depends only on t. The result now follows by letting C be the maximal parallel band containing B, and applying the above. $\diamondsuit$

We will write $v_1 = v_1(\kappa(\Sigma))$ for $\max\{v, v'\}$, where v and v' are the constants featuring in Lemma 5.3 and Corollary 5.5.

Lemma 5.6 : *Suppose that W is recursively t-taut and satisfies (W1). Let $a = X_-(W)$ and $b = X_+(W)$. Suppose there exists $s \in \mathbf{N}$ and $k \geq 1$ such that for all $\Phi \in \mathcal{F}$ with $\kappa(\Phi) > k$, we have $p(a, b; \Phi) \leq s$. Then Ψ is covered by a union of at most N bands of complexity at most k, where $N = N(\kappa(\Sigma), s, t)$.*

Proof : By Lemma 5.3 and Corollary 5.5, the hypothesis amounts to placing a bound on the heights of all bands complexity at least $k + 1$. If $k \geq \kappa(\Sigma)$, there is nothing to prove — Ψ is covered by Ψ itself. So we can assume that $k \leq \kappa(\Sigma) - 1$. In particular, $d_{\mathcal{G}(\Sigma)}(a, b) \leq s$. So by tautness, there is a ladder, $(\Omega_i)_{i=0}^n$, in Ψ with $n \leq s + t + 2$. Let O_i be as constructed in the proof of Lemma 5.3.

We proceed by induction on complexity, where the hypothesis is reformulated as a bound on heights of bands. Note that $O_i \setminus \Omega_i$ is a (disjoint) union of a bounded number of bands of complexity less than $\kappa(\Sigma)$. Using Corollary 5.5, we have a bound on the heights of these bands. Hence, by induction, O_i itself, is a union of a bounded number of bands of complexity at most k. Since the number, $n + 1$, of such O_i is bounded, the result follows. $\diamondsuit$

Lemma 5.7 : *Suppose that W, a, b are as in Lemma 5.6, and suppose that there is some $s \geq 0$ such that $p(\Phi; a, b) \leq s$ for all $\Phi \in \mathcal{F}$. Then $|\mathcal{W}|$ is bounded in terms of $\kappa(\Sigma)$, s and t.*

Proof : By Lemma 5.6, Ψ is covered by a bounded number of bands of complexity 0 and 1. But (by Lemma 5.3 and Corollary 5.5 again) the height of any complexity 1 band is bounded, and so it contains a bounded number of annuli of $\mathcal{W}$. The statement now follows easily. $\diamondsuit$

Lemma 5.8 : *With the same hypotheses as Lemma 5.6, there is some $N_1 = N_1(\kappa(\Sigma), s, t)$ such that there are at most N_1 maximal bands of complexity k.*

Proof : It's not hard to see that any band of complexity at most k can meet only a bounded number of (depending on $k \leq \kappa$) of maximal bands of complexity (at least) k. Since, by Lemma 5.6, Ψ is covered by a bounded number of bands of complexity at most k, the result follows. $\diamondsuit$

We can now prove Proposition 4.6.

Proof of Proposition 4.6 : We are supposing that a and b contain complete multicurves, say α and β, and that $\iota(a)$ and $\iota(b)$ are uniformly bounded.

Let W be a recursively t-taut annulus system satisfying (W1), as given by Proposition 5.1, with $X_-(W) = \alpha \subseteq a$ and $X_+(W) = \beta \subseteq b$.

By Lemma 5.3 and Corollary 5.5, we can assume that $v_0(\kappa)$ is large enough so that any $\Phi \in \mathcal{F}$ with $p(\Phi; a, b) \geq v_0(\kappa)$ is the base surface of a band, and that any $\gamma \in X(\Sigma)$ with $p(\gamma; a, b) \geq v_0(\kappa)$ lies in $X(W)$.

Now if $k \geq 1$, then Lemma 5.8 gives us a bound on the number of maximal bands of complexity k in W. If $k = -1$, then Lemma 5.7 gives us a bound on $|\mathcal{W}|$.

This gives the result. $\diamondsuit$

We conclude this section with some further discussion of how an annulus system is used to construct models. We have already described the constructions of $P(W)$, with a bounded geometry metric on $\Lambda(W) \subseteq P(W)$. In this paper, we do not need to worry about the structure on the interiors of model tubes, $T(\Omega)$, though we have seen that $\partial T(\Omega) = \Delta(\Omega)$ is a euclidean torus.

Constructions of bi-infinite models along these lines were discussed in [Bow4]. Here there is a complicating factor in that model tubes may meet the boundary of $P(W)$. For some purposes it is convenient to use an "augmented model", $\hat{P}(W)$ by gluing collars, $\Psi_\pm \cong \Sigma \times [0, 1]$ to $\partial_\pm \Psi$.

We say that a hyperbolic structure, σ, on Σ is a *pants structure* (*based on* a) if there is some multicurve, a, all of whose elements have length 1 in Σ. Note that a determines σ up to twisting along these curves. If b is a complete marking of a, then there is a unique such pants structure such that each of the marking curves has minimal length. We denote this structure by σ_b. (The essential points in these definitions is that Σ should have injectivity radius uniformly bounded below, and that all the curve mentioned have length bounded above in terms of $\kappa(\Sigma)$. Indeed, if σ is another structure with these properties for some marking b, then σ and σ_b are bilipschitz equivalent, with constant depending only on $\kappa(\Sigma)$.)

We can now describe the "augmented model". Suppose that $W \subseteq \Psi$ is an annulus system. Let $\hat{I} \supseteq I$ be a compact interval containing I in its interior, and let $\hat{\Psi} = \Sigma \times \hat{I} \supseteq \Psi$. We write $\Psi_\pm$ for the collars, i.e. the closures of the components of $\hat{\Psi} \setminus \Psi$. Thus $W \subseteq \Psi$ is an annulus system in Ψ (except that it won't be vertically generic at $\partial_\pm \Psi$. We also have a brick decomposition, $\hat{\mathcal{B}}(W) \supseteq \mathcal{B}(W)$, where Ψ_+ and Ψ_- are considered single bricks. To put a geometric structure on them, we need extra combinatorial information, namely, markings, $b_\pm$, of the complete multicurves $a = X_\pm(W) = \omega(\partial_\pm I)$. These markings determine hyperbolic structures, $\sigma_\pm = \sigma_{b_\pm}$ on Σ. We now give $\Psi_\pm$ the product structure, $\Sigma \times [0, 1]$. This gives us an augmented model, $\hat{P}(W)$.

We can now similarly define the twisting, $H(\Omega)$, of an annulus Ω meeting $\partial_\pm \Psi$. The vertical meridian, μ_0, again have length in $\hat{P}(W)$ bounded by $L(\Omega) + \frac{1}{2}$. One way to view this is as follows. If Ω meets $\partial_- \Psi$ we can construct another annulus system, W', by pushing Ω slightly off $\partial_- \Psi$, and inserting another annulus whose base curve is the marking curve, β, of $\alpha = \pi_\Sigma \Omega$ in b. Now $H(\Omega)$ can be viewed as the twisting of Ω in W' (at least up to an additive constant). If Ω does not meet $\partial_+ \Psi$, then as in Corollary 5.5, we see that $|H(\Omega) - \vec{p}(\gamma; b_-, b_+)|$ is bounded. (Note that it is conceivable that β might be the base curve of another annulus in W', but that doesn't matter for the relevant part of the argument.) We can perform similar constructions if Ω meets $\partial_+\Psi$, or indeed $\partial_-\Psi$ and $\partial_+\Psi$.

We note:

Lemma 5.9 : *Suppose that $b_\pm$ are complete markings of multicurves $a_\pm$, and let W be recursively t-taut and satisfy (W1), and with $X_\pm(W) = a_\pm$. Then for all $\Omega \in \mathcal{W}$, we have $|H(\Omega) - \vec{p}(\gamma; b_-, b_+)| \leq s$, where s depends only $\kappa(\Sigma)$ and t. Here $H(\Omega)$ refers to the twisting in the augmented model, $\hat{P}(W)$.*

Proof : If $\Omega \cap \partial_H \Psi = \emptyset$ this is covered by Corollary 5.5. Otherwise, it is covered by the above discussion. ◊

We remark that we can use this to reprove Proposition 4.3.

Proof of Proposition 4.3 : Let W be a t-taut annulus system satisfying (W1), as given by Proposition 5.1, where $t = t(\kappa(\Sigma))$, and let $\hat{P}(W)$ be the augmented model constructed as above using the markings a and b. We want to show that $\iota(a, b)$ is bounded.

Now Lemma 5.7 tells us that $|\mathcal{W}|$ is bounded in terms of $\kappa(\Sigma)$ and q. Applying Lemmas 5.3 and 5.9, we see that the twisting, $H(\Omega)$, is bounded for each $\Omega \in \mathcal{W}$. Now $|\mathcal{B}(\Sigma)|$ is bounded in terms of $|\mathcal{W}|$, and so there are only boundedly many ways of gluing the bricks together along their horizontal boundaries in order to reconstruct $\Lambda(W)$. Moreover, for each $\Omega \in \mathcal{W}$ with $W \cap \partial\Psi = \emptyset$, there are only boundedly many choices of the meridian curve. Thus, there are only boundedly many ways to reconstruct $P(W)$. But $\iota(a, b)$ is completely determined by this information, and so must be bounded. The result follows. ◊

As remarked in Section 4, this can also be deduced from the projection formula given in [MaM2]. (Indeed an elaboration of the above argument can be used to give another proof of this formula.)

6. Hyperbolic 3-manifolds.

In this section, we will use the theory of hyperbolic 3-manifolds to prove the following purely combinatorial statement. It will will be used in Section 7 to give us a basepoint with respect to which we can measure bending along subsurfaces, as we outlined in Section 2.

Lemma 6.1 : *$(\forall \kappa)(\exists r_1, s_1)(\forall r_2)(\exists r_3, r_4)$ such that the following holds. Suppose that Σ is a closed surface with $\kappa(\Sigma) = \kappa$, and that $\phi : \Sigma \longrightarrow \Sigma$ is a pseudoanosov mapping class. Then one of the following situations arises. Either:*

(1) There is a complete multicurve, a, in Σ, within distance s_1 of an axis of ϕ in $\mathcal{G}(\Sigma)$, and there is some $\Phi \in \mathcal{F}(\Sigma)$ with $p(\Phi; a, \phi a) \geq r_2$. Moreover, if $F \in \mathcal{F}(\Sigma)$ with $p(F; a, \phi a) \geq r_1$ and $p(F; \phi^n a, \phi^{n+1} a) \geq r_1$, then $n = 0$. Or:

(2) There is a marking, b, of Σ, within distance s_1 of an axis of ϕ in $\mathcal{G}(\Sigma)$, and satisfying the following. For all $\Phi \in \mathcal{F}$, we have $p(\Phi; b, \phi b) \leq r_4$. Moreover, if $\gamma \in X(\Sigma)$ and $p(\gamma; b, \phi b) \geq r_3$ and $p(\gamma; \phi^n b, \phi^{n+1} b) \geq r_3$, then $n = 0$.

Recall that an axis of ϕ is a bi-infinite geodesic in $\mathcal{G}(\Sigma)$ that is moved a bounded Hausdorff distance. Necessarily, this distance depends only on the hyperbolicity constant, hence only on $\kappa(\Sigma)$.

We begin by recalling some basic facts about product hyperbolic 3-manifolds.

Suppose that $M = \Sigma \times \mathbf{R}$ is hyperbolic and has no cusps. There is no preferred homeomorphism. However, we will choose a "preferred" class of homotopy equivalence between Σ and M. Recall, from Section 3, that any $\gamma \in X(\Sigma)$ can be realised as a closed geodesic, γ_M, in M. We denote its length by $l_M(\gamma)$.

Suppose $\eta > 0$ is less than the usual 3-dimensional Margulis constant and let $\mathcal{T} = \mathcal{T}(\eta)$ be the set of Margulis tubes (with core curves less than η in length). Provided η is sufficiently small in relation to $\kappa(\Sigma)$. This set is unlinked in M [O2]. If $T \in \mathcal{T}$ then ∂T is a euclidean torus with a natural longitude and meridian defined up to free homotopy. These can be realised as geodesics in ∂T. The length of the longitude is bounded above and below in terms of $\kappa(\Sigma)$ and η. These statements again follow from the construction of [O2], see [Bow1] for some elaboration. There is a complex modulus associated to ∂T, where the marking of ∂T is chosen so that the longitude is 1 and the imaginary part of the meridian is positive. We write $L_M(\partial T)$ for the imaginary part of the complex modulus. This can be thought of as a measure of "vertical length" and is thus analogous to the quantity $L(\Omega)$, defined in Section 5, for the vertical length of an annulus in an annulus system. We shall write $l_M(T)$ for the length of the core geodesic of T.

In addition to the properties above, further requirements on the Margulis constant arise from Proposition 6.10 and Lemma 6.12 below, again depending on $\kappa(\Sigma)$. Subject to these, we will fix a Margulis constant, η_0. We will usually supress mention of it when stating the results.

We need to discuss bands in M. These are used in [Bow1] and [Bow4] and we refer to those papers for elaboration. Analagous notions of "scaffolds" are used in [Mi2,BrCM].

The terminology runs parallel to that of Section 5. (The topology of the set-up is similar to that of the model, $P(W)$, described there.) A *strip*, C, is the image, under a homeomorphism $\Sigma \times \mathbf{R} \longrightarrow M$, of a subset of the form $\Phi \times J$, where Φ is a subsurface of Σ, and $J \subseteq \mathbf{R}$ is a non-trivial compact interval. Taking this homeomorphism in the preferred class, we refer to $\Phi \in \mathcal{F}$ as the *base surface* of C, which we denote by $\pi_\Sigma C$. We write $\kappa(C) = \kappa(\pi_\Sigma C)$ for its *complexity*. We write $\partial_V C = \partial\Phi \times J$, $\partial_\pm C = \Phi \times \partial_\pm J$ and $\partial_H C = \Phi \times \partial J = \partial_+ C \sqcup \partial_- C$. By a *band* we mean a strip such that each component of $\partial_V C$ lies in ∂T for some $T \in \mathcal{T}$. (We refer to these tubes as *boundary tubes*.) Moreover, we assume that for all other $T \in \mathcal{T}$, $T \cap C$ is either empty or a solid torus, and that $T \cap \partial_\pm C$ is at most one annulus. We write $\mathcal{T}_I(C)$ for the set of tubes that meet C in a solid torus. By a *fibre* of C, we mean an embedding of Φ in $C \setminus \partial_H C$ in the preferred homotopy class. (Such a fibre cuts C into two subbands.) We say that two bands are *parallel* if they have the same base surface.

We can put a riemannian pseudometric, ρ_C, on C by shrinking each $T \in \mathcal{T}_I(C)$ to have diameter 0. We can then define the *height*, $H_M(C)$, of C to be the distance, in C, between $\partial_- C$ and $\partial_+ C$ in this pseudometric. It turns out that every point of C lies on a fibre whose ρ_C-diameter in the pseudometric is bounded above in terms of $\kappa(\Sigma)$ (and the Margulis constant) see for example Canary's Filling Theorem [C]. The following is a consequence:

Lemma 6.2 : *There is some $H_0 = H_0(\kappa(C))$ such that if $x, y \geq 0$ and C is a band with $H_M(C) \geq H_0 + x + y$ then it can be split into two horizontal bands, C_1 and C_2 (by the fibre $\partial_+ C_- = \partial_- C_+$) with $H_M(C_1) \geq x$ and $H_M(C_2) \geq y$.* ♢

If $C \subseteq C' \setminus \partial_H C'$ are parallel bands, we refer to $(C' \setminus C) \cup \partial_H C$ as *collars* of C. We call them *h-collars* if they have height at least h. We say that C is h-collared if it admits a pair of h-collars.

Lemma 6.3 : *There is some $h_0 = h_0(\kappa(\Sigma))$ such that any two non-parallel h_0-collared bands of the same complexity are disjoint.* ♢

We also note that there is a bound on the heights of all 3HS bands, which we may as well take to be the same h_0 (see [Bow1]).

The following is shown in [Bow1]. It also follows from the corresponding statement for annulus systems, via [Bow4].

Lemma 6.4 : *Given any $h \geq 0$, there is a some L, depending only on h and $\kappa(\Sigma)$ such that if every band has height at most h, then for all $T \in \mathcal{T}$ we have $L_M(\partial T) \leq L$.* ♢

Next we need to consider a few technical points about realising homotopy equivalences of Σ into M. All maps from Σ to M will be assumed to lie in our preferred homotopy class.

Definition : We say that a map $g : \Sigma \longrightarrow M$ is *(topologically) efficient* (with respect to c) if $g^{-1}(\bigcup \mathcal{T})$ is a regular neighbourhood of a (possibly empty) multicurve c in Σ.

After small adjustment, we can always assume that $g(\Sigma)$ crosses the boundary of any Margulis tube, T, transversely, that is, $f^{-1}(T)$ is either empty or consists of two curves.

Suppose that σ is a pants structure with respect to a complete multicurve $a \supseteq c$. We can assume that $g^{-1}(\bigcup \mathcal{T})$ is an arbitrarily small neighbourhood of c in Σ.

Definition : We say that an efficient map g is ξ-efficient if $g|\Sigma \setminus f^{-1}(\bigcup \mathcal{T})$ is ξ-lipschitz.

We again allow $c = \emptyset$, in which case we just have a ξ-lipschitz map into $M \setminus \bigcup \mathcal{T}$.

We can make similar definitions for homotopies. Suppose that $f : \Sigma \times [0, 1] \longrightarrow M$ is a homotopy between maps g_0, g which are efficient respectively with respect to the empty set and the multicurve, c.

Definition : We say that f is efficient if for each $\alpha \in c$ there is some $T \in \mathcal{T}$ such that $f^{-1}T$ is a regular neighbourhood of $\alpha \times [0, \frac{1}{2}]$.

The union of these regular neighbourhoods is therefore a regular neighbourhood, N, of $((\bigcup c) \times [0, \frac{1}{2}])$.

Definition : We say that such a homotopy, f, is ξ-*efficient* if $f|\Sigma \setminus N$ is ξ-lipschitz.

Note that we allow the homotopy f to meet Margulis tubes other than those associated with c. However, the lipschitz property means there is a bound on how deeply such a homotopy can enter such a tube.

Suppose that $g : \Sigma \longrightarrow M$ is efficient with respect to c. Let σ_0 be the pullback metric on Σ. Suppose that this has curvature at most -1 (allowing for cone angles at least 2π). We can find a complete multicurve $a \supseteq c$ such that the minimal length $l_{\sigma_0}(\alpha)$ of α is bounded above in terms of $\kappa(\Sigma)$ for all $\alpha \in a$. In fact, if we have a multicurve $a \supseteq c$ with $l_{\sigma_0}(\alpha) \leq l$ for all $\alpha \in a$, then there is a pants structure σ on Σ, based on a such that the map g is ξ-efficient where ξ depends only on $\kappa(\Sigma)$ and l. (Thus, if we have chosen a particular multicurve a, as above, then ξ will depend only on $\kappa(\Sigma)$.) We also note that σ is determined up to twisting on a. The twist parameters on c will be irrelevent, and those on $a \setminus c$ are determined within bounds depending on $\kappa(\Sigma)$ and l.

One source of examples of efficient maps arise from minimal surfaces (see Lemma 6.12). Another source arises by realising a multicurve geodesically in M.

Definition : We say that an efficient map $g : \Sigma \longrightarrow M$ is *straight* (based on a) if the induced metric on Σ has curvature at most -1 and if there is a complete multicurve $a \supseteq c$ such that each $\alpha \in a$ gets mapped locally injectively to a closed geodesic, α_M, in M.

Clearly if $\alpha \in c$, then this is the core of the Margulis tube.

Such surfaces arise as "pleated surfaces", see [T1,Bon]. Briefly we triangulate Σ so that $\bigcup a$ contains the 0-skeleton and is contained in the 1-skeleton. We realise Σ so that each element of a is mapped to a closed geodesic and each triangle is mapped in totally geodesically. The pull-back metric is hyperbolic with cone singularities. One can further "spin" around the curves to obtain a pleated surface where the pull-back metric is hyperbolic without singularities. Such a surface will always be efficient (provided the Margulis constant is chosen small enough in relation to $\kappa(\Sigma)$). This holds regardless of the lengths of the closed geodesics.

Definition : We say that a straight map based on a is l-straight if $l_M(\alpha) \leq l$ for all $\alpha \in a$.

In practice, this l will be bounded in terms of $\kappa(\Sigma)$.

We can homotope lipschitz maps to straight ones. The following gives some geometric control of over this process.

Lemma 6.5 : *Suppose that σ is a pants metric with respect to a complete multicurve, a, and $g : \Sigma \longrightarrow M \setminus \bigcup \mathcal{T}$ is a ξ-lipschitz map. Suppose that $l_M(\alpha) \leq l$ for all $\alpha \in a$. Then there is a ξ'-efficient homotopy from g to a straight map with respect to a, where $\xi' = \xi'(\xi, l)$ depends only on ξ, l. In fact, if $g_0 : \Sigma \longrightarrow M$ is any straight map based on a, then we can assume that the final stage of the homotopy factors through g_0.*

(We remark that in our applications, l will depend only on $\kappa(\Sigma)$, so ξ' will depend only on ξ and $\kappa(\Sigma)$.)

Note that the straight map constructed will necessarily be l-straight. The map g_0 will be efficient with respect to some multicurve $c \subseteq a$.

Proof : The construction of a straight map, g_0, based on a complete multicurve derived from [T1,Bon] was outlined above. We write Σ_0 for its domain so that there is a cononical homotopy equivalence from Σ to Σ_0. Write σ_0 for the induced path metric on Σ_0. We want to construct a homotopy, f, that is ξ'-efficient with respect to the original pants metric, σ. We only give an outline, since similar constructions appear elsewhere, for example in [Bow4]. For simplicity, we describe it mainly in M, though formally, the constructions are best carried out equivariantly in $\tilde{M} = \mathbf{H}^3$.

We will use the following observation. Suppose that β_M and γ_M are closed geodesics in M each with length bounded below by η_0 and above by some larger constant, l. Consider a relative homotopy class, δ, of paths from β_M to γ_M, where we are allowed to slide the endpoints of δ along the curves. Suppose we have two realisations of δ, both of length at most some other constant l'. Then in fact we can homotope between them, only sliding the endpoints a bounded distance along the curves. This can be seen by lifting to $\mathbf{H}^3$, and noting that $\tilde{\beta}_M$ and $\tilde{\gamma}_M$ must diverge uniformly, and the lifts of the two realisations of δ lie in a set of bounded diameter. The above remains true if one, or both, of β_M and γ_M are replaced by an η_0-Margulis tube. In this case, δ is reinterpreted as a homotopy class of paths in M with the corresponding endpoint constrained to move in the boundary of the tube.

We first consider the case where $l_M(\alpha) \geq \eta_0$ for all $\alpha \in a$, so that $c = \emptyset$, and we are looking for a uniformly lipschitz homotopy (without worrying about Margulis tubes).

Let $\alpha \in a$, and let α_M be the closed geodesic representative of the homotopy class of $g(\alpha)$. The lifts of $g(\alpha)$ and α_M are a bounded distance apart in $\mathbf{H}^3$ (since the former curve is uniformly quasigeodesic). This bound depends only on l and η_0. The nearest point projection is equivariant, and in M gives us a uniformly lipschitz homotopy from $g(\alpha)$ to α_M. After further lipschitz homotopy in α_M, we can assume that the final path has constant speed (that is parametrised proportionally to the arc length in α induced by the metric σ). Performing this for each such α, we get a lipschitz homotopy from $g(\bigcup a)$ to a map $g' : \bigcup a \longrightarrow M$ that maps each $\alpha \in a$ to a constant speed geodesic.

We claim we can extend g' to a uniformly lipschitz parameterisation of $g_0(\Sigma_0)$, more precisely a map with domain Σ that factors through a uniformly bilipschitz map from (Σ, σ) to (Σ_0, σ_0). Let F be a component of $\Sigma \setminus \bigcup a$. We cut this into two congruent right angled hexagons by three geodesic arcs connecting boundary components (in the metric σ). Suppose ϵ is one such arc, connecting $x \in \beta \in a$ to $y \in \gamma \in a$, say. Combining with the trajectories of the homotopy from $g(\beta \cup \gamma)$ to $\beta_M \cup \gamma_M$, we get a path, δ, of bounded length from $g'(x) \in \beta_M$ to $g'(y) \in \gamma_M$. Let $h = g_0^{-1} \circ g : \bigcup a \longrightarrow \bigcup a$, and let $x' = h(x) \in \beta$ and $y' = h(y) \in \beta$. Since σ_0 has curvature at most -1, there is an arc, δ_0, of bounded length connecting β_M to γ_M in (Σ_0, σ_0). Applying the earlier observation, we can homotope it to δ in M, sliding its endpoints bounded distances in β_M and γ_M. It follows that there is a path, ϵ', in Σ_0 of bounded length connecting x' to y' and homotopic to δ fixing these endpoints. We can find such an arc for every arc ϵ, giving us a collection of arcs ϵ' which together with $\bigcup a$ cuts Σ_0 into hexagons. We now extend h so that it sends each ϵ linearly to ϵ', and then over the complementary hexagons to give a bilipschitz map from Σ to Σ_0. Let $g' = g_0 \circ h : \Sigma \longrightarrow M$. This agrees with g on $\bigcup a$. We now claim there is a lipschitz homotopy from g to g'. For example, note that for each path ϵ above, the trajectories of the homotopies of x and y together with $g(\epsilon)$ and $g'(\epsilon)$ form a closed curve of bounded length in M which we can span by a lipschitz disc. We do this for each such ϵ. We now have the homotopies defined on the complement of a set of 3-balls in $\Sigma \times [0, 1]$, and we extend over these to get a uniformy lipschitz homotopy $\Sigma \times [0, 1] \longrightarrow M$.

For the general case, suppose we have $\beta \in a$ with $l_M(\beta) < \eta_0$. Let $T = T_M(\beta) \in \mathcal{T}$ be the Margulis tube with core curve β_M. By assumption, $g(\beta) \cap T = \emptyset$, so we can find a uniformly lipschitz homotopy of $g(\beta)$ to a constant speed euclidean closed geodesic β' in ∂T, by first taking nearest point projection (in the universal cover) and then straightening it out by a lipschitz homotopy in ∂T. We can also assume that $g_0(\Sigma_0)$ meets ∂T in two constant speed euclidean geodesics curves (again after performing a lipschitz homotopy in ∂T). We claim that we can further homotope these curves a bounded distance in ∂T to β'. In fact, suppose that F is a component of $\Sigma \setminus \bigcup a$, and with $\beta \subseteq \partial F$, and let $\gamma \subseteq \partial F$ be another boundary curve. Let δ be homeomorphism from β to γ as before. We can realise this as a path of bounded length connecting β' to γ_M or to the corresponding curve γ' in $\partial T_M(\gamma)$ (depending on the length of γ_M). It can also be realised as a curve, δ_0, of bounded length in Σ_0. From the earlier observation, we can homotope δ to δ_0 sliding the endpoints a bounded distance in $\partial T_M(\beta)$ and in γ_M or $\partial T_M(\gamma)$. We can therefore find a lipschitz homotopy from the corresponding component of $g_0(\Sigma_0) \cap \partial T_M(\beta)$ to β', so we may as well assume these curves are equal. A similar argument applies to γ. We can now connect the endpoints of δ by a homotopic curve of bounded length in Σ_0, and the argument then proceeds similarly as before. $\diamondsuit$

We can generalise the terminology to a subsurface, Φ, of Σ. We define a map $g : \Phi \longrightarrow M$ to be *efficient* if $g^{-1}(\bigcup \mathcal{T})$ is a union of $\partial \Phi$ and disjoint non-homotopic essential annuli. Note that $g(\partial \Phi) \subseteq \partial(\bigcup \mathcal{T})$. We say that it is ξ-*efficient* if its restriction to $g^{-1}(\Phi \setminus \bigcup \mathcal{T})$ is ξ-lipschitz with respect to a pants metric on Φ associated to a multicurve that includes the core curves of $g^{-1}(\bigcup \mathcal{T})$.

We note

Lemma 6.6 : *For all $\xi \geq 0$ there exists $h(\xi) \geq 0$ such that if C is a band with $h(\xi)$-collars $C_\pm$ and $g : \Sigma \longrightarrow M$ is a ξ-efficient map with $g(\Sigma) \cap C \neq \emptyset$. Then $g|\Phi$ is efficient and $g(\Phi) \subseteq C_- \cup C \cup C_+$, where Φ is (homotopic to) the base surface of C.*

Proof : Let $C' = C_- \cup C \cup C_+$. The diameter of $g(\Sigma) \cap C'$ with respect to the pseudometric $\rho_{C'}$ is bounded above in terms of $\kappa(C) \leq \kappa(\Sigma)$ and ξ. Thus if $h(\xi)$ is sufficiently large, it cannot cross between the horizontal boundaries of C_- or of C_+. This means it can only exit via the vertical boundary, into the boundary Margulis tubes. Since g is assumed efficient on Σ, these tubes must correspond to components of the associated multicurve. In particular, the components of $\partial\Phi$ lie in the multicurve. It also follows that $g^{-1}C$ is homotopic to Φ in Σ, so we can take it to be equal to Φ. By definition, it follows that $g|\Phi$ is efficient. (For a more detailed argument, see for example, [Bow4].) $\diamondsuit$

Note that we have only really used the fact that any component of $g(\Sigma) \cap C$ has bounded ρ_C-diameter.

We can also extend the terminology of a "straight map" to a subsurface $\Phi \subseteq \Sigma$. We will only use this terminology when each component of $\partial\Phi$ corresponds to a core of a Margulis tube. To be consistent with the terminology of efficient maps, let $\Phi_0 \supseteq \Phi$ be a homotopic surface so that $\Phi_0 \setminus \Phi$ is a regular neighbourhood of $\partial\Phi_0$. By a *straight map* of Φ, we mean a map $g : \Phi \longrightarrow M$, which can be extended to Φ_0 so that each component of $\partial\Phi_0$ gets mapped to the core of a Margulis tube, and so that the induced metric on $\partial\Phi_0$ has curvature at most -1. As with straight maps of Σ, we can assume that $g|\Phi$ is efficient. In particular, $g|\Phi_0 \setminus \Phi$ homotopes boundary curves of Margulis tubes to core geodesics.

If C is a band in M, with base surface, Φ, we say that an efficient map g "maps into" C if $g(\Phi) \subseteq C \setminus \partial_H C$. Such a map is a relative homotopy equivalence from $(\Phi, \partial\Phi)$ to $(C, \partial_V C)$.

We note the following generalisation of Lemma 6.5, which follows by a similar argument:

Lemma 6.7 : *Suppose that Φ is a subsurface of Σ, and σ is a pants metric on Φ based on a complete muticurve, a, in Φ and $g : \Sigma \longrightarrow M$ is a ξ-lipschitz map, with $g^{-1}(\bigcup \mathcal{T}) = \partial\Phi$. Suppose that $l_M(\alpha) \leq l$ for all $\alpha \in a$. Then there is a ξ'-efficient homotopy from g to a straight map with respect to a, where ξ' depends only on ξ, l and the Margulis constant η.* $\diamondsuit$

Every sufficiently long band will contain a straight map:

Lemma 6.8 : *There is some $h_2 = h_2(\kappa(\Sigma))$ and $\xi_0 = \xi_0(\kappa(\Sigma))$ such that if C is a band of height at least h_2, with base surface Φ, then there is an ξ_0-straight map, $g : \Phi \longrightarrow M$ into C. Moreover, if $\mathcal{C}(C)$ is a collection of h_2-collared bands in C, not parallel to C, then $g(\Phi)$ can be assumed to be disjoint from $\bigcup \mathcal{C}(C)$.*

Proof : The first statement is a simple consequence of [Bow4]. (For example, note that every point in C is a bounded ρ_C-distance from a closed geodesic, γ_M in C which is homotopic to a simple closed curve, γ, in Φ. We can now construct a straight map whose associated multicurve includes γ (cf. [Bon]). We now take a multicurve of bounded length in the induced metric on Φ, and use this to construct a new straight map. This will be a bounded ρ_C-distance from the original, and now all the associated geodesics have uniformly bounded length.)

In fact, by choosing h_2 large enough we can arrange that $g(\Phi)$ is a long way from $\partial_H C$ in ρ_C. Now if $D \in \mathcal{C}(C)$ with collars $D_\pm$ of height h_2, then the ρ_C-diameter of $(D_- \cup D \cup D_+) \cap C$ is bounded. Thus, in particular, we can suppose that if $g(\Phi)$ meets $D \in \mathcal{C}(C)$ then $D_- \cup D \cup D_+ \subseteq C \setminus \partial_H C$. If such a D exists, we choose it to have maximal complexity. Now since $g(\Phi)$ cannot cross the collars of D (assuming h_2 is big enough), then $g|F$ is also straight, where F is the base surface of D (since it cannot cross the collars, $D_\pm$, cf. Lemma 6.6). Indeed, $g|F$ maps into $D_- \cup D \cup D_+$. We now modify g on F so that it maps into D_+, using the first statement of Lemma 6.8 applied to D_+. This reduces the number of elements of $\mathcal{C}(C)$ of maximal complexity meeting $g(\Phi)$. We now proceed by downward induction on complexity to eliminate all such intersections. $\diamondsuit$

Note that, for each curve α in the associated multicurve, we have $l_M(\alpha) \leq \xi_0$.

The following is a purely topological observation.

Lemma 6.9 : *Suppose C is a band, and let $g : F \longrightarrow C$, be a relative homotopy equivalence from $(F, \partial F)$ to $(C, \partial_V C)$. Suppose that $\mathcal{C}(C)$ is a collection of bands in C and that $g(F) \cap \bigcup \mathcal{C}(C) = \emptyset$. Then g is homotopic in $C \setminus \bigcup \mathcal{C}(C)$, rel $\partial_V C$, to an embedding, $(F, \partial F) \hookrightarrow (C, \partial_V C)$.*

Proof : Let N be a regular neighbourhood of $g(F)$ in $C \setminus \bigcup \mathcal{C}(C)$. Since g is a relative homotopy equivalence to C, it follows from the construction of [FHS] that g is homotopic in C to an embedding $g' : F \longrightarrow N \subseteq C$ (cf. [Bon]). In fact, their argument shows that there is a sequence of maps $g = g_1, \ldots, g_n = g'$ into N such that each g_i is homotopic in C to g_{i+1} via a homotopy supported on a disc or an annulus in F. Such a homotopy can be pushed off any band in $\mathcal{C}(C)$. $\diamondsuit$

We now bring the model space into play. Suppose that W is an annulus system in $\Psi = \Sigma \times I$, and $P(W) \equiv \Sigma \times I$ is constructed as in Section 5. Thus $\Lambda(W) = P(W) \setminus \bigcup_{\Omega \in \mathcal{W}} \operatorname{int} T(\Omega)$. We won't need to worry about the metric on the interior of model tubes here.

Definition : We say that a map, $f : P(W) \longrightarrow M$, in the preferred homotopy class is *efficient* if, for all $T \in \mathcal{T}$, there is some $\Omega \in \mathcal{W}$ with $f^{-1}(T) \subseteq T(\Omega)$. We say that it is ξ-efficient if, in addition, $f|\Lambda(W)$ is ξ-lipschitz. We say that it is straight if for each $\Omega \in \mathcal{W}$, $T(\Omega)$ gets mapped either into an element of $\mathcal{T}$ or to a closed geodesic.

The following notation will be convenient. If $\gamma \in X(\Sigma)$ corresponds to $\Omega \in \mathcal{W}$ and $T \in \mathcal{T}$ in the above definition, we write $\Omega = \Omega$, $T_P(\gamma) \subseteq T(\Omega(\gamma))$, $\Delta_P(\gamma) = \Delta(\Omega) = \partial T_P(\gamma)$, $T_M(\gamma) = T$, $\Delta_M(\gamma) = \partial T$. We can assume that $f^{-1}\Delta_M(\gamma) = \Delta_P(\gamma)$. We will also write $H_P(\gamma) = H_P(\Omega) = H(\Omega)$, $H_P(B) = H(B)$ and $L_P(\Omega) = L(\Omega)$ when referring to the combinatorial quantities defined for an annulus system in Section 5.

We recall from Section 5 the construction of a complete annulus system, W, and the associated model, $P(W)$, starting from two complete multicurves.

Proposition 6.10 : $(\forall l, \kappa)(\exists \xi = \xi(l, \kappa))$ *with the following property. Suppose that* M *is as above with* $\mathcal{T}$ *the set of Margulis tubes. Suppose that* a_+, a_- *are complete multicurves in* Σ *with* $l_M(\alpha) \leq l$ *for all* $\alpha \in a_+ \cup a_-$. *Suppose that* $W \subseteq \Psi$ *is a complete annulus system connecting* a_+ *and* a_-, *and that* $P(W)$ *is the associated model. Then there is a* ξ*-efficient straight map* $f : P(W) \longrightarrow M$ *with the following properties:*

(1) $(\forall h)(\exists h' = h'(h, l, \kappa))$ *if* $B \subseteq \Psi$ *is a band with* $H(B) \geq h'$, *then there is a band* $C \subseteq M$ *with the same base surface, with* $H(C) \geq h$ *and* $C \cap f(\partial P(W)) = \emptyset$ *and such that* $f|f^{-1}C$ *maps to* C *with non-zero degree.*

(2) $(\forall h)(\exists h'' = h''(h, l, \kappa))$ *if* $C \subseteq M$ *is a band with with* $H(C) \geq h''$ *and* $C \cap f(\partial P(W)) = \emptyset$, *and such that* $f|f^{-1}C$ *maps to* C *with non-zero degree, then there is a band* $B \subseteq \Psi$ *with the same base surface and with* $H(B) \geq h$.

Moreover, $f|\partial_\pm P(W)$ *is straight with respect to* $a_\pm$. *Indeed, we can take it to be any given straight map with respect to* $a_\pm$. ◊

Proposition 6.10 follows from the constructions of [Bow4]. There it was described in terms of a bi-infinite geodesic, though the same construction can be applied to a finite segment. In this case, the boundary components get mapped to straight surfaces.

Proposition 6.10 says nothing explicit about the geometry of Margulis tubes. To deal with this, the following observation will be useful.

Let Δ be a euclidean torus with preferred longitude, λ, of length 1, and preferred meridian, μ. Write the complex modulus as $L + iH$, where $L \geq 0$ and $H \in \mathbf{R}$. (Note that if $\Delta = \Delta_P(\gamma)$ above, then $L = L_P(\gamma)$ and $|H - H_P(\gamma)| \leq \frac{1}{2}$.) Let $S = \mathbf{R}/\mathbf{Z}$ be the unit length circle.

Lemma 6.11 : *Suppose that* $f : \Delta \longrightarrow S$ *be a* ξ*-lipschitz map with* $f|\lambda$ *of degree 1 and* $f|\mu$ *homotopically trivial. Then* $|H| \leq \xi L$.

Proof : The proof is elementary (a similar argument is used in [Bow4]). ◊

We now move on to consider surface bundles over circles. Let $\phi \in \mathcal{M}$ be any pseudoanosov mapping class. (We allow ϕ to be orientation reversing.) Let M_ϕ be the mapping torus of ϕ, which is hyperbolic by Thurston's hyperbolisation theorem (see Theorem 1.1). Let $M \cong \Sigma \times \mathbf{R}$ be the cyclic cover which has an action of $\mathbf{Z}$ generated by ϕ. This induces a free action of $\mathbf{Z}$ on $\mathcal{T}$, and we can view the quotient, $\mathcal{T}_\phi = \mathcal{T}/\mathbf{Z}$ as a set of Margulis tubes in $M_\phi = M/\langle \phi \rangle$. (One can show that it accounts for all Margulis tubes in M_ϕ provided η_0 is chosen small enough in relation to $\kappa(\Sigma)$, though we don't need to know that here.)

We say that a map $g : \Sigma \longrightarrow M_\phi$ is *efficient* (resp. straight etc.) if a lift to M is *efficient* (resp. straight etc.).

Lemma 6.12 : *There is* ξ_0*-efficient embedding* $j : \Sigma \longrightarrow M_\phi$ *(with respect to some pants metric,* σ*), where* ξ_0 *depends only on* $\kappa(\Sigma)$.

Proof : Let $j_0 : \Sigma \longrightarrow M_\phi$ be a minimal area embedding in the fibre class. Such an embedding exists, using [MeSY], or [SY] and [FHS]. The pullback metric, σ_0, on Σ has curvature at most -1.

Since Margulis tubes have locally convex boundaries, it follows that $j_0^{-1}(\bigcup \mathcal{T})$ is a disjoint union of discs and pairwise non-homotopic essential annuli. There is a bound on how deeply a disc can enter a Margulis tube, so we can push the discs off the tubes, while maintaining injectivity and a bound on the lipschitz constant with respect to σ_0. This gives us j.

We can now find a pants metric, σ, on Σ, uniformly bilipschitz equivalent to σ_0 on the complement of Margulis tubes, so that j is efficient. Indeed it will be ξ_0-efficient, where ξ_0 depends only on $\kappa(\Sigma)$. $\diamondsuit$

(Note that the argument requires η_0 to be sufficiently small in relation to $\kappa(\Sigma)$.)

We now fix some $h_3 \geq 0$ sufficiently large as will be determined below, and construct a set $\mathcal{C}$ of h_3-collared bands in M as follows. For each $\Phi \in \mathcal{F}$, we choose a h_3-collared band, C, with $H_M(C)$ maximal (or maximal to within an arbitrarily small positive constant), if such exists.

In fact, h_3 will be chosen large enough so that if $C \in \mathcal{C}$, then there are h_2-collars, $C^2_\pm$, of C as well as h_1-collars, $C^1_\pm$, of $C' = C^2_- \cup C \cup C^2_+$, such that $C'' = C^1_- \cup C' \cup C^1_+$ is h_0-collared. Here h_0 is as in Lemma 6.3, $h_1 = h(\xi_0)$ as in Lemma 6.6, where ξ_0 is the constant of Lemma 6.12, and h_2 is as in Lemma 6.8. In particular, any two distinct bands of the form C'' of the same complexity will be disjoint. We can achieve the above using Lemma 6.2, and setting $h_3 = h_2 + h_1 + h_0 + 2H_0$.

We can also assume that we have constructed the set of bands, $\mathcal{C}$, in a $\mathbf{Z}$-invariant manner. Since non-parallel h_0-collared bands of the same complexity do not intersect, we get a free action of $\mathbf{Z}$ on $\mathcal{C}$, and we can identify $\mathcal{C}_\phi = \mathcal{C}/\mathbf{Z}$ with a set of embedded bands on M_ϕ. Again, bands of the same complexity in $\mathcal{C}_\phi$ are disjoint (though bands of different complexities might intersect in quite complicated ways).

Proof of Lemma 6.1 :

Let M_ϕ be the hyperbolic manifold with monodromy ϕ. Let $\mathcal{C}_\phi$ the set of bands constructed above, and let $\mathcal{C}$ be the set of lifts to the infinite cyclic cover, M. Our starting point will be the map $j : (\Sigma, \sigma) \longrightarrow M_\phi$ given by Lemma 6.12. Recall that σ is a pants metric with respect to a multicurve a. Let $\tilde{j} : \Sigma \longrightarrow M$ be the lift of j. We choose a constant h_4, depending only on $\kappa(\Sigma)$ and r_2 to be described below. We split into two cases:

Case (1) : There is some $D \in \mathcal{C}$ with $H_M(D) > h_4$.

In this case, we aim to verify the statement of Case (1) of Lemma 6.1.

The first step is to construct a straight map $g : \Sigma \longrightarrow M_\phi \setminus \mathcal{C}_\phi$ which is homotopic in $M \setminus \mathcal{C}_\phi$ to an embedding $g' : \Sigma \longrightarrow M_\phi \setminus \mathcal{C}_\phi$. Here g is straight with respect to the complete multicurve, a. For each $\alpha \in a$, we want $l_M(\alpha) \leq l_1$, where l_1 depends only on $\kappa(\Sigma)$.

To this end, suppose that $\tilde{j}(\Sigma)$ meets some $C' = C^2_- \cup C \cup C^2_+$, where $C \in \mathcal{C}$. By Lemma 6.6, there is a subsurface, Φ, such that $\tilde{j}|\Phi$ is efficient, and $\tilde{j}|\Phi$ maps into C''. Let $\mathcal{F}'$ be the set of subsurfaces of maximal complexity that arise in this way. Any two elements of $\mathcal{F}'$ will be disjoint (since any two bands of the form C'' of the same complexity are disjoint). Now for each such Φ, we apply Lemma 6.8 to give us a new straight map

g_Φ of Φ, this time into C_+^2. This is homotopic to an embedded fibre, g'_Φ in C_+^2. We can assume that both g_Φ and g'_Φ and the homotopy between them miss all bands of $\mathcal{C}$ of smaller complexity. We can also assume that this construction is carried out equivariantly with respect to the action of $\langle\phi\rangle$ on M.

Let F be a non-annular component of the complement of $\bigcup \mathcal{F}'$ in Σ. By Lemma 6.5 (and the subsequent remark), we can find a $\xi'(\xi_0, l_1)$-efficient homotopy from $\tilde{j}|F$ to a straight map, g_F. Here l_1 is bound depending only on ξ_0 and $\kappa(\Sigma)$, hence ultimately only on $\kappa(\Sigma)$ so that each closed geodesic in M associated to the straight map has length at most l_1. We can assume that h_2 is chosen large enough so that no such homotopy can cross any element of C_2^+, and so the homotopy is disjoint from $\bigcup \mathcal{C}$.

We now assemble the straight maps g_Φ and g_F obtained above to give us a straight map $\tilde{g} : \Sigma \longrightarrow M \setminus \bigcup \mathcal{C}$. Assembling the maps g'_Φ and $\tilde{j}|F$ gives us a map $\tilde{g}' : \Sigma \longrightarrow M \setminus \bigcup \mathcal{C}$, homotopic to $\tilde{g}$ in $M \setminus \bigcup \mathcal{C}$.

We can project back down to M_ϕ to give us maps $g, g' : \Sigma \longrightarrow M_\phi$, which are homotopic in $M_\phi \setminus \bigcup \mathcal{C}_\phi$. Here g is a straight map with respect to a mulicurve a. On the complement of $\bigcup \mathcal{F}'$, the map g' agrees with j and it therefore an embedding. Moreover its image cannot meet any $g'(\Phi)$ since we arranged that $j(F) \cap C = \emptyset$ for all complementary surfaces F.

The monodromy of M may identify different bands on which we have constructed such homotopies. In other words, it is possible that we may have a collection of subsurfaces, $\Phi_1, \Phi_2, \dots, \Phi_n$, with $g|\Phi_i$ and $g'|\Phi_i$ all mapping into the same band, C_+^2, for some $C \in \mathcal{C}_\phi$. Since we arranged the modifications in the bands to be equivariant, the images $g'(\Phi_i)$ are all equal. Now, the original maps, $j|\Phi_i$ are all fibres of C'', and can be assumed to occur in the order $\Phi_1, \dots, \Phi_n$ in C''. We can therefore push the maps $j|\Phi_i$ off each other in a small regular neighbourhood, so that their images are all disjoint and also occur in the same order. Since there were no other parts of $j(\Sigma)$ between $j(\Phi_1)$ and $j(\Phi_n)$ in C'', there is an isotopy of the original fibres, $j(\Phi_i)$, to the new ones in C'' not meeting any other part of $j(\Sigma)$. This means we can join together the various pieces of g' by annuli in Margulis tubes, in such a way that g' is injective in M_ϕ.

In summary, we have achieved our goal of constructing g and g' as required above.

Now write $g_- = \tilde{g}$, $g_+ = \phi g_-$, $g'_- = \tilde{g}'$, $g'_+ = \phi g_+$, $a_- = a$ and $a_+ = \phi a_-$. Let $W \subseteq \Psi$ be the complete annulus system from a_- to a_+, and let $f : P(W) \longrightarrow M$ be the map given by part (1) of Proposition 6.10. In applying Proposition 6.10, we note that $l = l_1$ depends only on κ. The functions h' and h'' then depend only on $\kappa(\Sigma)$ and the argument h. Let v' be the constant of Corollary 5.5. Since the tautness of W depends only on $\kappa(\Sigma)$, so does v'.

Let Q be the region of M between the embedded surfaces $g'_-(\Sigma)$ and $g'_+(\Sigma)$. This is a fundamental domain for M under the $\langle\phi\rangle$-action. Note that $\partial Q \cap \bigcup \mathcal{C} = \emptyset$. Let $\mathcal{C}_0 = \{C \in \mathcal{C} \mid C \subseteq Q\}$. This is an orbit transversal for $\mathcal{C}$ under $\langle\phi\rangle$. Since $g_\pm$ is homotopic to $g'_\pm$ in $M \setminus \bigcup \mathcal{C}$, it follows that $\mathcal{C}_0$ is precisely the set of bands to which f maps with degree 1.

We can now retrospectively make our choice of h_4 to be the constant h'' as given by Proposition 6.10 when $h = r_2 + v'$. This depends only on $\kappa(\Sigma)$ and r_2. We can assume that $C_0 \in \mathcal{C}_0$. Since $H(C_0) \geq h_4 = h''$, it follows that there is some band $B \subseteq \Psi$, with

$H_P(B) \geq r_2+v'$. By Corollary 5.5, we have $p(\Phi; a_-, a_+) = p(\Phi; a, \phi a) \geq (r_2+v')-v' = r_2$, where Φ is the base surface of B.

We similarly set $r_1 = h' + v'$, where h' is the constant arising from part (1) of Proposition 6.10, when $h = 0$. This depends only on $\kappa(\Sigma)$. Suppose that $F \in \mathcal{F}(\Sigma)$ with $p(F; a, \phi a) \geq r_1$ and $p(F; \phi^n a, \phi^{n+1} a) \geq r_1$. Then, by Corollary 5.5 there is a band B in Ψ with base surface F and $H_P(B) \geq (h' + v') - v' = h'$, and so by part (1) of Proposition 6.10, there is a band $C \in \mathcal{C}$ also with base surface F so that $f|f^{-1}C$ maps to C with non-zero degree. In other words $C \in \mathcal{C}_0$. Similarly, since $p(\phi^{-n}F; a, \phi a) \geq r_1$, there is a band $C' \in \mathcal{C}$ with base surface $\phi^{-n}F$, with $f|f^{-1}C'$ mapping to C with non-zero degree. Thus $C' \in \mathcal{C}_0$. By construction, $\mathcal{C}$ is invariant under the covering transformation, ϕ, and distinct elements of $\mathcal{C}$ have distinct base surfaces. We must therefore have $C = \phi^n C'$. It follows that $\mathcal{C}_0 \cap \phi^n \mathcal{C}_0 \neq \emptyset$. Since $\mathcal{C}_0$ is an orbit transversal of the covering action, we have $n = 0$.

Finally note that for each $\alpha \in a$, we have $l_M(\alpha) \leq l_1$ and so, by Lemma 3.5, a is a bounded distance from an axis of ϕ on $\mathcal{G}(\Sigma)$. We write s_1 for this bound. This again depends only on $\kappa(\Sigma)$.

We have thus verified that Case (1) of Lemma 6.1 holds.

Case (2). For all $C \in \mathcal{C}$ we have with $H_M(C) \leq h_4$.

This time we verify the statement of Case (2) of Lemma 6.1.

Applyng Lemma 6.2, it follows from the construction of $\mathcal{C}$ that any band in M has height at most $h_4 + 2h_3 + 2H_0$, which depends only on $\kappa(\Sigma)$ and r_2. (Recall that $h_3 = h_2 + h_1 + h_0 + 2H_0$.) By Lemma 6.4, there is some $L \geq 0$ depending only on $\kappa(\Sigma)$ and r_2 such that $L_M(\partial T) \leq L$ for all $T \in \mathcal{T}$.

We again start with our ξ_0-efficient embedding $j : \Sigma \longrightarrow M_\phi$. Note that $l_M(\alpha) \leq l_0$ for all $\alpha \in a$. We can homotope this in M_ϕ to another embedding $g' : \Sigma \longrightarrow M_\phi \setminus \bigcup \mathcal{T}_\phi$, by pushing it off each Margulis tube. Note that if $T \in \mathcal{T}_\phi$, then $j(\Sigma) \cap T$ consists of a number of disjoint annuli. The lengths of the boundary curves in ∂T are all bounded. We can therefore maintain a bound on the lipschitz constant for g' depending only on L and $\kappa(\Sigma)$, and hence, ultimately, only on r_2 and $\kappa(\Sigma)$.

The structure on the domain, Σ, remains unchanged. Namely, it is a pants structure based on the complete multicurve, a. We extend this to a complete marking, $b \supseteq a$, of bounded length in σ. Since $l_M(\alpha) \leq \xi_0$ for all $\alpha \in a$, Lemma 6.5 gives us a ξ_2-efficient homotopy from g' to a ξ_0-straight map $g : \Sigma \longrightarrow M$. Here, ξ_2 depends on the lipschitz constant of g', hence on r_2 as well as $\kappa(\Sigma)$.

We now lift to M giving us maps $g_- = \tilde{g}$, $g_+ = \phi g_-$, $g'_- = \tilde{g}'$, $g'_+ = \phi g'_-$ as in Case (1). We write $a_- = a$ and $a_+ = \phi a_-$. Let $W \subseteq \Psi$ be the complete annulus system from a_- to a_+, and let $f : P(W) \longrightarrow M$ be the map given by Proposition 6.10. Again, as in Case (1), the constant $l = l_1$ depends only on $\kappa(\Sigma)$. The functions h' and h'' then depend only on $\kappa(\Sigma)$ and the argument h.

By assumption, every band in M has height at most h_4 so by Proposition 6.10, part (1), every band, B in Ψ has height at most $h'(h_4, l_1, \kappa)$. This, therefore places a bound on $p(\Phi; a, \phi a) = p(\Phi; a_-, a_+)$, and so therefore also on $p(\Phi; b, \phi b)$ (since these are all defined, and each curve of b has bounded intersection with a). We denote this bound by r_4. It only depends on $\kappa(\Sigma)$ and h_4, and hence on $\kappa(\Sigma)$ and r_2. Note that this also places a bound on $L_P(\Omega)$ for each annulus, Ω, of W.

Now let $\hat{P}(W)$ be the augmented model, constructed using markings $b_+ \supseteq a_+$ and $b_- \supseteq a_-$. We extend $f : P(W) \longrightarrow M$ given above to a map $f : \hat{P}(W) \longrightarrow M$, with $f|\partial_\pm \Psi = g'_\pm$ as follows. First note that $f|\partial_\pm \Psi$ is straight. Identifying $\Sigma = \partial_+ \Psi$, it is ξ_0-straight with respect to a pants metric on Ψ based on a where ξ_0 depends only on $\kappa(\Sigma)$. By the second part of Lemma 6.5, we can assume that it agrees with $g_+ : \Sigma \longrightarrow M$, with the same pants metric. We now use the homotopy from g_+ to g'_+ given by Lemma 6.12 to extend over Ψ_+. This homotopy is ξ_3-efficient, where $\xi_3 = \xi'(\xi_2)$ depends only on $\kappa(\Sigma)$ and r_2. To do this, we need to push W slightly into Ψ_+. More precisely, extend the annuli of W with base curves in c by adding in $\bigcup c \cup [0, \frac{1}{2}] \subseteq \Psi_+ = \Sigma \times [0,1]$. We also do a similar construction for Ψ_-, so as to give us our map f on $\hat{P}(W)$. This map is ξ_3-efficient.

Let $\mathcal{T}_0 \subseteq \mathcal{T}$ be the set of η_0-Margulis tubes lying between the disjoint embedded surfaces $g'_-(\Sigma)$ and $g'_+(\Sigma)$. Thus $\mathcal{T}_0$ is an orbit transversal of $\mathcal{T}$ under the covering action. Since the homotopy from $g_\pm$ to $g'_\pm$ is ξ_3-efficient, it cannot cross (i.e. map with non-zero degree to) any η_1-Margulis tube, where $\eta_1 \in (0, \eta_0)$ is a constant only depending on ξ_3 hence on $\kappa(\Sigma)$ and r_2. Let $\mathcal{T}' \subseteq \mathcal{T}$ be the set of tubes whose core curves have length at most η_1. The set $\mathcal{T}'_0 = \mathcal{T}_0 \cap \mathcal{T}'$ is then an orbit transversal of $\mathcal{T}'$. It is the set of tubes of $\mathcal{T}'$ to which f from $\hat{P}(W)$ maps with degree one.

Suppose that $\gamma \in X(\Sigma)$ with $p(\gamma; b, \phi b) \geq r_3$, where r_3 will be chosen sufficiently large below. Suppose that $p(\gamma; b, \phi b) \geq r_3$. By Lemma 5.3 and Corollary 5.5, there is an annulus $\Omega(\gamma)$ in $\hat{P}(W)$ with $H_P(\gamma) \geq r_3 - v_1$, where $v_1 = \max\{v, v'\}$ depends only on κ. Let the modulus of $\Delta = \partial T(\gamma)$ be $L + iH$. Thus, L is bounded above, and $|H - H_P(\gamma)| \leq \frac{1}{2}$. Now f maps Δ either to a closed geodesic γ_M, whose length is bounded below by the Margulis constant, η_0, or else to the boundary of some η_0-Margulis tube, $T = T(\gamma) \in \mathcal{T}$. Consider the first case. After rescaling the geodesic by a factor at most $1/\eta_0$, so as to have unit length, we are in the situation described by Lemma 6.11, with lipschitz constant ξ_3/η_0, and so we get $|H| \leq \xi_3 L/\eta_0$, which is bounded. So if we choose r_3 sufficiently large, depending on $\kappa(\Sigma)$ and r_2, we can eliminate this case. We therefore have a ξ_3-lipschitz map from Δ to ∂T. Suppose $T \notin \mathcal{T}'$, that is the core curve has length at least η_1. Postcomposing with the nearsest point projection, we get a ξ_3-lipschitz map to the core curve, and then rescaling we get a ξ_3/η_1-lipschitz map to the unit-length circle. Again, choosing r_3 big enough (this time in relation to r_2 as well as $\kappa(\Sigma)$), we get a contradiction. Thus we can suppose that $T \in \mathcal{T}'$. Since f maps with degree 1 to T, this means that $T \in \mathcal{T}'_0$.

Suppose also that $p(\gamma; \phi^n b, \phi^{n+1} b) = p(\phi^{-n}\gamma; b, \phi b) \geq r_3$. Then we also have $T(\phi^{-n}\gamma) \in \mathcal{T}_0$, so $T = T(\gamma) \in \phi^n \mathcal{T}'_0$. Thus $\mathcal{T}'_0 \cap \phi^n \mathcal{T}'_0 \neq \emptyset$. Since this is an orbit transversal, it follows that $n = 0$.

As in Case (1), we note that each $\alpha \in a$, $l_M(\alpha) \leq l_1$ and so, by Lemma 3.5, a is a bounded distance from an axis of ϕ on $\mathcal{G}(\Sigma)$. We write s_1 for this bound. This again depends only on $\kappa(\Sigma)$.

We have verified the statement of Case (2), finally proving Lemma 6.1. ♢

7. Proof of the main theorem.

We shall give a proof of Theorem 0. We begin by describing a fairly general construction. This will be a formal presentation of the argument involving "bent" edges as outlined in Section 2. It is in a form that can be generalised further, as we discuss in Section 8.

Suppose that χ is a (topological) 1-vertex triangulation of the base surface, S. We write $E(\chi)$ for the set of edges of χ. We write Υ for the 3-valent graph dual to χ — so $V(\Upsilon)$ is identified with the set of triangles of χ, and we can identify $E(\Upsilon)$ with $E(\chi)$. We can lift this picture to get a triangulation, $\tilde{\chi}$, of the universal cover, $\tilde{S}$, with dual graph, $\tilde{\Upsilon}$.

Suppose that Ξ is another finite graph such that every vertex has degree 2 or 3, and such that $\theta : \Xi \longrightarrow \Upsilon$ is a locally injective map of graphs. If $\Xi \neq \emptyset$, it must contain an embedded circuit, ω. We also make the following observation:

Lemma 7.1 : *Suppose that ϵ is some edge of χ, and ω is a circuit in Ξ. If $\theta(\omega)$ is homotopically trivial in S, then any lift of $\theta(\omega)$ to $\tilde{\Upsilon} \subseteq \tilde{S}$ crosses some lift, $\tilde{\epsilon}$, of ϵ in at least two distinct edges.*

Proof : Note that the subgraph of $\tilde{\Upsilon}$ consisting of those edges that do not cross any lift of ϵ is a forest. But the lift of $\theta(\omega)$ is an embedded circuit, and so the result follows easily. $\diamondsuit$

Here is an abstract set-up in which one can construct such a graph.

Let X be a set on which $\Gamma = \pi_1(S)$ acts. (In practice this will be a subset of $X(\Sigma) \cup \mathcal{F}(\Sigma)$.) Suppose that to each edge $e \in E(\tilde{\chi})$ we have associated a finite subset $Y(\epsilon) \subseteq X$ in a Γ-equivariant fashion (i.e. $Y(g\epsilon) = gY(\epsilon)$ for all $g \in \Gamma$). Let us suppose, for any such component, $\tilde{\epsilon}_0$, of the lift, ϵ_0, we have:

(Y1) If $\{\epsilon, \epsilon', \epsilon''\} \subseteq E(\tilde{\chi})$ is a triangle in $\tilde{\chi}$, then $Y(\epsilon) \subseteq Y(\epsilon') \cup Y(\epsilon'')$.

We construct a graph, $\tilde{\Xi}$, as follows. For each triangle, $\tau = \{\epsilon, \epsilon', \epsilon''\}$ in $\tilde{\chi}$, we set $Y_\tau = Y(\epsilon) \cup Y(\epsilon') \cup Y(\epsilon'')$. We set $V(\tilde{\Xi})$ to be the disjoint union of the Y_τ as τ ranges over all triangles. We denote an element of Y_τ by (x, τ), where $x \in Y(\epsilon) \cup Y(\epsilon') \cup Y(\epsilon'')$. We deem two elements (x, τ) and (y, τ') of $V(\tilde{\Xi})$ to be adjacent if τ and τ' are distinct triangles meeting along $\epsilon \in E(\tilde{\chi})$ and if $x = y \in Y(\epsilon)$. This defines a graph, $\tilde{\Xi}$ with a natural locally injective map $\tilde{\theta} : \tilde{\Xi} \longrightarrow \tilde{\Upsilon}$. Note that this is $|Y(\epsilon)|$-to-1 along each edge ϵ. There is also a natural map $\tilde{\Xi} \longrightarrow X$, which sends any component of $\tilde{\Xi}$ to a point. Since the construction was Γ-equivariant, it projects to a map $\theta : \Xi \longrightarrow \Upsilon$. Note that (Y1) implies that every vertex of Ξ has degree 2 or 3.

Suppose that ϵ_0 is some preferred edge of $\tilde{\chi}$, and let $\tilde{\epsilon}_0$ be a component of the lift of ϵ_0 to $\tilde{\Xi}$ consisting of a bi-infinite arc. This is translated by some $g \in \Gamma$ (corresponding to the element of $\pi_1(S) \cong \Gamma$ given by the loop ϵ_0). Let us suppose that:

(Y2): If δ, δ' are edges of $\tilde{\epsilon}_0$, with $Y(\delta) \cap Y(\delta') \neq \emptyset$, then $\delta = \delta'$.

This can be rephrased by saying that, for some (or indeed any) edge, δ, of ϵ_0, $Y(\delta) \cap g^n Y(\delta) \neq \emptyset$ implies $n = 0$. Now, using Lemma 7.1, we see that $\tilde{\Xi}$ is a forest. In other words, if ω is any circuit in Ξ, then $\theta(\omega)$ is non-trivial in S. Using the fact that $Y(\epsilon)$ is finite for all edges ϵ, we see that the some power of monodromy around such a loop fixes a point of X (corresponding to a component of $\tilde{\Xi}$). For the existence of such a loop, we need that $\Xi \neq \emptyset$, i.e.:

(Y3) For some $\epsilon \in E(\tilde{\chi})$, $Y(\epsilon) \neq \emptyset$.

To see how this might be applied, suppose that $\Sigma \longrightarrow \Pi \longrightarrow S$ is a surface bundle, and that $X = X(\Sigma)$. We get an action of $\Gamma = \pi_1(\Sigma)$ on X via the monodromy homomorphism $\phi : \Gamma \longrightarrow \mathcal{M}$. Suppose that Y is equivariant and satisfies (Y1), (Y2) and (Y3) above. Then any circuit, ω, in Ξ gives us an essential loop, $\theta(\omega)$ in S (not necessarily embedded), and some $\alpha \in X(\Sigma)$ that is fixed by some power of the monodromy around $\theta(\omega)$. This in turn gives us an essential torus in Π.

Essentially the same argument applies if we replace X by $\mathcal{F}_k = \{\Phi \in \mathcal{F}(\Sigma) \mid \kappa(\Phi) = k\}$, where $1 \leq k \leq \kappa(\Sigma) - 1$. This time, we get an invariant subsurface. We may have to run around $\theta(\omega)$ several times to get an essential torus in Π.

Here is another observation needed for our construction.

Lemma 7.2 : *Suppose that $r > 0$, and that Q is a set of at most n elements. Suppose that we have a map $\psi : Q \longrightarrow [0, \infty)$ and that $\psi(Q) \cap [q, \infty) \neq \emptyset$, where $q > 3^n r$. Then there is some non-empty $Z \subseteq Q$ with $\min \psi(Z) > 3 \max(\psi(Q \setminus Z) \cup \{r\})$.*

Proof : Indeed we can define Z canonically by starting with $y \in Q$ such that $\psi(y)$ is maximal, and work downwards until we find the first successive ratio bigger than 3. $\diamondsuit$

Note that this also works if Q is a set with a Γ-action, with $|Q/\Gamma| \leq n$, and $\psi : Q \longrightarrow [0, \infty)$ is Γ-invariant. Since the construction is canonical, we get a Γ-invariant subset $Z \subseteq Q$, with $\min \psi(Z) > 3 \max(\psi(Q \setminus Z) \cup \{r\})$.

To apply this, suppose we have a map $\psi : X \times E(\tilde{\chi}) \longrightarrow [0, \infty)$ satisfying $\psi(gx, g\epsilon) = \psi(x, \epsilon)$ for all $x \in X$ and all $g \in \Gamma$. Let us assume that:

(Q1) There is some $r \geq 0$ such that if $\{\epsilon, \epsilon', \epsilon''\} \subseteq E(\tilde{\chi})$ is a triangle in $\tilde{\chi}$ and $x \in X$ with $\psi(x, \epsilon) \geq 3r$, then either $\psi(x, \epsilon') \geq \frac{1}{3}\psi(x, \epsilon)$ or $\psi(x, \epsilon'') \geq \frac{1}{3}\psi(x, \epsilon)$ (or both).

Given $\epsilon \in E(\tilde{\chi})$, we assume $\{x \in X \mid \psi(x, \epsilon) \geq r\}$ to be finite. Let $Q(\epsilon) = \{x \in X \mid \psi(x, \epsilon) \geq r\} \times \{\epsilon\}$, and let et $Q = \bigcup_{E(\tilde{\chi})} Q(\epsilon) \subseteq X \times E(\tilde{\chi})$. Set $n = |Q/\Gamma| and q = 3^{n+1}r$. Note that there is a Γ-invariant map $\psi : Q \longrightarrow [0, \infty)$. If $\psi(Q) \cap [q, \infty) \neq \emptyset$, Lemma 7.2 and the subsequent remark gives us a non-empty invariant subset $Z \subseteq Q$. Given $\epsilon \in E(\tilde{\chi})$, we write $Y(\epsilon)$ for the set of those $x \in X$ that correspond to some element of Z. Thus Y is Γ-equivariant. From the construction of Z, we have $\min\{\psi(x, \epsilon) \mid \epsilon \in E(\tilde{\chi}), x \in Y(\epsilon)\} > 3 \max(\{\psi(y, \delta) \mid \delta \in E(\tilde{\chi}), y \in X \setminus Y(\delta)\} \cup \{r\})$.

Now, property (Q1) tells us that (Y1) holds, so we can construct Ξ and a map $\theta : \Xi \longrightarrow \Upsilon$ as before.

To translate (Y2) into these terms, we suppose:

(Q2) If $\delta \in E(\tilde{\chi})$ is a lift of the preferred edge ϵ_0 and g the corresponding element of Γ, then if $x \in X$ with $\psi(x, \delta) \geq r$ and $\psi(g^n x, \delta) \geq r$, then $n = 0$.

Finally, for (Y3), we need:

(Q3) There is some $\epsilon \in E(\tilde{\chi})$ and $x \in X$ with $\psi(x, \epsilon) \geq q$.

We can now formally begin the proof of Theorem 0.

Let us suppose that $\Sigma \longrightarrow \Pi \longrightarrow S$ is atoroidal. By Proposition 3.4, there is standard system of generators, $g_1, \ldots, g_n$, for $\Gamma \cong \pi_1(S)$, and some $\alpha \in X(\Sigma)$ such that for all i, $d_{\mathcal{G}(\Sigma)}(\alpha, g_i\alpha) \leq s$, where $s = s(\kappa(\Sigma))$.

Now a standard generating set corresponds to a collection of loops in S, all meeting at the basepoint, but with their interiors disjoint. Their complement is a $2n$-gon, and by adding diagonals, we obtain a 1-vertex triangulation, χ, of S. We lift this to $\tilde{S}$, and choose some lift of the basepoint, $x_0 \in V(\tilde{\chi})$. Note that this determines an identification of $\Gamma = \pi_1(\Sigma)$ with the covering group. We see that α is moved some bounded distance in $\mathcal{G}$ by each of the monodromies of the edges of $\tilde{\chi}$ incident on x_0.

Let ϵ_0 be any edge of χ, which we refer to as our "preferred edge". Let δ_0 be the edge of $\tilde{\chi}$ with vertex at x_0 that projects to ϵ_0. Let $\phi_0 \in \mathcal{M}$ be the corresponding monodromy.

We now apply Lemma 6.1. Let $r_1 = r_1(\kappa(\Sigma))$ and let r_2 be as determined below. To reduce us to Case (2), we first suppose we are in Case (1), and then derive a contradiction, assuming we have chosen r_2 big enough. (This will only depend ultimately on $\kappa(\Sigma)$ and $\kappa(S)$.)

In other words, we are assuming that there is a multicurve, a, in Σ, within distance $s_1 = s_1(\kappa(\Sigma))$ of an axis of ϕ_0 in $\mathcal{G}(\Sigma)$, and with the following properties. There is some $\Phi \in \mathcal{F}$ with $p(\Phi; a, \phi_0 a) \geq r_2$. Moreover, if $F \in \mathcal{F}$, with $p(F; a, \phi_0 a) \geq r_1$ and $p(F; \phi_0^n a, \phi_0^{n+1} a) \geq r_1$, then $n = 0$.

Now since α is translated a bounded distance, it is a simple exercise in hyperbolic spaces to show that α is a bounded distance from $\phi_0^m a$ for some $m \in \mathbf{Z}$. Now $\phi_0^m a$ satisfies the same conditions as a, and so, without loss of generality, we can assume that $d(\alpha, a)$ is bounded above by some constant depending only on $\kappa(\Sigma)$. From this we get:

Lemma 7.3 : *There is some $s_0 = s_0(\kappa(\Sigma))$ such that for each edge, $\epsilon \in E(\tilde{\chi})$ incident on x_0, we have $d_{\mathcal{G}(\Sigma)}(a, ga) \leq s_0$, where $g \in \Gamma$ corresponds to ϵ.* $\diamondsuit$

Now χ has only one vertex, so if $x \in V(\tilde{\chi})$, there is a unique $k \in \Gamma$ with $x = kx_0$. We can therefore set $a(x) = ka$. In other words, we associate to each $x \in V(\tilde{\chi})$ a multicurve, $a(x)$ in Σ, in a Γ-equivariant fashion.

Given $k \in \{1, \ldots, \kappa(\Sigma)\}$, write $\mathcal{F}_k = \{\Phi \in \mathcal{F} \mid \kappa(\Phi) = k\}$. Given $\Phi \in \mathcal{F}$ and $\epsilon \in E(\tilde{\chi})$, write $\psi(\Phi, \epsilon) = p(\Phi; a(x), a(y))$, where $x, y \in V(\tilde{\chi})$ are the endpoints of ϵ. Note that $\psi(\Sigma, \epsilon) = p(\Sigma; a(x), a(y)) = d_{\mathcal{G}(\Sigma)}(a(x), a(y)) = d_{\mathcal{G}(\Sigma)}(a, \phi a)$, where $\phi \in \mathcal{M}$ is the monodromy corresponding to the projection of the edge ϵ. Lemma 7.3 tells us that this is bounded, so rephrasing it, we see:

Lemma 7.4 : *For all $\epsilon \in E(\tilde{\chi})$, $\Psi(\Sigma, \epsilon) \leq s_0$, where s_0 depends only on $\kappa(\Sigma)$.* $\diamondsuit$

We can also rephrase Proposition 4.6 in these terms:

Lemma 7.5 : *There is a constant, $v_0 = v_0(\kappa(\Sigma))$ and a function, $N = N_\kappa$, depending only on $\kappa = \kappa(\Sigma)$, with the following property. Given $\epsilon \in E(\tilde{\chi})$ and $k \geq 1$, suppose that for all $F \in \mathcal{F}$ with $\kappa(F) > k$ we have $\psi(F,\epsilon) \leq v$, then there are at most $N(v)$ surfaces $\Phi \in \mathcal{F}_k$ with $\psi(\Phi,\epsilon) \geq v_0$.* ♢

Now we can assume (to simplify notation) that $r_1 \geq v_0$, and we let $q_0 = \max\{9, s_0, r_1\}$. Recall that v_0 is obtained from Lemma 7.5, s_0 from Lemma 7.3 and r_1 from Lemma 6.1. Each of these depends only on $\kappa = \kappa(\Sigma)$. Let $e = |E(\chi)|$, which depends only on $\kappa(S)$. We define constants, $q_0 < q_1 < q_2 < \cdots < q_{\kappa-1}$ and $n_1 < n_2 < \cdots < n_{\kappa-1}$ inductively as follows. Set

$$\begin{aligned} n_1 &= eN(q_0), & q_1 &= 3^{n_1+1}r_1 \\ n_2 &= eN(q_1), & q_2 &= 3^{n_2+1}r_1 \\ n_3 &= eN(q_2), & q_3 &= 3^{n_3+1}r_1 \\ &\vdots & &\vdots \\ n_{\kappa-1} &= eN(q_{\kappa-1}), & q_{\kappa-1} &= 3^{n_{\kappa-1}+1}r_1. \end{aligned}$$

Thus, for each k, the constants q_k and n_k depend only on $\kappa(\Sigma)$ and $\kappa(S)$. (We can assume that $N(v) \geq v$ for all v, so that n_i and q_i are indeed increasing in i.)

Now, by Lemma 7.4, we know that $\psi(\Sigma,\epsilon) \leq q_0$ for all $\epsilon \in E(\tilde{\chi})$. But now Lemma 7.5 tells us that for each $\epsilon \in E(\tilde{\chi})$, there are at most $N(q_0)$ surfaces, $\Phi \in \mathcal{F}_{\kappa-1}$ with $\psi(\Phi,\epsilon) \geq r_1$. Call this (possibly empty) set $Q(\epsilon)$, and let Q be the disjoint union of the sets $Q(\epsilon)$ as ϵ ranges over $E(\tilde{\chi})$. Note that $|Q/\Gamma| \leq eN(q_0) = n_1$.

We now use this as the basis for the construction of a graph, Ξ, as described above, with $r = r_1$ and with $X = \mathcal{F}_{\kappa-1}$. Property (Q1) follows from the coarse triangle inequality $\psi(\Phi,\epsilon) \leq \psi(\Phi,\epsilon') + \psi(\Phi,\epsilon'') + 3$, given that $q_0 \geq 9$. To verify (Q2), suppose that δ is any lift of ϵ_0, with initial vertex $x \in V(\tilde{\chi})$, and with $g \in \Gamma$ the corresponding element. Case (1) of Lemma 6.1 translates as saying that if $\psi(\Phi,\delta) \geq r_1$ and $\psi(\Phi,g^n\delta) \geq r_1$, then $n = 0$. By Γ-equivariance, this now holds for every lift of ϵ_0. This verifies (Q2). If $\Xi \neq \emptyset$, then we get the contradiction that Π is not atoroidal. Thus, we can assume that $\Xi = \emptyset$, in other words (Q3) fails. This means that $\psi(\Phi,\epsilon) \leq q_1 = 3^{n_1+1}r_1$ for all $\epsilon \in E(\tilde{\chi})$. Since $N(q) \geq q$, it follows that $q_1 \geq q_0$.

Now we apply the same argument with $\mathcal{F}_{\kappa-2}$ replacing $\mathcal{F}_{\kappa-1}$. For any $\epsilon \in E(\tilde{\chi})$, there are at most $N(q_1)$ subsurfaces, Φ with $\kappa(\Phi) \geq \kappa - 1$, and with $\psi(\Phi,\epsilon) \geq r_1$. Thus, with $X = \mathcal{F}_{\kappa-2}$, we get $|Q/\Gamma| \leq eN(q_1) = n_2$. We conclude that for all $\Phi \in \mathcal{F}_{\kappa-2}$ and all $\epsilon \in E(\tilde{\chi})$, we have $\psi(\Phi,\epsilon) \leq q_2 = 3^{n_2+1}r_1$.

We proceed inductively to show that for all $k \in \{1,\ldots,\kappa-1\}$, we have $\psi(\Phi,\epsilon) \leq q_k$ for all $\Phi \in \mathcal{F}_{\kappa-k}$ and all $\epsilon \in E(\tilde{\chi})$. We have shown:

Lemma 7.6 : *If Π is atoroidal, then for all $\Phi \in \mathcal{F}$, and all $\epsilon \in E(\tilde{\chi})$, we have $\psi(\Phi,\epsilon) \leq q_{\kappa-1}$.* ♢

But now, interpreting the first statement of Case (1) of Lemma 6.1, we see that there is some $F \in \mathcal{F}$ with $\psi(F,\delta) \geq r_2$, where δ is a lift of our preferred edge, ϵ_0, and where r_2 is a constant at least r_1, that we are otherwise free to choose. Thus, if we set r_2 bigger than $q_{\kappa-1}$ we get a contradiction. In other words, there is some $r_2 = r_2(\kappa(\Sigma), \kappa(S))$ such that Case (1) cannot arise if Π is atoroidal.

We are therefore in Case (2). In other words, there is a marking, b, of Σ within distance s_0 of an axis of $\phi_0 \in \mathcal{M}$, in $\mathcal{G}(\Sigma)$ (recalling that ϕ_0 is the monodromy associated with ϵ_0) and with the following properties. If $F \in \mathcal{F}$, we have $p(F; b, \phi_0 b) \leq r_4$. Moreover, if $\gamma \in X(\Sigma)$ and $p(\gamma; b, \phi_0 b) \geq r_3$ and $p(\gamma; \phi_0^n b, \phi_0^{n+1} b) \geq r_3$, then $n = 0$. Note that r_3 and r_4, depended only on $\kappa(\Sigma)$ and r_2, and hence now, only on $\kappa(\Sigma)$ and $\kappa(S)$.

We may assume that $r_3 \geq \max\{r_4, v_0\}$ so that the second statement with "$\gamma \in X(\Sigma)$" replaced by "$F \in \mathcal{F}$" becomes vacuously true. This is all we need for the subsequent argument. Note that, as in the previous case, we may assume that $d_{\mathcal{G}(\Sigma)}(b, \alpha)$ is bounded, and so $p(\Sigma; b, gb) \leq s_0 = s_0(\kappa(\Sigma))$ for all g corresponding to those edges of $E(\tilde{\chi})$ adjacent to the basepoint x_0.

We now proceed as in Case (1), defining an equivariant map $[x \mapsto b(x)]$ that associates a marking, $b(x)$, of Σ to each $x \in V(\tilde{\chi})$. For this, we start with $b(x_0) = b$. We now define $\psi(\Phi, \epsilon) = p(\Phi; b(x), b(y))$, where x, y are the endpoints of ϵ. We have $\psi(\Sigma, \epsilon) \leq s_0$ for all $\epsilon \in E(\tilde{\chi})$. In other words, Lemma 7.4 holds, with $b(x)$ replacing $a(x)$. Since our choice of b depended on r_2, hence on $\kappa(S)$, we need to add dependence on $\kappa(S)$ as well as on $\kappa(\Sigma)$. Property (Q2) holds, with $r \geq r_3$. Lemma 7.5 remains valid (with the same constant v_0).

We now set $q'_0 = \max\{9, s_0, r_3\}$. We define

$$\begin{aligned}
n'_1 &= eN(q'_0), & q'_1 &= 3^{n'_1+1} r_3 \\
n'_2 &= eN(q'_1), & q'_2 &= 3^{n'_2+1} r_3 \\
n'_3 &= eN(q'_2), & q'_3 &= 3^{n'_3+1} r_3 \\
&\ \vdots & &\ \vdots \\
n'_{\kappa-1} &= eN(q'_{\kappa-1}), & q'_{\kappa-1} &= 3^{n'_{\kappa-1}+1} r_3.
\end{aligned}$$

We now repeat the same argument as before (using $r = r_3$ instead of $r = r_1$) and conclude that for all $\Phi \in \mathcal{F}$ and all $\epsilon \in E(\tilde{\chi})$, we have $\psi(\Phi, \epsilon) \leq q'_{\kappa-1}$.

This is the same as Lemma 7.6 except that we have now replaced the multicurve, a, by a marking, b. This allows us to pass to the final step, namely to bound the twisting of curves. In summary, we set $n = eN(q'_{\kappa-1})$ and $q = 3^{n+1} r_3$. We deduce that for all $\Phi \in \mathcal{F}$, all $\gamma \in X(\Sigma)$, and all $\epsilon \in E(\tilde{\chi})$, we have $\psi(\Phi, \epsilon) \leq q$ and $\psi(\gamma, \epsilon) \leq q$.

Let us translate this information back to our standard generating set $g_1, \ldots, g_n$, of Γ which we used for the construction of χ. They correspond to certain edges of $E(\tilde{\chi})$ incident on the basepoint $x_0 \in V(\tilde{\chi})$. If $\phi_i \in \mathcal{M}$ is the monodromy associated to g_i, then for all i we have $p(\Phi; b, \phi_i b) \leq q$ for all $\Phi \in X(\Sigma) \cup \mathcal{F}(\Sigma)$. But now Lemma 4.5 gives us a finite subset $\mathcal{M}_0$, depending only on q, hence ultimately only on $\kappa(\Sigma)$ and $\kappa(S)$, such that $\phi_i \in \mathcal{M}_0$ for each i.

We have now verified the hypotheses of Lemma 1.6, where $\mathcal{P}$ is the set of all atoroidal bundles.

This proves Theorem 0.

8. A generalisation to one-ended groups

As observed in the introduction, the main argument of this paper goes through if Γ is a finitely presented one-ended group — though we have to make a somewhat unnatural assumption — essentially the conclusion of Proposition 3.4. This assumption is justified in [DaF] and in [Bow5], where versions of Proposition 3.4 are proven. Thus Theorem 0 can be generalised to this situation.

To give a more precise statement of what we can show, let us suppose that Γ is finitely presented and that $A \subseteq \Gamma$ is a finite generating set. We define the *complexity*, $c(A)$, of A as the minimal sum of the lengths of a set of relators with this generating set. Note that we can construct a one-vertex simplicial complex S with $\Gamma = \pi_1(S)$, with the number of 2-simplices, bounded in terms of $c(A)$, and so that each element of A corresponds to some edge of S. (Formally one can define a "one-vertex simplicial complex" in terms of its universal cover together with the action of Γ by covering translation.)

Proposition 8.1 : *Suppose that Γ is a one-ended finitely presented group, and that $\phi : \Gamma \longrightarrow \mathcal{M} = \mathcal{M}(\Sigma)$ is a purely pseudoanosov homomorphism, giving an induced action on $\mathcal{G}(\Sigma)$. Suppose that $A \subseteq \Gamma$ is a generating set and that there is some $\alpha \in X(\Sigma)$ and $l \geq 0$ such that $d(\alpha, g\alpha) \leq l$ for all $g \in A$. Then there is some $\theta \in \mathcal{M}$ such that the word length of each $\theta\phi(g)\theta^{-1}$ in $\mathcal{M}$ is bounded above in terms of $c(A)$ and l for all $g \in A$.*

Here the word length is, of course, measured with respect to some fixed set of generators of $\mathcal{M}$. Note that the hypotheses imply that Γ is torsion-free.

The argument is essentially already included in the paper, except that now S need no longer be a surface. This only really featured in Section 7, where a few comments are in order.

We need to reintepret the statements about the graphs Υ and Ξ. We let $\Upsilon \subseteq S$ be the dual to the triangulation of S, more precisely, it is the union of those 1-cells in the first barycentric subdivision of S that do not meet the vertex. Note that now Υ is bipartite with one set of vertices at the midpoint of edges, and other at the centres of triangles. (It is a subdivision of the dual graph we decribed earlier in the case of a surface.) We also arbitrarily choose some edge, ϵ_0, of S which represents a non-trivial element of $\pi_1(S) = \Gamma$. Note that each component of the lift of ϵ_0 is a bi-infinite path.

Suppose now that Ξ is a non-empty finite graph and that $\theta : \Xi \longrightarrow \Upsilon$ is a locally injective map that is a local homeomorphism near those vertices that get mapped to midpoints of edges of S. Suppose that all other vertices of Ξ (those sent to centres of triangles) have degree at least 2, and suppose that the image of $\pi_1(\Xi)$ in $\pi_1(S)$ trivial. Thus we get a lift $\tilde{\theta} : \Xi \longrightarrow \tilde{\Upsilon}$, where $\tilde{\Upsilon} \subseteq \tilde{S}$ is the preimage of Υ in the universal cover, $\tilde{S}$, of S. Let $\Xi_0 = \tilde{\theta}(\Xi) \subseteq \tilde{\Upsilon}$. We claim that Ξ_0 meets some bi-infinite lift of ϵ_0 in at least 2 distinct edges. (Note that Ξ_0 track in the sense of [Du], though generalised to allow for branching at the centres of 2-simplices.)

To see this, we first note that Ξ_0 must separate $\tilde{S}$. In fact, if N is a small regular neighbourhood of Ξ_0 in $\tilde{S}$, then ∂N cannot be connected. If it were, the image of $\pi_1(\partial N)$ in $\pi_1(N)$ would have to be a proper subgroup, since these groups are both free. (It is here we use the fact that each vertex of Ξ_0 in the centre of a triangle has degree at least 2, so that the map of $\pi_1(\partial N)$ into $\pi_1(N)$ is injective.) By van-Kampen's theorem, we would get the contradiction $\pi_1(S)$ is non-trivial. Now again, by van-Kampen's theorem, different components of ∂N lie in different components of $\tilde{S} \setminus \mathrm{int}\, N$. Thus, Ξ_0 separates $\tilde{S}$. Since $\tilde{S}$ is one-ended, there must be some bounded component of the complement. This must contain a vertex of $\tilde{S}$, and we see that the lift of ϵ_0 passing through this vertex must meet Ξ_0 in at least two distinct edges, thereby proving the claim.

These observations are enough to make the previous argument of Section 7 go through. The hypotheses of Proposition 8.1 give us our complex S. If there were no bound on word length, we would, at some point, be able to construct a non-empty finite graph, Ξ, and a map $\theta : \Xi \longrightarrow \Upsilon$, as a above. The purely pseudoanosov hypothesis means that the image of $\pi_1(\Xi)$ in S must be trivial. We proceed to get a contradiction as before. This then proves Proposition 8.1.

References.

[Ba] J.Barnard, *Bounding surface actions on hyperbolic spaces* : preprint, Norman 2007.

[Bon] F.Bonahon, *Bouts des variétés hyperboliques de dimension 3* : Ann. of Math. **124** (1986) 71–158.

[Bow1] B.H.Bowditch, *Systems of bands in hyperbolic 3-manifolds* : Pacific J. Math. **232** (2007) 1–42.

[Bow2] B.H.Bowditch, *Length bounds arising from tight geodesics* : Geom. Funct. Anal. **17** (2007) 1001–1042.

[Bow3] B.H.Bowditch, *Tight geodesics in the curve complex* : Invent. Math. **171** (2008) 281–300.

[Bow4] B.H.Bowditch, *Geometric models for hyperbolic 3-manifolds* : preprint, Southampton 2005.

[Bow5] B.H.Bowditch, *One-ended subgroups of mapping class groups* : preprint, Warwick 2007.

[BrCM] J.F.Brock, R.D.Canary, Y.N.Minsky, *Classification of Kleinian surface groups II: The ending lamination conjecture* : preprint 2004.

[C] R.D.Canary, *A covering theorem for hyperbolic 3-manifolds and its applications* : Topology **35** (1996) 751-778.

[Da] F.Dahmani, *Accidental parabolics and relatively hyperbolic groups* : Israel J. Math. **153** (2006) 93–127.

[DaF] F.Dahmani, K.Fujiwara, *Copies of a one-ended group in a mapping class group* : Groups Geom. Dyn. **3** (2009) 359–377.

[De] T.Delzant, *L'image d'un groupe dans un groupe hyperbolique* : Comment. Math. Helv. **70** (1995) 267–284.

[Du] M.J.Dunwoody, *The accessibility of finitely presented groups* : Invent. Math. **81** (1985) 449-457.

[EE] C.J.Earle, J.Eells, *A fibre bundle description of Teichmüller space* : J. Differential Geometry **3** (1969) 19–43.

[FHS] M.Freedman, J.Hass, P.Scott, *Least area incompressible surfaces in 3-manifolds* : Invent. Math. **71** (1983) 609–642.

[GoH] G.González-Díez, W.J.Harvey, *Surface groups inside mapping class groups* : Topology **38** (1999) 57–69.

[Gr] M.Gromov, *Hyperbolic groups* : in "Essays in Group Theory" (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263.

[Ha] W.J.Harvey, *Boundary structure of the modular group* : in "Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference" (ed. I.Kra, B.Maskit), Ann. of Math. Stud. No. 97, Princeton University Press (1981) 245–251.

[Hi] J.A.Hillman, *Four-manifolds, geometries and knots* : Geometry & Topology Monographs, Vol. 5 (2002).

[J] F.E.A.Johnson, *A rigidity theorem for group extensions* : Arch. Math. **73** (1999) 81-89.

[K] M.Kapovich, *On normal subgroups in the fundamental groups of complex surfaces* : preprint, Salt Lake City 1998.

[LR] C.Leininger, A.W.Reid, *A combination theorem for Veech subgroups of the mapping class group* : Geom. Funct. Anal. **16** (2006) 403–436.

[MaM1] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103-149.

[MaM2] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.

[Mc] C.T.McMullen, *From dynamics on surfaces to rational points on curves* : Bull. Amer. Math. Soc. **37** (2000) 119-140.

[MeSY] W.Meeks, L.Simon, S.T.Yau, *Embedded minimal surfaces, exotic spheres, and manifolds with positive Ricci curvature* : Ann. of Math. **116** (1982) 621–659.

[Mi1] Y.N.Minsky, *Bounded geometry for Kleinian groups* : Invent. Math. (2001) **146** 143–192.

[Mi2] Y.N.Minsky, *The classification of Kleinian surface groups I : Models and bounds* : preprint, Stony Brook 2002.

[Mo] L.Mosher, *Geometric survey of subgroups of mapping class groups* : Handbook of Teichmüller theory, Vol. I, IRMA Lect. Math. Theor. Phys. No. 11, European Math. Soc. Zürich (2007) 387–410.

[O1] J.-P.Otal, *Le théorème d'hyperbolisation des variétés fibrées de dimension 3* : Astérisque No. 235, Société Mathématique de France (1996).

[O2] J.-P.Otal, *Les géodésiques fermées d'une variété hyperbolique en tant que noeuds* : in "Kleinian groups and hyperbolic 3-manifolds" (ed. Y.Komori, V.Markovic, C.Series), London Math. Soc. Lecture Notes Series No. 299, Cambridge University Press (2003) 95–104.

[R] A.W.Reid, *Surface subgroups of mapping class groups* : in "Problems on mapping class groups and related topics" (ed. B. Farb), Proc. Sympos. Pure Math. No. 74, Amer. Math. Soc. (2006) 257–268.

[SY] R.Schoen, S.T.Yau, *Existence of incompressible minimal surfaces and the topology of three-dimensional manifolds with nonnegative scalar curvature* : Ann. of Math. **110** (1979) 127–142.

[T1] W.P.Thurston, *The geometry and topology of 3-manifolds* : notes, Princeton (1979).

[T2] W.P.Thurston, *Hyperbolic structures on 3-manifolds II: Surface groups and manifolds which fiber over the circle* : preprint, Princeton 1986.

[W] H.C.Wang, *Topics in totally disconnected groups* : in "Symmetric spaces" (ed. W.M.Boothby, G.L.Weiss), Marcel Dekker (1972) 460–485.