

Geometric models for hyperbolic 3-manifolds

Brian H. Bowditch

School of Mathematics, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.

<http://www.maths.soton.ac.uk/staff/Bowditch/>

[First draft: August 2005. Revised: October 2006]

0. Introduction.

In [Mi4,BrocCM1], Minsky, Brock and Canary gave a proof of Thurston's Ending Lamination Conjecture for indecomposable hyperbolic 3-manifolds. In this paper, we offer another approach to this which was inspired by the original. Many of the key results are similar, though the overall logic is somewhat different. A possible advantage is that much of the relevant theory of "hierarchies" and other, more technical, parts of the proof can be simplified. We hope that it will help to render some of these ideas more generally accessible. In another paper [Bow5], we show how to adapt these arguments to deal with the general (decomposable) case of a tame 3-manifold. Brock, Canary and Minsky have also stated that they can similarly adapt argument in this regard. Some brief remarks to this effect are given in [BroCM1], and a sequel is promised as [BroCM2]. We also note that an alternative approach to the Ending Lamination Conjecture has recently been proposed by Brock, Bromberg, Evans and Souto [BroBES]. Their approach does not involve the construction of a model.

In the late 1970s Thurston made several important conjectures relating to the geometry of 3-manifolds. Probably the most famous is the "Geometrisation Conjecture" recently announced by Perelman, but the "Tameness Conjecture" and "Ending Lamination Conjecture", have also been central to the development of hyperbolic geometry, and together amount to a complete classification of finitely generated torsion free Kleinian groups.

Let M be a complete hyperbolic 3-manifold with $\pi_1(M)$ finitely generated. The Tameness Conjecture classifies the "geometric" ends of M into two sorts "geometrically finite" and "simply degenerate". This was shown by [T,Bon,Can] to be equivalent to Marden's earlier conjecture [Mar], namely that M is topologically finite, i.e. the interior of a compact manifold. Tameness was proven in the "indecomposable" case by Bonahon [Bon], and recently in general by Agol [Ag] and by Calegari and Gabai [CalG] (see also [So] and [Bow6]).

The Ending Lamination Conjecture goes on to say that M is completely determined up to isometry by a finite set of "end invariants". A number of special cases of this conjecture were established by Minsky [Mi1,Mi2,Mi3], and a proof for all indecomposable 3-manifolds is given in [Mi4,BroCM1]. This makes much use of the work of Masur and Minsky on the curve complex, in particular, hyperbolicity [MasM1] and the theory of tight geodesics and hierarchies [MasM2].

We should note that this work has had implications beyond the theory of 3-manifolds. The Ending Lamination Conjecture is closely related to the large scale structure of Teichmüller space and other geometries associated to a surface, and hence has applications

to the mapping class groups and analysis etc. The wider applications of these ideas have already produced results in a number of different directions. Some further references can be found in the expository article [Bow1].

We should give more careful statement of the Ending Lamination Conjecture. Let M be a complete hyperbolic 3-manifold. For simplicity of exposition we assume it to be orientable. The Margulis lemma gives us a decomposition of M into thick and thin parts. The thin part, where the injectivity radius is small, consists of cusps (\mathbf{Z} -cusps and $\mathbf{Z} \oplus \mathbf{Z}$ -cusps) and tubes (solid tori). We write $\Psi(M)$ for the non-cuspidal part of M , that is M minus the interior of the cusps. Its boundary, $\partial\Psi(M)$, is a disjoint union of euclidean tori and cylinders.

We now assume that $\pi_1(M)$ is finitely generated. Scott's theorem [Sco1,Sco2] gives us a compact submanifold or "core", $\Psi_0 \subseteq \Psi(M)$, whose inclusion into $\Psi(M)$ is a homotopy equivalence. In fact by the relative core theorem of McCullough [Mc] says we can take Ψ_0 so that it meets ∂M in precisely in the torus components of ∂M and in compact annular cores of the cylindrical ends. The ends of $\Psi(M)$ are in bijective correspondence with the relative boundary components of Ψ_0 in $\Psi(M)$. We write $\mathcal{E}(M)$ for the set of such ends — the "geometric" ends of M . Note that there are only finitely many cusps and geometric ends.

One class of end, namely the geometrically finite ends, have long been well understood. Such an end expands exponentially and has a canonical ideal boundary in the form of a Riemann surface — the geometrically finite end invariant. The remaining ends have come to be understood more gradually, but are now known to be "simply degenerate" and have associated an "ending lamination". These notions were defined by Thurston [Th] and clarified by Bonahon [Bon]. The Ending Lamination Conjecture says that M is completely determined by its topology and end invariants — one for each geometric end. The geometrically finite case, where all ends are geometrically finite, had already been dealt with by the deformation theory developed by Ahlfors, Bers, Kra, Marden, Maskit, Sullivan etc. (The collection [BeK] gives an overview of the state of the art prior to Thurston's work.) Thus, what remains to large extent boils down to understanding the simply degenerate ends.

This can also be interpreted in terms of Kleinian groups. We can write $M = \mathbf{H}^3/\Gamma$ where Γ acts properly discontinuously on hyperbolic 3-space \mathbf{H}^3 . We get a decomposition of $\partial\mathbf{H}^3$ into the limit set $\Lambda(\Gamma)$ and discontinuity domain, $D(\Gamma)$. Ahlfors's finiteness theorem [Ah] tells us that $D(\Gamma)/\Gamma$ is a (possibly empty) finite disjoint union of Riemann surfaces of finite type. Here we shall focus on the indecomposable case. In this case, $\Lambda(\Gamma)$ is connected, and is conjectured to be locally connected, though this remains unproven in general.

The approach of [Mi4,BrocCM1] to prove the Ending Lamination Conjecture is to build a "model" riemannian manifold, P , which depends only on the topology and end-invariants of M , and then construct a bilipschitz map from P to M . If M' were another manifold with the same topology and end invariants as M , we get another such map from P to M' and hence a bilipschitz map to between M and M' . This lifts to an equivariant bilipschitz map of universal covers, namely \mathbf{H}^3 , and extends to an equivariant quasiconformal homeomorphism of $\partial\mathbf{H}^3$. It then follows from the "classical" deformation

theory referred to above that this map must be conformal and the result follows.

To construct a model for a simply degenerate end, Minsky [Mi4] uses the curve complex as defined by Harvey [Har], and developed in [MasM1] and [MasM2] into a sophisticated theory of “hierarchies”. From the end invariant, he constructs such a hierarchy, and uses it to find a set of unlinked tubes (solid tori) in $\Sigma \times [0, \infty)$. Each of these tubes is given the structure of a Margulis tube, and locally bounded geometry is put on the remainder. He then defines a lipschitz map to an end of M . A key ingredient of this is the “a-priori bounds” result which says that geodesics in M corresponding to curves in the hierarchy have bounded length. The paper [BroCM1], which includes some of the most technically difficult arguments, shows that this map can be made bilipschitz.

This is of course a very brief summary of a very substantial piece of work. To a large extent it could also serve as a summary of what we do here. Our main result is weaker, in that we only claim that the map from P to M is lipschitz and its lift to the universal covers is a quasi-isometry. (For this we invent the term “sesquilipschitz”, see Section 5.) This, in turns, gives us an equivariant quasi-isometry between the universal covers of M and M' , which is sufficient for a quasiconformal extension to the boundary. Many of the results here have analogous statements in [Mi4, BroCM1], though the proofs are somewhat different and are approached in a different order. We will start from a version of the a-priori bounds theorem [Mi4], a more direct argument for which is given in [Bow4]. We can then bypass most of the theory of hierarchies, and make do with a simpler version, the essential points of which are laid out in Section 4. Given this, most of the construction of the lipschitz map uses fairly standard hyperbolic geometry. Again the technically most complicated part involves showing that this map is a quasi-isometry of universal covers. For this we use a notion of “band” which allows us to use induction over the complexity of the base surface. There is an analogy with the notion of “scaffolds” used in [BroCM1].

We will present most of the argument in the specific context of a doubly degenerate manifold, namely, where $\Psi(M)$ is a topological product, $\Psi(M) \cong \Sigma \times \mathbf{R}$, and both ends are simply degenerate. We do this for several reasons.

Firstly it greatly simplifies the exposition. Nearly all the main ideas can be seen in this context. What remains for the general indecomposable case is largely a matter of describing how the various bits fit together in a more complicated situation.

Secondly, as far as applications of these ideas to Teichmüller theory etc. are concerned, one only really needs to worry about such product manifolds. In the case of a doubly degenerate group, we get a somewhat cleaner, and stronger statement. In particular, one can show that the quasi-isometry constants are uniform, in that they depend only on the topology of the base surface. (A similar uniformity in this case is obtained in [BroCM1].) This is lost (at least without more work) in the general indecomposable case.

A third, though relatively minor, reason is that one is obliged to give some special consideration to the double limit case, since there we have to check that each end of the model gets sent to “right” end of M — a fact that is automatic from the topology in all other situations. (Indeed, precisely this issue caused Bonahon a certain amount of strife in [Bon].)

Many of the ideas for this paper were worked out while visiting the Max Planck Institute in Bonn. The remainder and most of the write-up was carried out when I was

visiting the Tokyo Institute of Technology at the invitation of Sadayoshi Kojima. I thank both institutions for their generous hospitality. I thank Dave Gabai for his many helpful comments on an earlier draft of this paper, and in particular, for the argument that appears at the end of Section 3.

1. Outline of proof.

Before the outline of the proof, we describe the main ingredients we are using.

We shall deal specifically with the case of an indecomposable hyperbolic 3-manifold. We need the fact that such a manifold is topologically finite as well as some of the geometrical principles of tameness as described in [Bon]. In particular, there is a sequence of pleated surfaces going out any simply degenerate end. Here it is sufficient to interpret a “pleated surface” as a uniformly Lipschitz map of a finite-area hyperbolic surface into the 3-manifold in the right homotopy class. Given topological finiteness, we will not need to make direct reference to the core theorems of Scott and McCullough [Sc1,Sc2,Mc], though of course, these were the starting point for Bonahon’s work.

A key topological ingredient is the procedure of Freedman, Hass and Scott [FHS] for replacing a map of a surface in a 3-manifold by an embedding in the same homotopy class. The hypothesis is that such a homotopic embedding should exist somewhere in the 3-manifold. Although the procedure in [FHS] is phrased in terms of minimal surfaces, the relevant part of the argument is a purely combinatorial, and as observed by Bonahon [Bon] one can take the embedded surface in an arbitrarily small neighbourhood of the image of the original map. (However, caution is needed in that this result does not give any a-priori geometric control on the homotopy between these maps.)

We will need some results about topological products. The cobordism theorem of Waldhausen [He1] tells us that two disjoint homotopic proper π_1 -injective embeddings of the same compact surface, Σ into a 3-manifold bound a topological product region $\Sigma \times [0, 1]$ (possibly with fake 3-balls surgered in, which certainly cannot arise in our situation). A result of Brown [Brow] tells that an injective homotopy equivalence of Σ into $\Sigma \times \mathbf{R}$ is ambient isotopic to $\Sigma \times \{t\}$.

While the topology of an end is just a product, we will need to deal with the thick part of such an end, where Margulis tubes are removed. A result of Otal [O2] says that such tubes are unlinked. We shall simply quote this result (though we go some way towards reproving it). The topology of the complement of a set of unlinked curves can nevertheless be quite complicated. In particular we shall need the fact that a proper degree one map between two such manifolds is a homotopy equivalence. This can be seen directly from the Hopfian property [He2]. This in turn uses Thurston’s hyperbolisation of Haken manifolds (see [O1,K]) to reduce to an algebraic theorem of Malcev. An alternative argument that avoids hyperbolisation was suggested to me by Dave Gabai, and is included, in outline, in Section 3. Once this is established, we can go on to use the fact [Wal] that such a map is homotopic to a homeomorphism. This is a now classical fact. The proof is sophisticated but purely topological.

The main combinatorial ingredient is the curve graph, \mathcal{G} , of Harvey [Har] associated

to a surface. We use the fact [MaM1] that this is Gromov hyperbolic [Gr1]. Other proofs of this are given in [Bo1] and [Ham2]. To relate our end invariants to the more standard notion, we use the description of its boundary in terms of arational laminations [Kl] see also [Ham1], (though this could be bypassed if we were simply to define an end invariant as a boundary point of the curve complex). We use the notion of tight geodesics in \mathcal{G} , and some finiteness properties thereof [MasM2]. The relevant results are reproven in [Bow4] and in [Sh]. The only direct result needed from [MasM2], other than inspiration, is the existence of a tight geodesics between any two curves. This is a relatively short combinatorial argument, and logically independent of the remainder of that paper.

To relate all this to our 3-manifold we need a version of the a-priori bound theorem of Minsky [Mi4]. We give a proof of this in [Bow4].

We will be using some standard, mostly elementary facts about Kleinian groups. We use the thick-thin decomposition, arising from the Margulis lemma, and described by Thurston [T]. For the finish, we will use, as many before us, some now “classical” facts from deformation theory of Kleinian of groups as developed by Ahlfors, Bers, Kra, Marden, Maskit, Sullivan etc.

We can now set about our outline proof.

First, we describe the construction of the model manifold. Let us deal with the doubly degenerate case. Thus, $\Psi(M) \cong \Sigma \times \mathbf{R}$ where Σ is compact orientable surface. For the purposes of exposition let us suppose that Σ is closed (though the inductive nature of the proof obliges us to consider non-closed surfaces, even if we start with a closed surface). In this case, $M = \Psi(M)$, and we will write $\Theta(M)$ for the thick part of M .

Let $\mathcal{G}(\Sigma)$ be the curve graph associated to Σ as described in Section 2. The end invariants of M give us two distinct points, $a, b \in \partial\mathcal{G}(\Sigma)$ (Section 11). We now have a tight bi-infinite geodesic in $\mathcal{G}(\Sigma)$ from a to b . This consists of a sequence of curves, $(\alpha_i)_{i \in \mathbf{Z}}$, in Σ with $\alpha_i \cap \alpha_{i+1} = \emptyset$ for all i . One way to imagine this would be to construct what we call a “ladder” in $\Sigma \times \mathbf{R}$. This is a sequence of annuli, $\Omega_i = \alpha_i \times I_i$, where $I_i \subseteq \mathbf{R}$ is a closed interval, and $I_i \cap I_j$ is a non-trivial interval if $|i - j| \leq 1$ and $I_i \cap I_j = \emptyset$ otherwise. A consequence of the a-priori bounds theorem (Theorem 11.4) proven in [Mi4] and in [Bow4], is that the closed geodesics in M in the corresponding homotopy classes all have bounded length. Thus, the ladder in some way reflects the geometry of M . However, we need more than this to determine a riemannian metric on our model. We can extend $\{\Omega_i \mid i \in \mathbf{Z}\}$ to a locally finite set, \mathcal{W} , of disjoint annuli of the form $\alpha \times I$, where $\alpha \subseteq \Sigma$ is a curve and $I \subseteq \mathbf{R}$ is a closed interval, and which is “complete”. Completeness means that for each $t \in \mathbf{R}$, $(\Sigma \times \{t\}) \cap \bigcup \mathcal{W}$ is either a pants decomposition of Σ , or a pants decomposition missing one curve. In the former case, each complementary component is a 3-holed sphere. In the latter case there will be one component that is either a 4-holed sphere or a 1-holed torus. (Combinatorially this is essentially the same as a path in the pants graph. In [MasM2], the analogous procedure is expressed in terms of a resolution of a hierarchy.)

There are many ways one might construct such a complete system of annuli. The only properties we need are laid out in Theorem 4.1. (Some additional conditions are added in Lemmas 4.2 and 4.3 for the purposes of giving an inductive construction, but these are not needed for applications.) We have (P1) that the annuli arise from an iteration of the tight geodesic construction. This is needed in order to obtain the a-priori bound on the length

of the corresponding curves in M . We require (P2) that no two annuli are homotopic, i.e. have the same base curve. This ensures that no two tubes in the model will correspond to the same Margulis tube in M . This is in turn needed to ensure that the map constructed between thick parts is a homotopy equivalence. We require a “tautness” condition (P3) and (P4), expressed in terms of ladders, which says that the annuli follow a geodesic in the curve graph up to bounded distance. This will ensure that the map from our model to M does not crumple up or fold back over large distances. For the inductive structure of the proof we also require this to hold on a class of subsurfaces of Σ , in the appropriate sense. This is incorporated into the statement of properties (P3) and (P4).

Given our annulus system \mathcal{W} , the construction of our model space, P , is relatively straightforward. (When Σ is closed, there is no distinction between P and its “non-cuspidal part” $\Psi(P)$.) First, we cut open each annulus of \mathcal{W} so as to give us a manifold $\Lambda = \Lambda(\mathcal{W})$ with a toroidal boundary component for each annulus. The fact that the local combinatorics of \mathcal{W} are bounded means that we can give Λ a riemannian metric that is natural up to bilipschitz equivalence, and such that each toroidal boundary is euclidean. The exact construction doesn’t much matter, but a precise prescription is given in Section 7. We can now glue in a “model” Margulis tube to each toroidal boundary component. This gives us our model space $P \cong \Sigma \times \mathbf{R}$. (This is essentially just a variant of [Mi4], except that our combinatorial requirements are weaker.)

The extension to a map $f : P \rightarrow M$ now uses the “a-priori bound”. Each model Margulis tube gets sent either to a Margulis tube in M , or to a closed geodesic of bounded length (possibly still quite long though). We write $\mathcal{T}(P)$ for the former set of model tubes, and write $\Theta(P) = P \setminus \text{int} \bigcup \mathcal{T}(P)$ for the “thick” part of the model space. The construction gives us that $\Theta(P) = f^{-1}\Theta(M)$ and that $f|_{\Theta(P)}$ is lipschitz. These constructions are described in Sections 6 and 7, and make use of a result (Lemma 5.9) from Section 5. (It is possible that $\Theta(P)$ may depend on M , but it doesn’t matter what strategy we adopt to construct a map from the model space, once the model space has been defined.)

Topological considerations described in Section 3 now imply that the map $f : \Theta(P) \rightarrow \Theta(M)$ is a homotopy equivalence, and so we get an equivariant map $\tilde{f} : \tilde{\Theta}(P) \rightarrow \tilde{\Theta}(M)$ of universal covers. Some energy is now invested in showing that \tilde{f} is a quasi-isometry. Given this, the argument can be completed as follows. We first use the fact (Lemma 4.8) that, when lifted to an appropriate cover, the boundary of a model tube is quasi-isometrically embedded. This then implies that the map between this boundary and the corresponding boundary from M is a quasi-isometry in the induced euclidean path metrics. This in turn gives us a means (Section 5) of arranging that f is lipschitz on each tube, and a quasi-isometry between the universal covers of tubes. We thus now have a lipschitz map $P \rightarrow M$, and the fact that the lift to the universal covers is a quasi-isometry is relatively straightforward given what we have shown. The details are described in Section 11.

It still remains to explain why the lift, $\tilde{f} : \tilde{\Theta}(P) \rightarrow \tilde{\Theta}(M)$ is a quasi-isometry. This is where tautness comes into play. One can show (see Section 9) that two curves of bounded length and a bounded distance apart in M are also a bounded distance apart in the curve graph, $\mathcal{G}(\Sigma)$. Tautness gives us some control on how far apart such curves can be in the model space. This is a start, but is not sufficient. We need to construct topological barriers in P so that points separated by a barrier in P get mapped to points separated by

a similar barrier in M — taking appropriate account of homotopy classes of paths, since we are really interested in the lift to universal covers. Our barriers are called “bands”. A band in $\Psi(M)$ is a product, $\Phi \times I$, where Φ is a subsurface of Σ , and I a compact interval, and such that $\partial\Phi \times I$ lies in the boundary of a model tube. Its intersection with $\Theta(M)$ is a “band” in $\Theta(M)$. We shall usually insist that $\partial\Phi \times \{0, 1\} \subseteq \Theta(M)$. Much of the second half of Section 4 and Section 10 are devoted to analysing such bands.

To give an idea of how this works, we consider a very simple case. There is a vague sense in which a point, x , of M approximately determines a “fibre” hyperbolic structure on Σ — the domain of a lipschitz “pleated surface” with image close to x . In the bounded geometry case, where there are no Margulis tubes, this structure progresses at roughly uniform rate in the \mathbf{R} direction. (Indeed it stays close to a Teichmüller geodesic [Mi1,Mi2].) A slightly more complicated situation is where we have just one unknotted Margulis tube, T , corresponding to a curve $\alpha \subseteq \Sigma$. Let us suppose that α separates Σ into two subsurfaces Φ_1 and Φ_2 . One possibility is that the area of ∂T is bounded. As we cross the tube, we will be twisting our fibre structure along α , but doing little to the structure on Φ_1 and Φ_2 . Alternatively, while we are crossing α , it may be that the structure on Φ_1 changes a lot. In this case, the length of ∂T in the direction transverse to the longitude of T , becomes very large. We get a band, $B_1 \cong \Phi_1 \times I \subseteq \Theta(M)$, with $B_1 \cap \partial T = \partial\Phi_1 \times I$. This will be very long in the I direction, where the change in the structure takes place. We might get another similar band, $B_2 \cong \Phi_2 \times I$. In this case, B_1 and B_2 together act as barriers between the two ends of $\Theta(M)$. If there is no band on the Φ_2 side of α , then we could sneak around B_1 on the other side of T . However, this path will be in the “wrong” homotopy class, and the lift of B_1 to the cover, $\tilde{\Theta}(M)$, will still serve as a barrier there. In general we may get a very complicated system of nested bands. A general decomposition of M into bands is discussed in [Bow3], though we won’t be needing any result from that paper here.

The logic of our argument is different from the motivation of last paragraph. We construct our bands first in the model space (Section 4). We then show that they correspond to bands in M (Lemma 9.6). The manner in which the form “barriers” is rephrased in terms of pushing around paths and discs (Section 10). We need some general principles of bounded geometry to complete the argument. These are discussed in Section 8.

Finally in Section 12, we discuss how all this works in the general indecomposable case. The only really new ingredient needed is an analysis of the geometry of geometrically finite ends, but this is something pretty well understood.

We remark that most of the arguments presented here should be adaptable to the case of pinched negative curvature (cf. [Can]), though of course, the final rigidity conclusion is no longer valid. One will need to reinterpret various things, for example, the boundaries of Margulis tubes and cusp will no longer be euclidean. A more serious obstruction seems to be the fact that the a-priori bounds result (see [Bow4]) makes use of trace identities to deal with the cases of the 1-holed torus and 4-holed sphere (necessary to start induction). In principle these ought to be amenable to purely geometric proofs, though it would require a somewhat different argument.

2. The curve graph.

We recall some basic facts about the curve graph.

Let Σ be a compact surface with (possibly empty) boundary, $\partial\Sigma$. Throughout this paper, we will use the abbreviations 3HS, 4HS and 1HT for the 3-holed sphere, 4-holed sphere and 1-holed torus respectively.

Let $X(\Sigma)$ be the set of essential (i.e. non-trivial non-peripheral) homotopy classes of simple closed curves in Σ (usually just referred to as *curves*). The *curve complex* as defined by Harvey [Har] is a simplicial complex with vertex set, $X(\Sigma)$, where a subset $A \subseteq X(\Sigma)$ bounds a simplex if its elements can be realised disjointly in Σ . We are only interested in its 1-skeleton, $\mathcal{G}(\Sigma)$, the *curve graph*. We write $V(\mathcal{G}) = X(\Sigma)$, and $d = d_X$ for the combinatorial metric. Given $\alpha, \beta \in \mathcal{G}$, we write $\iota(\alpha, \beta)$ for their geometric intersection number. One can show that $d(\alpha, \beta) \leq \iota(\alpha, \beta) + 1$.

For many purposes it will be convenient to fix some hyperbolic structure on Σ with geodesic boundary, $\partial\Sigma$. In this way, each curve is realised uniquely as a closed geodesic in Σ , and we can use the same notation for a curve and its realisation. This serves purely to simplify the description of certain combinatorial constructions, and bears no relation to the various geometric structures in which we have a genuine interest.

A *multicurve*, γ , in Σ is a non-empty disjoint union of curves. We write $X(\gamma) \subseteq X(\Sigma)$ for the set of components of γ . A multicurve is *complete* if it has maximal cardinality. In this case, each component of $\Sigma \setminus \gamma$ is a 3HS. (It thus gives us a “pants decomposition” of Σ .) We write $\kappa(\Sigma) = 3 \text{genus}(\Sigma) + |\partial\Sigma| - 3$ for the *complexity* of Σ , i.e. the number of curves in any complete multicurve. If $\kappa(\Sigma) = 0$, then Σ is a 3HS. If $\kappa(\Sigma) = 1$, then Σ is either 4HS or a 1HT. In the case where $\kappa(\Sigma) = 1$, we define the *modified curve graph* $\mathcal{G}'(\Sigma)$ with vertex set $V(\mathcal{G}') = X(\Sigma)$ and α, β are adjacent if $\iota(\alpha, \beta)$ is minimal (i.e. 1 for a 1HT and 2 for a 4HS). In both cases, this is a Farey graph.

It was shown by Masur and Minsky [MasM1] that if $\kappa(\Sigma) \geq 2$, then $\mathcal{G}(\Sigma)$ is Gromov hyperbolic (see also [Bow2, Ham2]). We write $\partial\mathcal{G}$ for its Gromov boundary. This has been investigated in [Kl, Ham1].

A *multigeodesic* is a sequence $(\gamma_i)_i$ of multicurves such that for all i, j and all $\alpha \in X(\gamma_i)$ and $\beta \in X(\gamma_j)$, $d(\alpha, \beta) = |i - j|$. It is *tight* if for all non-terminal indices, i , any curve that crosses γ_i must also cross either γ_{i-1} or γ_{i+1} . A *tight geodesic*, $(\alpha_i)_i$ is a sequence of curves such that there exists a tight multigeodesic $(\gamma_i)_i$ such that $\alpha_i \in X(\gamma_i)$ for all i . (Note that this “tight” terminology is now standard, if a little confusing: A tight geodesic is a geodesic in \mathcal{G} in the usual sense. However it need not be a tight as a multigeodesic, in the sense defined. In most cases we will be talking about tight geodesics. We will only need to specify tight multigeodesic for a construction in Section 4: see Lemma 4.2.) The notion of tightness was introduced in [MasM2]. They show that the set of tight geodesics between any two curves in non-empty and finite. Other arguments for finiteness are given in [Bow4] and [Sh], the latter giving explicit bounds.

One can make a stronger statement, for example:

Lemma 2.1 : *Suppose that $(\alpha_i)_{i \in \mathbb{N}}$ and $(\beta_i)_{i \in \mathbb{N}}$ are sequences of curves, each converging to a point of either $\partial\mathcal{G}(\Sigma)$ or $X(\Sigma)$ then for any bounded set $A \subseteq X$, there is a finite subset,*

$B \subseteq A$, such that for all sufficiently large i, j any curve in A also lying on any tight geodesic from α_i to β_j must lie in B .

Proof : Let $a, b \in V(\mathcal{G}) \cup \partial\mathcal{G}$ be the limits of $(\alpha_i)_i$ and $(\beta_i)_i$. If $a, b \in V(\mathcal{G})$, then the result is immediate from the finiteness of tight geodesics between two points (see [MaM2], [Bow4] or [Sh]). Suppose $a, b \in \partial\mathcal{G}$. If $a = b$, we can take $B = \emptyset$. If $a \neq b$, the result follows from other finiteness results for tight geodesics. For example, in [Bow4], it is shown that if $\alpha, \beta \in X(\Sigma)$ and $r \in \mathbf{N}$, then there is some finite $C \subseteq X(\Sigma)$ such that if $(\gamma_i)_{i=0}^p$ is a tight geodesic with $d(\alpha, \gamma_0) \leq r$ and $d(\beta, \gamma_p) \leq r$, then $\gamma_i \in C$ for all i with $12r \leq i \leq p - 12r$. The rest is just an exercise in hyperbolic spaces. We can take α and β arbitrarily close to a and b , and $r \geq 0$ so that for all i and j sufficiently large, any geodesic from α_i to β_j meets both $N(\alpha, r)$ and $N(\beta, r)$. Choosing α and β far away from our bounded set A , the result now follows. Finally the case where $a \in V(\mathcal{G})$ and $b \in \partial\mathcal{G}$ follows by a variation on the above result, namely if $\gamma_0 = \alpha$ and $d(\beta, \gamma_p) \leq r$ and then γ_i lies in a finite subset for all $i \leq p - 12r$. \diamond

This result allows us to use diagonal sequence arguments. For example, we obtain the fact [MasM2] that any two boundary points are connected by a bi-infinite tight geodesic.

By a *subsurface* of Σ we mean the closure, in Σ , of a non-empty connected open subsurface $\text{int}(\Phi)$ of Σ with geodesic boundary, $\partial\Phi$. We write $\partial^\Sigma\Phi = \partial\Phi \setminus \partial\Sigma$. We express it in this way since we want to allow the possibility of two boundary curves in $\partial^\Sigma\Phi$ being identified to a single curve in Σ . (We could homotope Φ to an embedded surface with a complementary annulus so that these boundary curves become genuinely distinct, though for most purposes, it will be convenient to realise things with respect to some fixed hyperbolic structure.) We are allowing Σ as a subsurface of itself. A subsurface cannot be a disc or an annulus. Note that if Φ' is a proper subsurface of Φ then $\kappa(\Phi') < \kappa(\Phi)$.

The following definitions arise out of the discussion of ‘‘hierarchies’’ in [MasM2]. Let $\kappa = \kappa(\Sigma)$. Given $Q \subseteq X(\Sigma)$ and $k \in \mathbf{N}$, let $Y_k(Q)$ be Q together with all those curves $\gamma \in X(\Sigma)$ such that there is some subsurface $\Phi \subseteq \Sigma$ with $X(\partial^\Sigma\Phi) \subseteq Q$ and $2 \leq \kappa(\Phi) \leq \kappa - k + 1$, and two curves $\alpha, \beta \in Q \cap X(\Phi)$, such that γ lies on some tight geodesic in $\mathcal{G}(\Phi)$ from α to β . Note that, for $k \geq \kappa$ there is no such subsurface, so $Y_k(Q) = Q$. For any k , we set $Y^k(Q) = Y_k Y_{k-1} Y_{k-2} \cdots Y_1(Q)$, and set $Y^\infty(Q) = Y^\kappa(Q)$.

Note that this set contains the union, $Y_0(Q)$, of all tight geodesics between any pair of points of $Q \subseteq X(\Sigma)$. (For the first step, we are allowing $\Phi = \Sigma$). However, all constructions involving proper subsurfaces occur in a 1-neighbourhood in $X(\Sigma)$ of a curve already constructed. In particular, we see that $Y^k(Q) \subseteq N(Y_0(Q), k)$ and so $Y^\infty(Q) \subseteq N(Y_0(Q), \kappa)$. Note that these sets are all locally finite. Also, Lemma 2.1, tells us that:

Lemma 2.2 : *Suppose that $(\alpha_i)_{i \in \mathbf{N}}$ and $(\beta_i)_{i \in \mathbf{N}}$ are sequences of curves, each converging to a point of either $\partial\mathcal{G}(\Sigma)$ or $X(\Sigma)$ then for any bounded set $A \subseteq X$, there is a finite subset, $B \subseteq X$, such that for all sufficiently large i, j , $A \cap Y^\infty(\{\alpha_i, \beta_j\}) \subseteq B$. \diamond*

Thus, after passing to a subsequence, we can assume that $Y^\infty(\{\alpha_i, \beta_i\})$ stabilises on each bounded set to some finite set, thereby giving us a locally finite limit.

We have the following variation for subsurfaces of complexity 1. Given $Q \subseteq X(\Sigma)$, define $\bar{Y}(Q)$ as for $Y_k(Q)$, replacing the statement ‘‘ $2 \leq \kappa(\Phi) \leq \kappa - k + 1$ ’’ by the statement

“ $\kappa(\Phi) = 1$ ”, and replacing “ $\mathcal{G}(\Phi)$ ” by “ $\mathcal{G}'(\Phi)$ ”. (All geodesics in $\mathcal{G}'(\Phi)$ are deemed to be “tight” in this case.) We write $\bar{Y}^\infty(Q) = \bar{Y}(Y^\infty(Q))$. A similar discussion applies. In particular, Lemma 2.2 holds with \bar{Y}^∞ replacing Y^∞ .

A number of variations on the above definitions are possible, and would probably serve just as well. We have chosen a formulation that works well with our inductive constructions below, and in Section 4.

A *path of multicurves* is a sequence $(\gamma_i)_i$ of multicurves such that for each i , γ_{i+1} is obtained from γ_i by either adding or deleting one component. A path in $(\alpha_i)_i$ in \mathcal{G} determines a path of multicurves by inserting $\alpha_i \cup \alpha_{i+1}$ between α_i and α_{i+1} .

Lemma 2.3 : *Suppose that $\alpha, \beta \in X(\Sigma)$ with $d(\alpha, \beta) \geq 3$. Then there is a complete multicurve, γ , with $\alpha \subseteq \gamma$ and $X(\gamma) \subseteq Y^\infty(\{\alpha, \beta\})$.*

Proof : Let $\underline{\gamma} = (\gamma_i)_{i \in D}$ be a path of multicurves with indexing set $D = \{0, \dots, m\}$, so that $\gamma_0 = \alpha$ and $\gamma_m = \beta$. Given any $i \in D$, write $c_i = |X(\gamma_i)|$ for the number of components of γ_i . Thus $c_0 = c_m = 1$. We say that i is a “local minimum” if $c_{i-1} = c_{i+1} = c_i + 1$. We assume:

(*) If i is not a local minimum, then $X(\gamma_i) \subseteq Y^{c_i-1}(\{\alpha, \beta\})$.

Note that such a path of multicurves exists: just take any tight geodesic from α to β and construct a path of multicurves from it as described earlier.

Let $\kappa = \kappa(\Sigma)$. Given any $n \in \{1, 2, \dots, \kappa\}$ let $v_i(\underline{\gamma}) = |\{i \in D \mid c_i = n\}|$. $v_i(\underline{\gamma}) = (v_1(\underline{\gamma}), \dots, v_\kappa(\underline{\gamma})) \in \mathbf{N}^\kappa$. We order \mathbf{N}^κ lexicographically. It is thus well-ordered. We now choose $\underline{\gamma}$ so as to minimise $v(\underline{\gamma})$ among all paths of multicurves (of any length) satisfying (*) above.

Given any $i \in D$, let F_i be the union of γ_i and all those components of $\Sigma \setminus \gamma_i$ which are three-holed spheres.

First, we claim that if i is a local minimum, then $F_{i-1} = F_{i+1}$. To see this let $\gamma_{i-1} = \gamma_i \cup \delta$ and $\gamma_{i+1} = \gamma_i \cup \epsilon$. Note that $i-1$ and $i+1$ are not local minima, and so δ, ϵ and all the components of γ_i lie in $Y^{c_i}(\{\alpha, \beta\})$. Now δ and ϵ must cross, otherwise we could replace γ_i by $\gamma_i \cup \delta \cup \epsilon$, thereby decreasing $v(\underline{\gamma})$. Thus, δ and ϵ lie in the same component, Φ , of $\Sigma \setminus \gamma_i$. If $\kappa(\Phi) > 1$, then we can connect δ to ϵ by some tight geodesic $\delta = \delta_0, \delta_1, \dots, \delta_p = \epsilon$ in $\mathcal{G}(\Phi)$. Now each $\delta_j \in Y_{c_i+1}(Y^{c_i}(\{\alpha, \beta\})) = Y^{c_i+1}(\{\alpha, \beta\})$. We can now replace γ_i by the sequence

$$\gamma_i \cup \delta_0 \cup \delta_1, \quad \gamma_i \cup \delta_1, \quad \gamma_i \cup \delta_1 \cup \delta_2, \quad \dots, \quad \gamma_i \cup \delta_{p-1}, \quad \gamma_i \cup \delta_{p-1} \cup \delta_p.$$

To verify (*), note that the $\gamma_i \cup \delta_j$ are all at local minima, and that $\gamma_i \cup \delta_j \cup \delta_{j+1}$ has $c_i + 2$ components. We have therefore again reduced $v(\underline{\gamma})$. We conclude that $\kappa(\Phi) = 1$. From this it follows that $F_{i-1} = F_i \cup \Phi = F_{i+1}$, as claimed.

Next, we claim that there does not exist any i with $c_{i+2} = c_{i+1} + 1 = c_i + 2 = c_{i-1} + 1$. For in such a case, we have $F_{i-1} = F_{i+1}$ by the above, and so we could replace γ_i by $\gamma_{i-1} \cup \delta$ and γ_{i+1} by $\gamma_i \cup \delta$, where $\delta = \gamma_{i+2} \setminus \gamma_{i+1}$. Note that since γ_{i+1} and γ_{i+2} were not local minima in the original path, all curves lie in $Y^{c_i+1}(\{\alpha, \beta\})$. We have reduced $v(\underline{\gamma})$ giving a contradiction. By the same argument, we cannot have any i with $c_{i-2} = c_{i-1} + 1 = c_i + 2 = c_{i+1} + 1$.

We have thus shown that c_i increases monotonically from $c_0 = 1$ up to a maximum value $c_p = k$ and then alternates between k and $k - 1$ before decreasing monotonically from $c_q = k$ down to $c_m = 1$. By the first observation, we see that $F_p = F_q$. But $\alpha \subseteq \gamma_p \subseteq F_p$ and $\beta \subseteq \gamma_q \subseteq F_q$. Since $\alpha \cup \beta$ fills Σ , it follows that $F_p = F_q = \Sigma$. Thus γ_i is a complete multicurve (i.e. $c_p = k = \kappa$). Moreover, $\alpha \subseteq \gamma_p$ and $X(\gamma_p) \subseteq Y^{\kappa-1}(\{\alpha, \beta\})$. We can thus set $\gamma = \gamma_p$. \diamond

3. Topology of products.

In this section, we give some discussion to the topology of a product $\Psi = \Sigma \times \mathbf{R}$, where Σ is a compact surface. As in Section 2, it is convenient to fix a hyperbolic structure on Σ , so that $X(\Sigma)$ can be realised as a set of geodesic curves in Σ . (This structure bears no relation to the geometry we are interested in.) We will write $\partial\Psi = \partial\Sigma \times \mathbf{R}$. We write $\pi_\Sigma : \Psi \rightarrow \Sigma$ and $\pi_V : \Psi \rightarrow \mathbf{R}$ for the vertical and horizontal projections respectively.

Definition : A *horizontal fibre* in Ψ is a subset of the form $\Sigma \times \{t\}$ for some $t \in \mathbf{R}$.

A horizontal curve is a subset of the form $\gamma \times \{t\}$ where $\gamma \subseteq \text{int } \Sigma$ is a simple closed geodesic, and $t \in \mathbf{R}$.

More generally we have:

Definition : A *fibre*, $S \subseteq \Psi$, in Σ is the image $f(\Sigma)$ in Ψ of an embedding $f : \Sigma \rightarrow \Psi$ with $f^{-1}\partial\Psi = \partial\Sigma$ such that $\pi_\Sigma \circ f$ is homotopic to the identity on Σ (rel $\partial\Sigma$).

If S, S' are two fibres, we write $S < S'$ if they are disjoint and S' can be homotoped to the positive end of Ψ in $\Psi \setminus S$.

A theorem of Brown [Brow] tells us that if S is any fibre in Ψ , then there is an ambient isotopy sending S to a horizontal fibre. Inductively, we see that if S_1, S_2, \dots, S_n are disjoint fibres, then there is an ambient isotopy of Ψ sending each S_i to a horizontal fibre, $S \times \{t_i\}$. After permuting the indices, we can assume that $t_1 < t_2 < \dots < t_n$, and so $S_1 < S_2 < \dots < S_n$. From this, we see easily that $<$ defines a total order on any locally finite set of disjoint fibres. If $S < S'$ we write $[S, S']$ for the compact region bounded by S and S' .

By a *homotopy fibre*, we mean a map $f : \Sigma \rightarrow \Psi$ with $f^{-1}\partial\Psi = \partial\Sigma$ and with $\pi_\Sigma \circ f$ homotopic to the identity. We will usually take such a map to be in general position. We will also sometimes refer to its image as a homotopy fibre. The result of [FHS] tells us that any neighbourhood, U , of $f(\Sigma)$ contains a fibre. It can also be assumed to contain any given point in $f(\Sigma)$. We can also define an order on homotopy fibres by writing $f(\Sigma) < f(\Sigma')$ if $f(\Sigma) \cap f(\Sigma') = \emptyset$ and $f(\Sigma')$ can be homotoped to positive infinity in $\Psi \setminus f(\Sigma)$. This is again a total order on any disjoint, locally finite collection. Moreover, if S, S' are fibres close to $f(\Sigma)$ and $f(\Sigma')$, then $S < S'$ if and only if $f(\Sigma) < f(\Sigma')$.

Definition : By a *curve* in Ψ we mean a simple closed curve that is not homotopic to a point or into $\partial\Psi$.

A *horizontal curve* is a subset of the form $\gamma \times \{t\}$ where $\gamma \subseteq \text{int}\Sigma$ is a simple closed geodesic and $t \in \mathbf{R}$.

Suppose that \mathcal{L} is a locally finite collection of curves in Ψ .

Definition : We say that \mathcal{L} is *unlinked* if for every $\alpha \in \Sigma$ there is some fibre, $S(\alpha)$, containing α , so that $S(\alpha) \cap S(\beta) = \emptyset$ for all distinct $\alpha \neq \beta \in \mathcal{L}$.

A curve α is *unknotted* if $\{\alpha\}$ is unlinked.

Clearly an unknotted curve is homotopic to a curve in $X(\Sigma)$. In the definition of an unlinked set, one can take the fibres $(S(\alpha))_{\alpha \in \mathcal{L}}$ to be a locally finite. (For our purposes, one could add this as a hypothesis.) Note that, by the earlier remarks, we can find an isotopy of Ψ sending each of the fibres to a horizontal one. We can then homotope each curve so that it becomes geodesic. Thus, a set of curve is unlinked if and only if there is an ambient isotopy sending it to a set of horizontal curves.

Suppose $\gamma \subseteq M$ is a curve. We can locally compactify $\Psi \setminus \gamma$ by adjoining a toroidal boundary component $\Delta(\gamma)$. We can think of $\Delta(\gamma)$ as the unit normal bundle to γ . We write $\Lambda(\gamma)$ for the resulting manifold. Note that it comes equipped with a natural homotopy class of meridian curve, $m(\gamma)$. We can recover Ψ up to homeomorphism by gluing in a solid torus, $T(\gamma)$, along $\Delta(\gamma)$ so that the meridian bounds a disc in $T(\gamma)$. If γ is unknotted, it also has a natural class of longitude. It can be defined as a simple curve that can be homotoped to infinity in $\Lambda(\gamma)$. It can also be determined by sitting the curve in some (indeed any) fibre. More generally, if Λ is any locally finite set of curves, we can form the manifold $\Lambda(\mathcal{L})$ by adding toroidal boundaries to $\Psi \setminus \bigcup \mathcal{L}$. Thus $\partial\Lambda = \partial\Psi \cup \bigcup_{\gamma \in \mathcal{L}} \Delta(\gamma)$.

We will need the following:

Proposition 3.1 : *Suppose that \mathcal{L} and \mathcal{L}' are unlinked collections of curves. Suppose that no two elements of \mathcal{L} are homotopic in Ψ , and similarly for \mathcal{L}' . Suppose that $f : \Lambda(\mathcal{L}) \rightarrow \Lambda(\mathcal{L}')$ is a proper degree-one map with $f^{-1}\partial\Lambda(\mathcal{L}') = \partial\Lambda(\mathcal{L})$. Suppose that $f|_{\partial\Lambda(\mathcal{L})}$ is a homotopy equivalence and sends the meridian and longitude of each toroidal boundary component to the meridian and longitude of its image. Then f is homotopic to a homeomorphism from $\Lambda(\mathcal{L})$ to $\Lambda(\mathcal{L}')$.*

Note that f necessarily extends to a degree-one map from Ψ to itself. Since the induced homomorphism on the surface group $\pi_1(\Psi) \cong \pi_1(\Sigma)$ is surjective, it follows from the residual finiteness for such groups [Sc3] (via the hopfian property) that this extension is a homotopy equivalence of Ψ . (In fact, this will be immediate from the construction in our applications.) Thus, there is no loss in taking the extension to be homotopic to the identity on Ψ .

The issue of degree-one maps between 3-manifolds has been investigated by a number of people (see, for example, [Wan] for a survey). There are certainly many examples which are not homotopy equivalences. Some positive results are also known, but I know of no

result that directly implies the statement given above.

If we can show that $\Lambda(\mathcal{L})$ and $\Lambda(\mathcal{L}')$ are homeomorphic, then we are in reasonably good shape. Consider the case where \mathcal{L} is finite. Then Λ is a topologically finite and admits a hyperbolic structure hyperbolic by the work of Thurston [O1,Ka]. Thus $\pi_1(\Lambda)$ is residually finite. Since it is finitely generated it is hopfian by a result of Malcev. It follows that f induces an isomorphism of fundamental groups (see [He2]). Since $\pi_2(\Lambda)$ is trivial (by the sphere theorem), f is a homotopy equivalence. Now using the work of Waldausen [Wal], it follows that f is homotopic to a homeomorphism. In the case where \mathcal{L} is infinite, we will need a bit more explanation as to why the map on fundamental groups is injective, but this is relatively simple. Of course, there is a highly non-trivial input into this. An argument, suggested by Gabai, that bypasses hyperbolisation will be outlined at the end of this section.

Thus, most of the additional work is involved in showing that $\Lambda(\mathcal{L})$ and $\Lambda(\mathcal{L}')$ are indeed homeomorphic. For this we need to define a partial order on the link components, and show that this is preserved. It all seems fairly intuitive, but the details are a bit subtle.

Suppose $\alpha, \beta \subseteq \Psi$ are unlinked curves. Write $\alpha \approx \beta$ if they they do not cross homotopically Σ , i.e. if $d(\pi_\Sigma \alpha, \pi_\Sigma \beta) \leq 1$. This is equivalent to asserting that there is some fibre of Ψ containing both α and β . We will write $\alpha \preceq \beta$ (respectively $\alpha \succeq \beta$) if β can be homotoped out the positive (respectively negative) end of Ψ in $\Psi \setminus \alpha$.

Suppose that $\alpha \preceq \beta$ and $\alpha \succeq \beta$. Then β can be homotoped from the negative to the positive end of Ψ without ever meeting α . Such a homotopy must intersect any fibre of Ψ in at least some (not necessarily embedded) curve homotopic to β . From this one can see that $\alpha \approx \beta$ by the above definition.

Lemma 3.2 : *Suppose α and β are unlinked. Then the following are equivalent:*

- (1) $\alpha \preceq \beta$
- (2) $\beta \succeq \alpha$
- (3) *We can find disjoint fibres $S \supseteq \alpha$ and $S' \supseteq \beta$ with $S < S'$.*

Proof : It's clearly enough to show that (1) implies (3). By hypothesis we have disjoint fibres, $Z \supseteq \alpha$ and $Z' \supseteq \beta$. If $Z' < Z$, then $\alpha \succeq \beta$, and so by the above observation, $\alpha \approx \beta$. It follows that α and β are contained in some common fibre. We can now push these fibres slightly so that they become disjoint in the order required. \diamond

We write $\alpha \prec \beta$ to mean that $\alpha \preceq \beta$ and $\alpha \not\approx \beta$. Thus, for any two unlinked curves, α and β , exactly one of the relations $\alpha \prec \beta$, $\beta \prec \alpha$ or $\alpha \approx \beta$ holds.

Suppose that $\gamma_1, \gamma_2, \dots, \gamma_n$ are unlinked curves, with $\gamma_i \prec \gamma_{i+1}$ for all i . By definition, we have disjoint fibres, $S_i \supseteq \gamma_i$, and by the above observation, we must have $S_i < S_{i+1}$ for all i . From this we can deduce that there does not exist any finite cycle of unlinked curves, $\gamma_0, \gamma_1, \dots, \gamma_n = \gamma_0$ with $\gamma_i \prec \gamma_{i+1}$ for all i .

If \mathcal{L} is any unlinked set of curves, we can let $<$ be the transitive closure of the relation \prec on \mathcal{L} . We see that this is a strict partial order on \mathcal{L} . Moreover, $\alpha < \beta$ implies $\alpha \preceq \beta$, and so if $\alpha \not\approx \beta$ we also get $\alpha \prec \beta$.

We aim to show that this order on \mathcal{L} , together with the natural map of \mathcal{L} to $X(\Sigma)$

determines \mathcal{L} up to ambient isotopy, and so, in particular, $\Lambda(\mathcal{L})$ up to homeomorphism. We use the following observation.

We say that a strict total order \ll is *compatible* with the partial order $<$ if $\alpha \ll \beta$ implies $\alpha < \beta$. We say that it is *discrete* if all intervals are finite.

Lemma 3.4 : *Suppose that \mathcal{L} is an unlinked set of curves. Suppose that \ll is a discrete total order on \mathcal{L} compatible with $<$. Then we can find a set of disjoint fibres, $(S(\gamma))_{\gamma \in \mathcal{L}}$ with $\Gamma \subseteq S(\gamma)$ for all γ , and $S(\alpha) < S(\beta)$ if and only if $\alpha \ll \beta$.*

Proof : Since \mathcal{L} is unlinked, we can find a locally finite disjoint collection of fibres $Z(\gamma) \supseteq \gamma$. Write $\alpha \ll' \beta$ to mean that $Z(\alpha) < Z(\beta)$. This defines another discrete total order compatible with $<$. We can now find a sequence, $(\ll_n)_n$, of discrete total orders, all compatible with $<$, with $\ll_0 = \ll'$, with \ll_n stabilising on \ll on any finite subset of \mathcal{L} , and with \ll_{n+1} obtained from \ll_n by interchanging the order on a pair of consecutive elements of \mathcal{L} . We suppose inductively that $Z_n(\gamma) \supseteq \gamma$ is a collection of disjoint fibres inducing the order \ll_n . Suppose that \ll_{n+1} is obtained by interchanging the order on α and β . These are consecutive, which means that the region $[Z_n(\alpha), Z_n(\beta)]$ contains no other curve of \mathcal{L} . Since both orders are compatible with \leq , we must have $\alpha \approx \beta$. We can now construct two new disjoint fibres, $Z_{n+1}(\alpha) \supseteq \alpha$ and $Z_{n+1}(\beta) \supseteq \beta$, both in $[Z_n(\alpha), Z_n(\beta)]$, but with the opposite order. We set $Z_{n+1}(\gamma) = Z_n(\gamma)$ for all $\gamma \neq \alpha, \beta$. Now the above process stabilises on any compact subset of Ψ , and so we eventually end up with a collection of fibres, $(S(\gamma))_\gamma$ inducing \ll as required. \diamond

Lemma 3.5 : *Suppose \mathcal{L} and \mathcal{L}' are unlinked set of curves with a bijection $[\gamma \mapsto \gamma']$ from \mathcal{L} to \mathcal{L}' . Suppose that γ and γ' are homotopic in Ψ for all γ . Suppose also that $\alpha \preceq \beta$ implies $\alpha' \preceq \beta'$. Then there is an end-preserving self-homeomorphism of Ψ sending γ to γ' for all γ .*

Proof : First note that by the condition on homotopies, we have $\alpha \approx \beta$ if and only if $\alpha' \approx \beta'$. By the trichotomy, we see that $\alpha \prec \beta$ if and only if $\alpha' \prec \beta'$, and so $\alpha < \beta$ if and only if $\alpha' < \beta'$.

Now, let $S'(\gamma') \supseteq \gamma'$ be a locally finite disjoint set of fibres of Ψ . Applying Lemma 3.4, we can find a locally finite disjoint set of fibres $S(\gamma) \supseteq \gamma$ such that $S(\alpha) < S(\beta)$ if and only if $S'(\alpha') < S'(\beta')$. We can now find an isotopy of Ψ sending each $S(\gamma)$ to $S'(\gamma')$. Since γ' is homotopic to γ in Ψ and hence in Σ , we can isotop γ to γ' in $S(\gamma)$ and extend to an ambient isotopy in a small neighbourhood of $S(\gamma)$. The resulting homeomorphism sends \mathcal{L} to \mathcal{L}' as required. \diamond

Lemma 3.6 : *Suppose that \mathcal{L} and \mathcal{L}' are unlinked sets of curves, and that $f : \Psi \rightarrow \Psi$ is an end-preserving homeomorphism with $f^{-1}(\bigcup \mathcal{L}') = \bigcup \mathcal{L}$, and with $f|_{\bigcup \mathcal{L}}$ a homeomorphism to $\bigcup \mathcal{L}'$. Then there is an end-preserving homeomorphism $g : \Psi \rightarrow \Psi$ homotopic to f in Ψ with $g|_{\bigcup \mathcal{L}} = f|_{\bigcup \mathcal{L}}$.*

Proof : We may as well suppose that f is homotopic to the identity on Ψ . Suppose $\alpha, \beta \in \mathcal{L}$. If $\alpha \preceq \beta$, then we can homotope β out the positive end of Ψ in $\Psi \setminus \alpha$. The image of this homotopy under f sends $f(\beta)$ out the positive end in $\Psi \setminus f(\alpha)$. Thus $f(\alpha) \preceq f(\beta)$. Lemma 3.7 now gives us a homeomorphism of Ψ sending each $\gamma \in \mathcal{L}$ to $f(\gamma)$. By isotopy in a neighbourhood of γ we can assume that $f|_{\gamma} = g|_{\gamma}$. \diamond

We remark that we do not in fact need to assume that \mathcal{L}' is unlinked in Ψ . This is a consequence of the other hypotheses. It can be seen by an argument using a refinement of the result of [FHS] mentioned above, and can be extracted from the discussion of Otal [O2], where he shows that “short curves” in a hyperbolic 3-manifold product are unlinked. However, rather than reproduce that argument here, we shall simply quote Otal’s result when applying this (see Proposition 7.3).

Lemma 3.7 : *Let $f : \Lambda(\mathcal{L}) \longrightarrow \Lambda(\mathcal{L}')$ be as in the hypothesis of Proposition 3.1. Then f is a homotopy equivalence.*

Proof : By extending over the tori $T(\gamma)$ we get a map satisfying the hypotheses of Lemma 3.5, and so it follows that $\Lambda(\mathcal{L})$ and $\Lambda(\mathcal{L}')$ are homeomorphic. Let $\Gamma = \pi_1(\Lambda(\mathcal{L}))$. The map f induces an epimorphism of Γ , and we need to show that this is also injective.

If \mathcal{L} were finite, then this follows from the fact that $\pi_1(\Lambda(\mathcal{L}))$ satisfies the hopfian property as described earlier.

In general suppose $g \in \Gamma$ were in the kernel. We can represent it by a closed curve $\delta \subseteq \Lambda(\mathcal{L})$, which lies between two fibres, say $S < Z$ in Ψ . We can take these disjoint from $\bigcup \mathcal{L}$. Let $\mathcal{L}_0 = \{\gamma \in \mathcal{L} \mid \gamma \subseteq [S, Z]\}$, and write \mathcal{L}'_0 for the corresponding subset of \mathcal{L}' . Now f extends to degree one map between $\Lambda(\mathcal{L}, \mathcal{L} \setminus \mathcal{L}_0)$ and $\Lambda(\mathcal{L}', \mathcal{L}' \setminus \mathcal{L}'_0)$. By the finite case, it now follows that δ bounds a disc in $\Lambda(\mathcal{L}, \mathcal{L} \setminus \mathcal{L}_0)$. But we can now push that disc into the region $[S, Z]$, and so we see that δ bounds a disc in $\Lambda(\mathcal{L})$. In other words, g is trivial in Γ .

The result now follows by Whitehead’s theorem, given that the higher homotopy groups are trivial. \diamond

Proof of Proposition 3.1 : We have shown (Lemma 3.7) that f is a homotopy equivalence. If \mathcal{L} is finite, since $\Lambda(\mathcal{L}) \cong \Lambda(\mathcal{L}')$ is Haken, the result then follows by Waldhausen [Wal].

To deal with the general case, we first show that if $F \subseteq \Lambda(\mathcal{L})$ is a surface with $\partial F = F \cap \partial \Lambda(\mathcal{L})$, then $f|_F$ is homotopic to an embedding in $\Lambda(\mathcal{L}')$ relative to $\partial \Lambda(\mathcal{L}')$. (In fact, this is all we need for our applications.) To see this, choose fibres $S, Z \subseteq \Lambda(\mathcal{L}')$, disjoint from $\bigcup \mathcal{L}'$, so that $f(F) \subseteq [S, Z]$. Let $\mathcal{L}'_0 = \{\gamma \in \mathcal{L}' \mid \gamma \subseteq [S, Z]\}$, and let $\mathcal{L}_0 \subseteq \mathcal{L}$ be the corresponding curves in \mathcal{L} . Now f is a homotopy equivalence from $\Lambda(\mathcal{L}, \mathcal{L} \setminus \mathcal{L}_0)$ to $\Lambda(\mathcal{L}', \mathcal{L}' \setminus \mathcal{L}'_0)$. By the finite case, $f|_F$ is homotopic to an embedding in $\Lambda(\mathcal{L}', \mathcal{L}' \setminus \mathcal{L}'_0)$. By [FHS] we can take this embedding in a small neighbourhood of $f(F)$, and so, in particular, in $[S, Z]$. But Ψ retracts onto $[S, Z]$, so we can now push the homotopy into $[S, Z]$, and so these surfaces are also homotopic in $\Lambda(\mathcal{L}')$.

To complete the proof, we can apply this result to a sequence of fibres, $(S_i)_{i \in \mathbf{Z}}$ in $\Lambda(\mathcal{L})$, whose images are disjoint, so as to find disjoint fibres Z_i in $\Lambda(\mathcal{L}')$, homotopic to

$f(S_i)$. We can now apply the finite case to the regions between these fibres. We omit the details, since we have already shown what we need for the following corollary. \diamond

As we have noted, the fact we really need is the following. It was proven in the course of Proposition 3.1, and can also be viewed as a combination of Proposition 3.1 and [FHS].

Corollary 3.8 : *Let $f : \Lambda(\mathcal{L}) \rightarrow \Lambda(\mathcal{L}')$ be as in the hypotheses of Proposition 3.1. Suppose that $F \subseteq \Lambda(\mathcal{L})$ is a properly embedded π_1 -injective compact surface ($F \cap \partial\Lambda(\mathcal{L}) = \partial F$). Let U be any neighbourhood of $f(F)$ in $\Lambda(\mathcal{L}')$. Then there is a proper embedding $g : F \rightarrow U$ such that $f|_F$ is homotopic in $\Lambda(\mathcal{L}')$ to g relative to ∂F . \diamond*

We conclude this section with an outline of how one can bypass the use of hyperbolisation in the proof of Proposition 3.1. It elaborates on a suggestion of Gabai, and I thank him for his permission to include it here.

We can suppose that $\mathcal{L} = \mathcal{L}'$, and that $f : \Lambda(\mathcal{L}) \rightarrow \Lambda(\mathcal{L})$ extends to a map homotopic to the identity on Ψ . We claim that f is homotopic to a homeomorphism (in fact, the identity) on $\Lambda(\mathcal{L})$.

We write $\Sigma_t = \Sigma \times \{t\}$. We can assume that $\mathcal{L} = \{\alpha_i \mid i \in I\}$, $I \subseteq \mathbf{Z}$ and each α_i is a curve in Σ_i . We can also assume that $\Sigma \times (\mathbf{Z} + \frac{1}{2}) \subseteq \Lambda(\mathcal{L})$.

We will first show that if $t \in \mathbf{Z} + \frac{1}{2}$, then $f|_{\Sigma_t}$ is homotopic in $\Lambda(\mathcal{L})$ to the inclusion $\Sigma_t \hookrightarrow \Lambda(\mathcal{L})$. Let J be the vertical range of $f(\Sigma_t)$, i.e. the compact interval $\{u \in \mathbf{R} \mid \Sigma_u \cap f(\Sigma_t) \neq \emptyset\}$. Suppose first that $I \cap J \cap [t, \infty) \neq \emptyset$. Let i be its maximal element, and let $A = \alpha_i \times [i, \infty)$. We can assume that $(f|_{\Sigma_t})^{-1}A$ consists of simple closed curves. These will be either trivial or homotopic to α_i . Since $\Lambda(\mathcal{L})$ is aspherical, after homotopy in $\Lambda(\mathcal{L})$, we can get rid of trivial curves, and since it is atoroidal, we can get rid of pairs of non-trivial curves with opposite orientations. We are left with either p positively oriented curves or $-p$ negatively oriented curves homotopic to α , where $p \in \mathbf{Z}$ is the number of times $f(\Sigma_t)$ wraps around α_i . More precisely, $p = \langle \omega, f(\Sigma_t) \rangle$, where $\omega \in H^2(\Lambda(\{\alpha_i\}))$ measures the intersection with a ray in $\Lambda(\{\alpha_i\})$ from α_i to $+\infty$. For sufficiently negative u , $f(\Sigma_u) \subseteq \Sigma \times (-\infty, n)$ and so $\langle \omega, f(\Sigma_u) \rangle = 0$. Since $f(\Sigma_u)$ is homotopic to $f(\Sigma_t)$ in $\Lambda(\{\alpha_i\})$ it follows that $p = \langle \omega, f(\Sigma_t) \rangle = 0$. In other words, we have pushed $f(\Sigma_t)$ off A . We can now homotope it (in $\Lambda(\mathcal{L})$) below Σ_i . We can continue inductively until $I \cap J \cap [t, \infty) = \emptyset$. Proceeding similarly below, we push $f(\Sigma_t)$ so that its vertical range lies in the component of $\mathbf{R} \setminus I$ containing t . After a further homotopy, we will get $f(\Sigma_t) \subseteq \Sigma_t$. Now, $f|_{\Sigma_t} : \Sigma_t \rightarrow \Sigma_t$ is homotopic to the identity in $\Sigma \times \mathbf{R}$ and hence in Σ_t . This proves the claim.

Performing such homotopies for all $t \in \mathbf{Z} + \frac{1}{2}$, we can assume that $f|_{\Sigma \times (\mathbf{Z} + \frac{1}{2})}$ is just the inclusion $\Sigma \times (\mathbf{Z} + \frac{1}{2}) \hookrightarrow \Sigma \times \mathbf{R}$.

Now let $P_n = \Sigma \times [n - \frac{1}{2}, n + \frac{1}{2}]$, and let $R_n = P_n \cap \Lambda(\mathcal{L})$. (In other words, R_n is P_n with at most one tube drilled out.) We next homotope f so that $f(R_n) = R_n$. The idea is to proceed as we did for the surfaces Σ_t . Given $i \in I \cap J \cap [n + 1, \infty)$ maximal, let $A_i = \alpha_i \times [i, \infty)$ as before. This time, we can assume that each component of $f^{-1}(A_i)$ is a surface in the interior of R_n . Moreover, by standard 3-manifold topology, we can assume it to be incompressible. It's not hard to see that any incompressible surface in R_n must be boundary parallel. Thus it must either be a fibre or be homotopic in P_n to α_n . But it must be homotopic in $\Sigma \times \mathbf{R}$ to α_i , giving a contradiction. We have arranged

that $f(R_n) \cap A_i = \emptyset$. Continuing as with Σ_t , we eventually homotope $f(R_n)$ into R_n as claimed.

Next, if $n \notin I$, then $R_n = P_n$, and by Waldhausen [Wal], we can homotope $f|_{R_n}$ fixing $R_{n-\frac{1}{2}} \cup R_{n+\frac{1}{2}}$ to a homeomorphism (in fact the identity). If $n \in I$, let $B = (\alpha_n \times [n, n + \frac{1}{2}]) \cap R_n$. Since R_n is atoroidal, we can assume $f(B) = B$. Again, we can assume that $f^{-1}(B) \setminus B$ consists of incompressible surfaces in $R_n \setminus B$, and therefore empty. In other words $f^{-1}(B) = B$. After a homotopy, holding the curve $\Sigma_{n+\frac{1}{2}} \cap B$ fixed (though perhaps rotating the other boundary component of B) we get $f|_B$ to be inclusion. Cutting R_n along B , we get a homeomorphic copy of $\Sigma \times [0, 1]$ and so, again, we are done by Waldhausen. Doing this for all n , we homotope f to the identity in $\Lambda(\mathcal{L})$.

4. Annulus systems.

Here we describe the basic combinatorial construction used to build our model space. It is analogous to the theory of hierarchies, though much less sophisticated.

Let Σ be a compact surface. We fix a hyperbolic structure on Σ with geodesic boundary. Let $\Psi = \Sigma \times \mathbf{R}$. We write $\pi_\Sigma : \Psi \rightarrow \Sigma$ and $\pi_V : \Psi \rightarrow \mathbf{R}$ be vertical and horizontal projections respectively. We refer to the two ends of Ψ as the *positive* and *negative* ends. We are only really interested in the total order on the vertical coordinate in \mathbf{R} . We are thus free to adjust by any orientation preserving homeomorphism of \mathbf{R} .

Definition : A *horizontal curve* in M is a subset of the form $\gamma \times \{t\}$ for some curve $\gamma \subseteq \Sigma$, and some $t \in \mathbf{R}$.

A *horizontal surface* in M is a subset of the form $\Phi \times \{t\}$ for some subsurface $\Phi \subseteq \Sigma$ and some $t \in \mathbf{R}$. If $\Phi = \Sigma$, it is called a *horizontal fibre*.

A *vertical annulus* is a subset of the form $\gamma \times I$ where $I \subseteq \mathbf{R}$ is a non-trivial compact interval.

A *strip* is a subset of the form $\Phi \times I$ where $\Phi \subseteq \Sigma$ and $I \subseteq \mathbf{R}$ is a non-trivial compact interval.

As before, in what follows we shall fix some hyperbolic structure on Σ with geodesic boundary, and realise curves as geodesics. In this way, curves will automatically intersect minimally, so it will simplify the combinatorial arguments.

Let $\Omega = \gamma \times I$ be a vertical annulus, with $I = [\partial_- I, \partial_+ I] \subseteq \mathbf{R}$. We write $\partial_\pm \Omega = \gamma \times \partial_\pm I$. These are horizontal curves. (We shall be a bit sloppy about distinguishing points of \mathbf{R} from 1-element subsets.) We write $\partial_H \Omega = \gamma \times \partial I = \partial_- \Omega \sqcup \partial_+ \Omega$.

Let $B = \Phi \times I$ be a strip. We write $\partial_\pm B = \Phi \times \partial_\pm I$, $\partial_H B = \Phi \times \partial I = \partial_- B \sqcup \partial_+ B$, $\partial_V B = \partial \Phi \times I$ and $\partial_V^\Sigma B = \partial^\Sigma \Phi \times I$. (Recall that $\partial^\Sigma \Phi = \partial \Phi \setminus \partial \Sigma$ is the relative boundary of Φ in Σ .) We refer to $\partial_H B$ as the *horizontal boundary* of B . It consists of two horizontal surfaces. We refer to $\partial_V B$ as the *vertical boundary*. Note that the relative boundary of B in Ψ is $\partial_V^\Sigma B \cup \partial_H B$. We define the complexity $\kappa(B)$ as B as $\kappa(\Phi)$. We refer to $\pi_\Sigma B$ as the *base surface* of B .

Definition : An *annulus system*, \mathcal{W} , in Ψ is a locally finite collection of disjoint vertical annuli.

Let $W = \bigcup \mathcal{W}$. Given $t \in \mathbf{R}$, let $\gamma_t = \pi_\Sigma(W \cap (\Sigma \times t))$. This is either empty or a multicurve in Σ . Clearly, W is completely determined by the piecewise constant map $[t \mapsto \gamma_t]$. For many purposes, it will be convenient to assume that \mathcal{W} is *vertically generic*, that is, if $\Omega, \Omega' \in \mathcal{W}$ with $\pi_V \partial_H \Omega \cap \pi_V \partial_H \Omega' \neq \emptyset$, then $\Omega = \Omega'$. This is easily achieved by pushing the horizontal boundaries of annuli up or down a little. In this case, Ω is combinatorially equivalent to sequence $(\gamma_i)_i$ where each γ_i is either empty or a multicurve, and γ_{i+1} is obtained from γ_i by either adding or deleting a multicurve. We say that Ω is *vertically full* if $\mathbf{R} = \pi_V W$. In this case, the γ_i are all multicurves. In Section 2, we referred to such a sequence (γ_i) as a *path of multicurves*. In other words, there is a bijective correspondence between paths of multicurves and vertically full (and generic) annulus systems up to vertical reparametrisation.

A *ladder* is a minimal vertically full annulus system. This corresponds to a path of multicurves where the number of components alternates between 1 and 2. Those with one component constitute a path in the curve complex. Thus a ladder is combinatorially equivalent to a (bi-infinite) path in $\mathcal{G}(\Sigma)$.

Definition : A strip, $B \subseteq \Psi$ is a *band* (with respect to \mathcal{W}), if $\partial_V^\Sigma B \subseteq W$.

We can view $W_B = (W \cap B) \setminus \partial_V B$ as a finite annulus system on B (at least if $\partial_H B \cap \partial_H \Omega = \emptyset$ for all $\Omega \in \mathcal{W}$), and a similar discussion applies. In this case if W_B is full then it corresponds to a finite path of multicurves in Φ . We write \mathcal{W}_B for the set of components of W_B .

Definition : An annulus system is *complete* if for any horizontal fibre $S \subseteq \Psi$, each component of $S \setminus W$ has complexity at most 1.

In this case, a *brick* is a maximal band whose interior does not meet W .

Any such brick has complexity at most 1. We refer a brick as *type 0* or *type 1* depending on whether its complexity is 0 or 1. Let $\mathcal{D} = \mathcal{D}(\mathcal{W})$ be the set of all bricks. One sees easily that this is a locally finite collection of bands with disjoint interiors, and that $\Psi = \bigcup \mathcal{D}$. If two bands meet in a horizontal surface, then one is of type 0 and the other of type 1. We can recover \mathcal{W} from \mathcal{D} as the set of components of $W = \bigcup_{B \in \mathcal{D}} \partial_V^\Sigma B$. We write $\mathcal{W} = \mathcal{W}(\mathcal{D})$.

Given an annulus system, \mathcal{W} , we can define $\Lambda = \Lambda(\mathcal{W})$ as the metric completion of $\Psi \setminus W$. Thus there is a toroidal boundary component, $\Delta(\Omega)$, of Λ , associated to each $\Omega \in \mathcal{W}$. Indeed, $\partial \Lambda = \partial \Psi \cup \bigcup_{\Omega \in \mathcal{W}} \Delta(\Omega)$. There is natural projection, $\pi_\Psi : \Lambda \rightarrow \Psi$ that is injective on $\text{int } \Lambda$. On each $\Delta(\Omega)$ it is injective on $\pi_\Psi^{-1} \partial_H \Omega$ and two-to-one elsewhere. As in Section 3, $\Delta(\Omega)$ comes equipped with a free homotopy class of longitude, denoted $l(\Omega)$ and meridian denoted $m(\Omega)$. We refer to this procedure as “opening out” the annuli of \mathcal{W} . We can also fill them back in again.

Given any subset, $\mathcal{W}_0 \subseteq \mathcal{W}$ let $\Lambda(\mathcal{W}, \mathcal{W}_0)$ be the manifold obtained from Λ by gluing

in a solid torus, $T(\Omega)$ to $\Delta(\Omega)$ so that the meridian bounds a disc. Thus, we recover Ψ up to homeomorphism as $\Lambda(\mathcal{W}, \mathcal{W})$. From a purely topological point of view this is a rather fruitless exercise, however, we will want to view these spaces as having different structures. We will be regarding Ψ together with \mathcal{W} as an essentially combinatorial object, whereas $\Lambda(\mathcal{W}, \mathcal{W})$ will be given a geometric structure.

Suppose that \mathcal{W} is a complete annulus system. We obtain a *brick decomposition* of $\Lambda = \Lambda(\mathcal{W})$ by lifting each brick in Ψ to a *brick* in Λ . Note that two vertical boundary components of such a (lifted) brick may become identified under the projection map, π_Ψ , but the projection is otherwise injective on bricks. By abuse of notation we will also denote this lifted brick decomposition by \mathcal{D} . In this case, if two bricks meet one is of type 0 and the other of type 1. They meet along a horizontal 3HS.

We want to describe a particular construction of complete annulus systems. Note that Lemma 2.3 gave us a means of connecting two curves by a path of multicurves giving us a complete annulus system in some compact region of Ψ . However this will not be sufficient for our purposes. We will require some additional properties. To describe these, we need some further definitions.

Let \mathcal{W} be an annulus system and $B \subseteq \Psi$ a band. We write $X_\pm(B) = \pi_\Sigma(W \cap \partial_\pm B \setminus \partial_V B) \subseteq X(\Phi) \subseteq X(\Sigma)$. Recall that a ladder in B is minimal set of annuli in \mathcal{W}_B which is vertically full in B (i.e. $\pi_V W_B = \pi_V B$).

Definition : The *height*, $H(B)$ of B is the minimal length of a ladder in B .

More intuitively, the height can be thought of as the minimal number of annuli we need to cross to get from one horizontal boundary component of B to the other, where we are allowed to jump between bands along paths in horizontal fibres.

We also write $H_0(B) = d_X(X_-(B), X_+(B))$. Note that $H_0(B) \leq H(B)$. In the case where \mathcal{W} is complete and $\kappa(B) \geq 2$, these quantities are finite. We define the *slackness* of B as $H(B) - H_0(B)$.

Definition : We say that B is *k-taut* if its slackness is at most k .

In the case where $\kappa(B) = 1$, \mathcal{W}_B consists of a sequence $\Omega_0, \dots, \Omega_n$ of annuli whose vertical projections are disjoint and occur in this order. In this case, we say that B is *taut* if $(\pi_\Sigma \Omega_i)_i$ is a geodesic segment in $\mathcal{G}'(\Phi)$.

The main result about existence of complete annulus systems in the bi-infinite case can be stated as follows.

Let Σ be a compact surface with complexity, $\kappa(\Sigma) \geq 2$. Suppose $a, b \in \partial\mathcal{G}(\Sigma)$ are distinct and that $(\alpha_i)_i$ and $(\beta_i)_i$ are sequences of curves converging on a and b respectively. As observed in section 2, the sets $Y^\infty(\{\alpha_i, \beta_i\})$ converge locally on some locally finite subset $Y \subseteq B$. This lies a bounded distance (depending only on $\kappa(\Sigma)$) from any bi-infinite geodesics from a to b in $\mathcal{G}(\Sigma)$.

Theorem 4.1 : *There is a constant c depending only on $\kappa(\Sigma)$ such that for any $a, b \in \partial\mathcal{G}$ and Y constructed as above, we can find a complete annulus system \mathcal{W} such that*

- (P1) $X(\mathcal{W}) \subseteq \bar{Y}^\infty(Y)$.
- (P2) If $\Omega, \Omega' \in \mathcal{W}$ with $\pi_\Sigma \Omega = \pi_\Sigma \Omega'$, then $\Omega = \Omega'$.
- (P3) Every band in Ψ of complexity at least 2 in Ψ is c -taut.
- (P4) Every band of complexity 1 is taut.

Note that in (P3) we are allowing bands with base surface Σ . Tautness then tells us that the path of multicurves associated to \mathcal{W} is quasi-geodesic.

Property (P1) will eventually serve to show that the map from the model space is lipschitz (using the ‘‘a-priori bounds’’, Theorem 11.4). Property (P2) is needed to show that the map has degree one on the ‘‘thick parts’’ of these space. Properties (P3) and (P4) are needed for the reverse coarse inequalities, to show that our map is a quasi-isometry.

Using the local finiteness of our sets, Y (Lemma 2.4), we can see that we can reduce to a finite case of Theorem 4.1. Here we consider only a finite band, $O = \Sigma \times I$. We interpret ‘‘completeness’’ to include the statement that $X_-(\mathcal{W})$ and $X_+(\mathcal{W})$ are both complete multicurves. In this case, we started with two curves α, β which fill Σ . We can replace Y by $Y^\infty(\{\alpha, \beta\})$, and insist that $\alpha \in X_-(\mathcal{W})$ and $\beta \in X_+(\mathcal{W})$. To get us started, we can apply Lemma 2.3 to give us our a multicurve that will serve as $X_-(\mathcal{W})$. (This is the only reason we need the two Y ’s in property (1). One can clearly formulate other versions that would not involve us in constructing quite so many tight geodesics, but there seems little point for our purposes.)

We now have the basis for proving Theorem 4.1. The argument will be by induction on complexity. We need to state the induction hypothesis in a different way. This formulation is mostly stronger than that already given. However, it is weaker in the sense that we are assuming we are given an initial complete multicurve, and we will also forget, for the moment about bands of complexity 1.

We say that a horizontal curve $\gamma \subseteq O$ is *compatible* with an annulus system $W = \bigcup \mathcal{W}$ if either $\gamma \subseteq W$ or $\gamma \cap W = \emptyset$. We say that a vertical annulus $\Omega \subseteq O$ is *compatible* with W if $\Omega \cap W$ is empty, a single vertical annulus or a boundary curve of Ω . (This implies that $W \cup \Omega$ is also an annulus system.)

Lemma 4.2 : *Suppose that $\kappa(\Sigma) \geq 1$, and that α, β are multicurves in Σ with α complete. Then there is a complete (vertically generic) annulus system, $W = \bigcup \mathcal{W} \subseteq O = \Sigma \times I$ satisfying:*

- (1) $\alpha = X_-(\mathcal{W})$ and $\beta \subseteq X_+(\mathcal{W})$.
- (2) If $\Omega \subseteq O$ is a vertical annulus with $\partial_-\Omega$ and $\partial_+\Omega$ both compatible with W , then Ω is compatible with W .
- (3) If Ω is a vertical annulus with $\partial_-\Omega$ compatible with W and with $\partial_+\Omega \subseteq \partial_+O$ and not crossing β , then Ω is compatible with W (so that by completeness, $W \cap \partial_+O \subseteq \Omega$).
- (4) If $B = \Phi \times J \subseteq O$ is a band of complexity at least 2, then B is $2(\kappa(\Sigma) - \kappa(\Phi) + 1)$ -taut.
- (5) $X(\mathcal{W}) \subseteq Y^\infty(X(\alpha) \cup X(\beta))$.

Proof : The proof will be induction on $\kappa(\Sigma)$.

In the case where $\kappa(\Sigma) = 1$, there is not much to be said. Here, α, β are just single curves, and we enlarge them slightly to be vertical annuli. We put them in a single annulus if they happen to be equal.

The case where $\kappa(\Sigma) = 2$ is not much harder. We choose $\gamma_0 \in X(\alpha)$ and $\gamma_n \in X(\beta)$ with $d(\gamma_0, \gamma_n) = d(X(\alpha), X(\beta)) = n$, say, and connect γ_0 to γ_n by a tight geodesic $\gamma_0, \gamma_1, \dots, \gamma_n$ in $\mathcal{G}(\Sigma)$. The path of multicurves,

$$\alpha, \quad \gamma_0, \quad \gamma_0 \cup \gamma_1, \quad \gamma_1, \quad \dots, \quad \gamma_{n-1} \cup \gamma_n \quad [, \quad \gamma_n, \quad \beta]$$

gives us a ladder in O , which in this case, is a complete annulus system. (The last two bracketed terms are omitted if β is a single curve.) The properties stated are all easily verified in this case.

Now suppose that $\kappa(\Sigma) \geq 3$ and that the lemma holds for all surfaces of smaller complexity. To apply the induction hypothesis, first consider the special case where $X(\alpha) \cap X(\beta) \neq \emptyset$. Let $\gamma \subseteq \alpha$ be the union of all components of α that do not cross any curve of β . Thus, γ and $\gamma \cup \beta$ are multicurves. Now, $\gamma \times I$ cuts O into subsets of the form $\Phi \times I$ where Φ is (the closure of) a component of $\Sigma \setminus \gamma$. If $\beta \cap \text{int } \Phi = \emptyset$, let $W_\Phi = (\alpha \cap \text{int } \Phi) \times I$. Otherwise, we apply the inductive hypothesis to $\Phi \times I$ and the curves $\alpha \cap \text{int } \Phi$ and $\beta \cap \text{int } \Phi$ to give us a complete annulus system W_Φ in $\Phi \times I$. We now set $W = (\gamma \times I) \cup \bigcup_\Phi W_\Phi$ as Φ ranges over all such subsurfaces. (We can modify W so that it becomes vertically generic.) It is complete, and all the above properties are easily verified. For (4), note that if a band B does not lie in any of the components, $\Phi \times I$, then it is crossed by a single annulus and so $H(B) = H_0(B) = 0$, so it is 0-taut.

The general construction is as follows. Let $n = d(X(\alpha), X(\beta))$. Similarly as in the complexity 2 case, we choose $\gamma_0 \in X(\alpha)$ and $\gamma_n \in X(\beta)$ with $d(\gamma_0, \gamma_n) = n$. Now connect γ_0 to γ_n by a tight multigeodesic $\gamma_0, \gamma_1, \dots, \gamma_n$. (This needs to be a bona fide tight multigeodesic, as discussed in Section 2.) Now write $I = [0, n+1]$ and set $O_i = \Sigma \times [i, i+1]$ for $i \in \{0, 1, \dots, n\}$. Thus $O = \bigcup_i O_i$ and $\partial_+ O_i = \partial_- O_{i+1}$. We apply the above construction to give us a complete annulus system, $W_0 \subseteq O_0$ with $\alpha = X_-(W_0)$ and $\gamma_0 \cup \gamma_1 \subseteq X_+(W_0)$. Let α_1 be the complete multicurve $X_+(W_0)$. We now do the same thing in O_1 to get a complete annulus system, $W_1 \subseteq O_1$ with $\alpha_1 = X_-(W_1)$ and $\gamma_1 \cup \gamma_2 \subseteq X_+(W_1)$. Set $\alpha_2 = X_+(W_1)$, and continue inductively. For the final step, O_n , we get $W_n \subseteq O_n$ with $\gamma_{n-1} \cup \gamma_n \subseteq \alpha_n = X_-(W_n)$ and $\beta \subseteq X_+(O_n)$. Note that for all i , $X_+(O_i) = X_-(O_{i+1})$ and so the annulus systems match up. We set $W = \bigcup_{i=0}^n W_n \subseteq O$, and let \mathcal{W} be the set of components of W .

By construction, W is complete annulus system satisfying (1). We verify the remaining properties in turn.

We need the following observation. Suppose Ω is any vertical annulus with $\partial_H \Omega$ compatible with W . If $\partial_- \Omega \subseteq O_i$ and $\partial_+ \Omega \subseteq O_j$, then $d(X(\gamma_i), \pi_\Sigma \Omega) \leq 1$ and $d(X(\gamma_j), \pi_\Sigma \Omega) \leq 1$ and so $j - i = d(X(\gamma_i), X(\gamma_j)) \leq 2$. In other words, Ω can enter at most three of the O_i .

(2) Suppose that $\partial_- \Omega \subseteq O_i$ and $\partial_+ \Omega \subseteq O_j$ are both compatible with W . By the above observation, there are three cases:

Case (2a): $j = i$ so $\Omega \subseteq O_i$, and we are done by the inductive procedure, i.e. (1) applied to W_i .

Case (2b): $j = i + 1$. Note that $\pi_\Sigma \Omega$ does not cross either γ_i or γ_{i+1} (by compatibility of $\partial_- \Omega$ with W_i and $\partial_- \Omega$ with W_{i+1} respectively). By (3) applied to W_i , we see that $\Omega \cap O_i$ is compatible with W_i , and that $\Omega \cap \partial_+ O \subseteq W$. We now apply (2) to W_{i+1} , showing that $\Omega \cap O_{i+1}$ is compatible with W_{i+1} . Thus Ω is compatible with W as required.

Case (2c): $j = i + 2$. In this case, π_Σ does not cross either γ_i or γ_{i+2} . By tightness of the multigeodesic $(\gamma_i)_i$ we see that it cannot cross γ_{i+1} either. As in (2b), we see that $\Omega \cap O_i$ is compatible with W_i , and that $\Omega \cap \partial_+ O_i \subseteq W$. Applying (3) to W_{i+1} , we see that $\Omega \cap O_{i+1} \subseteq W$. Finally applying (2) to W_{i+2} , $\Omega \cap O_{i+2}$ is compatible with W_{i+2} . Thus Ω is compatible with W as required.

(3) Suppose Ω is a vertical annulus with $\partial_- \Omega$ compatible with W and $\partial_- \Omega \subseteq \partial_+ O$ and compatible with β . Let $\partial_- \Omega \subseteq \partial_+ O$. As with (2), there are three possibilities.

Case (3a) $i = n$. We just apply (3) to W_n .

Case (3b) $i = n - 1$. As in (2b), we see that $\Omega \cap W_i$ is a vertical annulus meeting $\partial_+ O_i = \partial_- O_n$, and applying (3) to W_n , we see that $\Omega \cap O_n \subseteq W$.

Case (3c) $i = n - 2$. We argue as in (2c). This time, we get $\Omega \cap (O_{n-1} \cup O_n) \subseteq W$.

(4) Let B be a band. First, consider the case where $\pi_\Sigma(B) = \Sigma$. Let $\partial_- B \subseteq O_i$ and $\partial_+ B \subseteq O_j$. Thus, $X(\gamma_i) \subseteq X_-(W \cap B)$ and $X(\gamma_j) \subseteq X_+(W \cap B)$, and so $H_0(B) = d(X_-(W \cap B), X_+(W \cap B)) \geq d(X(\gamma_i), X(\gamma_j)) - 2 = j - i + 2$. By construction we have a ladder crossing in $W \cap B$ of length at most $j - i$. Thus $H(B) - H_0(B) \leq 2$. In other words this is 2-taut, so we are done.

For the other cases, we make the following observation. Suppose that a band B is a union of two subbands, $B = B_1 \cup B_2$, meeting at a common horizontal boundary. If $H(B_2) = 0$ (i.e. some annulus in \mathcal{W} crosses $B \setminus \partial_V B$) then $H(B) = H(B_1)$ and $H_0(B) \geq H_0(B_1) - 1$. Thus, if B_1 is k -taut, then B is $(k + 1)$ -taut.

Suppose now that B is a band with $\Phi = \pi_\Sigma B \neq \Sigma$. Since $\partial_V^\Sigma B$ lies in W , B can meet at most three O_i . Suppose $\partial_- B \subseteq O_i$ and $\partial_+ B \subseteq O_j$. We have three cases.

Case (4a) $j = i$. $B \subseteq O_i$, so we apply (4) to W_i .

Case (4b) $j = i + 1$. If $\gamma_{i+1} \cap \text{int } \Phi = \emptyset$, then (by (3) applied to W_i) $B \cap W_i$ is just a product: $(B \cap W \cap \partial_+ O_i) \times \pi_V(B \cap O_i)$ and so $B \cap W$ is combinatorially identical to $B \cap W_{i+1}$. By (4) applied to W_{i+1} , $B \cap W_{i+1}$ is $2(\kappa(\Sigma) - \kappa(\Phi))$ -taut, and so the same applies to $B \cap W$, and we are happy. In the case where $\gamma_{i+1} \cap \text{int } \Phi \neq \emptyset$, then $H(B \cap O_i) = 0$. By (4) applied to W_i , $B \cap W_i$ is $2(\kappa(\Sigma) - \kappa(\Phi))$ -taut, and so by above observation, $B = (B \cap O_i) \cup (B \cap O_{i+1})$ is $(2(\kappa(\Sigma) - \kappa(\Phi)) + 1)$ -taut.

Case (4c). $j = i + 2$. Now, $\gamma_i \cup \gamma_{i+2}$ is connected. Since neither multicurve can cross $\partial^\Sigma \Phi$, $(\gamma_i \cup \gamma_{i+2}) \cap \partial \Phi = \emptyset$. We are again reduced to two subcases. First, if $(\gamma_i \cup \gamma_{i+2}) \cap \text{int } \Phi$, then by tightness of γ_i , we also have $\gamma_i \cap \text{int } \Phi = \emptyset$. Applying (3) to $B \cap O_i$, and then again to $B \cap O_{i+1}$, we see that $B \cap W \cap (O_i \cup O_{i+1})$ is just a product. Thus, $B \cap W$ is combinatorially identical to $B \cap W_{i+2}$. By (4) applied to W_{i+2} , the latter is $2(\kappa(\Sigma) - \kappa(\Phi))$ -taut. The second subcase is when $\gamma_i \cup \gamma_{i+2} \subseteq \text{int } \Phi$. In this case, $H(B \cap O_i) = H(B \cap O_{i+2}) = 0$. Also, by (4) applied to W_{i+1} , $B \cap W_{i+1}$ is $2(\kappa(\Sigma) - \kappa(\Phi))$ -taut. Thus, the above observation applied twice tells us that B is $2(\kappa(\Sigma) - \kappa(\Phi)) + 2$ -taut, as required.

(5) By construction, $X(\gamma_i) \subseteq Y_1(X(\alpha) \cup X(\beta))$ for all i . By (5) applied to W_i , we

have $X(W_i) = Y_\kappa Y_{\kappa-1} \cdots Y_3 Y_2 (X(\gamma_{i-1}) \cup X(\gamma_i) \cup X(\gamma_{i+1}))$. Thus $X(W) = \bigcup_i X(W_i) \subseteq Y_\kappa \cdots Y_2 Y_1 (X(\alpha) \cup X(\beta)) = Y^\infty (X(\alpha) \cup X(\beta))$ as required. \diamond

We can now include the case of complexity 1 bands as an afterthought:

Lemma 4.3 : *Suppose that $\alpha, \beta \subseteq \Sigma$ are multicurves with α complete. Then there is a complete annulus system, $W = \bigcup \mathcal{W} \subseteq O$, satisfying:*

- (1) $\alpha = \pi_\Sigma(W \cap \partial_- O)$ and $\beta = \pi_\Sigma(W \cap \partial_+ O)$,
- (2) If $\Omega, \Omega' \in \mathcal{W}$ with $\pi_\Sigma \Omega = \pi_\Sigma \Omega'$, then $\Omega = \Omega'$.
- (3) If $B \subseteq O$ is a band with $\kappa(B) \geq 2$, then B is $2(\kappa(\Sigma) - \kappa(B) + 2)$ -taut.
- (4) If $B \subseteq O$ is a band with $\kappa(B) = 1$, then B is taut.
- (5) $X(\mathcal{W}) \subseteq \bar{Y}^\infty (X(\alpha) \cup X(\beta))$.

Proof : Let \hat{W} be the complete annulus system given by Lemma 4.2. Suppose that A is a type 1 brick with respect to \hat{W} , in other words, a band with $\kappa(A) = 1$ and with $W \cap A \setminus \partial_V A$ consisting of two curves, $\delta \subseteq \partial_- A$ and $\epsilon \subseteq \partial_+ A$. By property 4.2(1), $\pi_\Sigma \delta \neq \pi_\Sigma \epsilon$. Let $\delta = \gamma_0, \gamma_1, \dots, \gamma_n = \epsilon$ be a geodesic in $\mathcal{G}'(\pi_\Sigma A)$. Let $\Omega_i = \gamma_i \times I_i \subseteq A$ be disjoint annuli occurring in this order vertically. This give an annulus system, $W_A = \bigcup_i \Omega_i$ in A . We perform this construction for all type 1 bricks. Since these bricks are disjoint, we get an annulus system $W = \hat{W} \cup \bigcup_A W_A$, and A varies over all type 1 bricks. We need to verify the above properties.

- (1) Since $W \cap \partial_H O = \hat{W} \cap \partial_H O$, this follows by construction.
- (2) Suppose that $\Omega, \Omega' \in \mathcal{W}$, with $\pi_\Sigma \Omega = \pi_\Sigma \Omega'$. Let Ω'' be a vertical annulus connecting Ω to Ω' (so that $\Omega \cup \Omega'' \cup \Omega'$ is a vertical annulus). Now $\partial_H \Omega' \subseteq W$, so $\partial_\pm \Omega'$ are compatible with \hat{W} . It follows by 4.2(2) that Ω' is compatible with \hat{W} and so by the construction, the only way that can happen is if $\Omega = \Omega'$.
- (3) Let B be a band with $\kappa(B) \geq 2$. We can assume that $\partial_V^\Sigma B \subseteq \hat{W}$, for if a boundary component were in $W \setminus \hat{W}$, then B would be crossed by an annulus of W and so it would have height 0, and there is nothing to prove. In other words, B is band with respect to \hat{W} . With respect to \hat{W} is $2(\kappa(\Sigma) - \kappa(B) + 1)$ -taut. In passing to W , $H(B)$ can only increase. It is possible there may be some new curves in $\partial_\pm B \cap W$, but this could decrease $H_0(B)$ by at most 2. It follows that, with respect to W , B is $2(\kappa(\Sigma) - \kappa(B) + 2)$ -taut.
- (4) Let B be a band with $\kappa(B) = 1$. As in (3), we need only consider the case where $\partial B \subseteq \hat{W}$. We note that there is no component of \hat{W} contained in the interior of B . For if Ω were such a component, we could construct another annulus, Ω' in B with $\pi_\Sigma \Omega' \neq \pi_\Sigma \Omega$, so that $\partial_H \Omega' \cap \mathcal{W} = \emptyset$, and with $\pi_V \Omega$ contained in the interior of $\pi_V \Omega'$. Thus Ω' crosses Ω . But $\partial_\pm \Omega$ is compatible with \hat{W} , so this contradicts 4.2(2). We see that the only element of \mathcal{W} in the interior of B were those added in some brick, A , of our construction. Since these were made out a geodesic in $\mathcal{G}'(\pi_\Sigma B)$, it follows that B is, by definition, taut.
- (5) Clearly $X(\mathcal{W}) \subseteq \bar{Y}(X(\hat{W}))$. Since $X(\hat{W}) \subseteq Y^\infty (X(\alpha) \cup X(\beta))$, the result follows. \diamond

Putting Lemma 2.3 together with Proposition 3.3, we see that we have proven the finite analogue of Theorem 3.1. The bi-infinite case now follows using Lemma 2.2, as discussed earlier.

There are various combinatorial properties of bands that we will need. In what follows, we can define the *height* $H(B)$ of a complexity 1 (4HS or 1HT) band to be the number of elements of \mathcal{W} it contains. The height of a band of higher complexity is defined as above.

We shall say that two bands, $A, B \subseteq \Phi$ are *parallel* if they have the same base surface, $\pi_\Sigma A = \pi_\Sigma B$. We say that B is a parallel subband if also $B \subseteq A$. Note that in this case, the closures of $B \setminus A$ are parallel bands which we refer to as the *collars* of B (in A). We denote them by B_- and B_+ . We write $D(B, A) = \max\{H(B_-), H(B_+)\}$ for the (*combinatorial*) *depth* of B in A . We say that a band A is *maximal* if it is not contained in any larger parallel band. Every band B is contained in a unique maximal parallel band, $M(B)$. We write $D(B) = D(B, M(B))$. We say that B is *r-collared* if $D(B) \geq r$.

Lemma 4.4 : *Suppose that B, B' are 1-collared with base surfaces Φ and Φ' respectively. If $\text{int } B \cap \text{int } B' \neq \emptyset$, then Φ and Φ' are nested, i.e. either $\Phi \subseteq \Phi'$ or $\Phi' \subseteq \Phi$.*

Proof : Let $S \subseteq \Psi$ be a horizontal fibre through a point of $\text{int } B \cap \text{int } B'$. Now $S \cap B$ and $S \cap B'$ are fibres of B and B' respectively, and so the boundaries of Φ and Φ' cannot cross. If Φ and Φ' are not nested, then we can find curves $\alpha \subseteq \partial\Phi \setminus \partial\Phi'$ and $\alpha' \subseteq \partial\Phi' \setminus \partial\Phi$. Let $\Omega, \Omega' \in \mathcal{W}$ be the vertical annuli with $\pi_\Sigma \Omega$ and $\pi_\Sigma \Omega'$. Thus Ω and Ω' contain boundary components of the maximal bands, $M(B)$ and $M(B')$ respectively. By considering the horizontal projections of these bands to \mathbf{R} , we see easily that either Ω crosses one of the collars $M(B) \setminus B$ or Ω' crosses one of the collars $M(B') \setminus B'$. This contradicts the assumption that B and B' are 1-collared. \diamond

Let \mathcal{B} be the set of maximal bands in Ψ with $H(B) > 0$. Given $n \in \mathbf{N}$, If $A \in \mathcal{B}$ and $r \in \mathbf{N}$ with $H(A) \geq 2r + 1$. We can find a parallel subband, $B \subseteq A$, so that each of collars, B_\pm , have height $H(B_\pm)$ exactly r . For each $A \in \mathcal{B}$, and each such r , we choose such a band B , and write $\mathcal{B}(r)$ for the set of all bands that arise in this way for a fixed $r \in \mathbf{N}$.

Given a subsurface, Φ , of Σ , write $\mathcal{B}_\Phi(r) \subseteq \mathcal{B}(r)$ for those bands in $\mathcal{B}(r)$ whose base surface is a proper subsurface of Φ .

Let \mathcal{D} be the brick decomposition of Ψ described earlier. Given a subset $Q \subseteq \Psi$ we define the *size* of Q , denoted $\text{size}(Q)$, to be the number of bricks of \mathcal{D} whose interiors meet the interior of Q . We view $\text{size}(Q)$ as a combinatorial measure of the volume of Q .

Lemma 4.5 : *Given h, r, κ , there is some $\nu = \nu(h, r, \kappa)$ such that if B is a band with $H(B) \leq h$ and base surface Φ , then*

$$\text{size}(B \setminus \bigcup \mathcal{B}_\Phi(r)) \leq \nu(h, r, \kappa(\Phi)).$$

\diamond

Here we are allowing the case where $\Phi = \Sigma$, in which case, $\mathcal{B}_\Phi(r) = \mathcal{B}(r)$.

Proof : We proceed by induction on $\kappa(\Phi)$. If $\kappa(\Phi) = 1$, then we see explicitly that $\text{size}(B) \leq 3H(B) + 2$.

Now suppose that $\kappa(\Phi) \geq 2$. We can cut B into a set of $H(B) + 1$ parallel bands each of height 0. It is thus sufficient to deal with the case where $H(B) = 0$, in other words, some

vertical annulus of \mathcal{W} cuts through B . Now the set of all such annuli that cut through B cut B into set of bands of lower complexity. The number (possibly just 1) of such bands is bounded by $\kappa(\Phi)$. Let $A \subseteq B$ be such a band, and let $\Phi' = \pi_\Sigma A \subseteq \Phi$. Note that $\mathcal{B}_{\Phi'}(r) \subseteq \mathcal{B}_\Phi(r)$. If $H(A) \leq 2r$, then $\text{size}(A \setminus \bigcup \mathcal{B}_\Phi(r))$ is bounded (by $\nu(2r, r, \kappa(\Phi) - 1)$). If $H(A) \geq 2r + 1$, then $H(M(A)) \geq 2r + 1$, and so there is some $C \in \mathcal{B}(r)$ with base surface Φ' , so that each of the collars $M(A) \setminus C$ has height r . Now $C \in \mathcal{B}_\Phi(r)$ and $A \setminus C$ consists of at most two bands each of height at most r . Applying the inductive hypothesis again, we see that $\text{size}(A \setminus \bigcup \mathcal{B}_\Phi(r))$ is bounded (by $2\nu(r, r, \kappa(\Phi) - 1)$). Since there are at most $\kappa(\Phi)$ such bands A , this bounds $\text{size}(B \setminus \bigcup \mathcal{B}_\Phi(r))$, and the result follows by induction. \diamond

To each $\Omega \in \mathcal{W}$, we have associated a toroidal boundary component, $\Delta(\Omega)$, of $\Lambda(\mathcal{W})$ as described above. We can lift the brick decomposition, \mathcal{D} , to a brick decomposition of $\Lambda(\mathcal{W})$ which we also denote by \mathcal{D} . Let $\mathcal{D}(\Omega)$ be the set components of $D \cap \Delta(\Omega)$ which are annuli, as D runs over the set of bricks, \mathcal{D} . Thus $\Delta(\Omega) = \bigcup \mathcal{D}(\Omega)$ is a decomposition of this torus. We can view $|\mathcal{D}(\Omega)|$ as a combinatorial measure of its length.

Similarly suppose P is a non-compact boundary component of $\Lambda(\mathcal{W})$. This must be a bi-infinite cylinder, identified with a boundary component of Ψ . We get a decomposition of P into a collection $\mathcal{D}(P)$ of compact annuli, by taking the intersection with bricks.

We observed that we can recover Ψ up to homeomorphism by gluing a solid torus $T(\Omega)$ to each $\Delta(\Omega)$, so as to obtain the space $\Lambda(\mathcal{W}, \mathcal{W})$. We can describe this more explicitly as follows. Given $\Omega \in \mathcal{W}$, we choose an explicit homeomorphism of $T(\Omega) \setminus \partial_H \Omega$ with $S^1 \times [0, 1] \times (0, 1)$, and foliate $T(\Omega) \setminus \partial_H \Omega$ with annuli of the form $S^1 \times [0, 1] \times \{t\}$ for $t \in (0, 1)$. We set up the homeomorphism so that the two circles $S^1 \times \{0\} \times \{t\}$ and $S^1 \times \{1\} \times \{t\}$ are horizontal in $\Delta(\Omega)$ and get identified with the same horizontal circle in Ω , under the projection of $\Lambda(\Omega)$ to Ψ . We add in two degenerate leaves, $\partial_- \Omega$ and $\partial_+ \Omega$ to complete the foliation of $T(\Omega)$.

Now if S is a horizontal fibre in Ψ , then pulling back to $\Lambda(\mathcal{W})$ we get a union of ‘‘horizontal’’ surfaces. We can now use the foliations on the tori $T(\Omega)$ to complete this to a fibre of $\Lambda(\mathcal{W}, \mathcal{W})$. These fibres foliate $\Lambda(\mathcal{W}, \mathcal{W})$. We refer to them as *horizontal fibres*. We denote by $S(x) \subseteq \Lambda(\mathcal{W}, \mathcal{W})$ the fibre containing $x \in \Lambda(\mathcal{W}, \mathcal{W})$. There is a natural projection of $\Lambda(\mathcal{W}, \mathcal{W})$ to Ψ collapsing each torus $T(\Omega)$ to Λ , so that the fibres of $T(\Omega)$ are preimages of horizontal curves. The horizontal fibres of $\Lambda(\mathcal{W}, \mathcal{W})$ are preimages of horizontal fibres of Ψ . By a *band* in $\Lambda(\mathcal{W}, \mathcal{W})$ we mean the preimage of a band in $\Lambda(\mathcal{W}, \mathcal{W})$.

Suppose now that there is some $L \geq 0$ and a partition $\mathcal{W} = \mathcal{W}_0 \sqcup \mathcal{W}_1$ of \mathcal{W} such that $|\mathcal{D}(\Omega)| \leq L$ for all $\Omega \in \mathcal{W}_1$. (Such a situation will arise in Section 7 — see Theorem 7.1.) We write

$$\mathcal{T} = \{T(\Omega) \mid \Omega \in \mathcal{W}_0\}$$

,

$$\mathcal{T}_1 = \{T(\Omega) \mid \Omega \in \mathcal{W}_1\}.$$

Let

$$\Theta = \Lambda(\mathcal{W}, \mathcal{W}_1) = \Lambda(\mathcal{W}) \cup \bigcup \mathcal{T}.$$

(Thus Θ is homeomorphic to $\Lambda(\mathcal{W}_0)$.)

Note that Θ is made out of a collection \mathcal{D} of bricks and “tubes” \mathcal{T}_1 . We refer to the elements of $\mathcal{D} \cup \mathcal{T}_1$ collectively as the *building blocks* of Θ . Similarly, if R is a boundary component of Θ (either a cylinder or a torus, we refer to the elements of $\mathcal{D}(R)$ as the *building blocks* of R . If β is a path in Θ , or in R , we define the *combinatorial length* to be equal to the number of building blocks that it meets (counting multiplicities).

Note that (since we are assuming that $|\mathcal{D}(\Omega)| \leq L$ for all $\Omega \in \mathcal{W}_1$), each building block of Θ meets boundedly many others (in fact at most $\max\{8, L\}$). It follows that given any $x \in \Theta$ and any $r \in \mathbf{N}$, there is bound, depending on r , on the number of building blocks that can be connected to x by a path of combinatorial length at most r .

We want to make a couple of observations concerning the embedding of the boundary components of Θ into Θ . We begin with the non-compact components, since the description is somewhat simpler. The geometrical interpretation of these statements is made more apparent by Lemma 4.8, which will eventually be used in Section 11 (see Lemmas 11.8 and 11.9).

Lemma 4.6 : *Suppose that Π is a non-compact boundary component of Θ . Suppose that β is a path of combinatorial length n in Θ connecting two points, $x, y \in \Pi$, and homotopic into Π , relative to $\{x, y\}$. Then x and y are connected by a path in Π whose combinatorial length is bounded above by some uniform linear function of n .*

Proof : Let $C_x, C_y \in \mathcal{D}(\Pi)$ be annular blocks containing x and y respectively, and let $\mathcal{D}_{x,y} \subseteq \mathcal{D}(\Omega)$ be the set of annular blocks between C_x and C_y . We want to bound $|\mathcal{D}_{x,y}|$ linearly in terms of n .

Given $z \in \Pi$, recall that $S(z)$ is the horizontal fibre of $\Lambda(\mathcal{W}, \mathcal{W})$ containing z . Let $F(z)$ be the component of $S(z) \cap \Theta$ containing z . There is a bound, say l_0 , on the number of blocks that $F(z)$ can meet, depending only on $\kappa(\Sigma)$. Now if $z \in \bigcup \mathcal{D}_{x,y}$ we see that $F(z)$ must meet β (from the assumption that β is homotopic into Π). It follows that z is connected to β by a path in Θ of combinatorial length at most l_0 . By the earlier observation (on the uniform local finiteness of our system of building blocks) we see that this gives some bound on $|\mathcal{D}_{x,y}|$ in terms of n .

To make this a linear bound, let us fix our favourite positive integer, say 10, and let l_1 be the bound when n is at most $10 + 2l_0$. This means that if $z, w \in \Pi$ are separated by at least l_1 blocks, then if γ is any path from $F(z)$ to $F(w)$ in Θ , which can be homotoped into Π by sliding its endpoints along $F(z)$ and $F(w)$ respectively, then $l(\gamma) \geq 10$.

For the general case, we now choose a sequence of points $x = z_0, z_1, \dots, z_p = y$ in $\bigcup \mathcal{D}_{x,y}$ so that z_i and z_{i+1} are separated by at least l_1 annular blocks in $\mathcal{D}(\Omega)$, and with $|\mathcal{D}_{x,y}|$ bounded above by a fixed linear function of p . Now the path β must cross each of the surfaces $F(z_i)$ and so p is in turn bounded above by a linear function of $l(\beta)$. \diamond

We need a version of this where $\Delta = \Delta(\Omega)$ for $\Omega \in \mathcal{W}_0$ is a toroidal boundary component.

Lemma 4.7 : *Suppose that Δ is a compact boundary component of Θ . Suppose that β is a path of combinatorial length n in Θ connecting two points, $x, y \in \Delta$. If β is homotopic to a path in Δ , relative to $\{x, y\}$, then we can find such a path in Δ in this relative homotopy*

class whose combinatorial length is bounded above by some uniform linear function of n .

Proof : The argument is a slight refinement of that used for Lemma 4.6. If $z \in \Delta$, we can define the surface $F(z)$ exactly as in Lemma 4.6.

Note that $\partial_H \Omega$ cuts $\Delta = \Delta(\Omega)$ into annuli, A and A' , say. Suppose first that $x, y \in A$ and that β is homotopic into A relative to x, y . Essentially the same argument as before gives a bound on the number of building blocks separating x and y in A . We can therefore construct surfaces $F(z_i)$ as before so as to obtain a linear bound in terms of n .

The general case is complicated by the fact that β might wrap around Δ many times in the vertical direction (that of a meridian curve). However, we can construct surfaces, $F(z_i)$, on both sides of Δ (the annuli A and A'), and we note that β must cross all of these surfaces in sequence (counting multiplicities). We should note that nothing we have said excludes the possibility that the total vertical length, $\Delta(\Omega)$, is small, (maybe smaller than l_0 , for example) so we need to take at least one such surface. \diamond

We next want to translate some of these combinatorial observations into more geometrical terms. To this end, we shall put a riemannian metric, d , on $\Lambda(\mathcal{W}, \mathcal{W})$. The construction of d will be explained more carefully in Section 7, where we construct the model space. For the purposes of this section, we only care about the metric restricted to Θ . The key points (which can be taken as hypotheses for the moment) are as follows.

The local geometry of the decomposition of (Θ, d) is bounded. In particular, there is a uniform lower bound on the injectivity radius of (Θ, d) . Each building block of Θ (in $\mathcal{D} \cup \mathcal{T}_1$) has bounded diameter. Moreover, there is a positive lower bound on the d -distance between any two disjoint building blocks. We can also assume that each of the building blocks of any boundary component of Θ is a fixed isometry class of annulus, say $S^1 \times [0, 1]$. If $x \in \Theta$, the fibre $S(x)$ meets each block of Θ in a surface of bounded diameter. In particular, the diameter of each component of $S(x) \cap \Theta$ is bounded.

Note that an immediate consequence is that, at least up to homotopy, the combinatorial length of a path in Θ is bounded above by a linear function of its d -length. Since the building blocks are not simply connected, we do not have a converse statement keeping control of homotopy. However, distances between points, and diameters of sets, in the metric d are bounded above by a linear function of their combinatorial counterparts.

We can immediately translate Lemmas 4.6 and 4.7 into geometrical terms and express them in a unified fashion as follows.

Suppose that R is a boundary component of Θ , and let $H \subseteq \pi_1(\Theta)$ be the subgroup generated by a horizontal longitude. (Thus $H \equiv \pi_1(R)$ if R is a bi-infinite cylinder). Let $\hat{\Theta}$ be the cover corresponding to H . Thus R lifts to a bi-infinite cylinder, $\hat{R} \subseteq \hat{\Theta}$ (so that $\hat{R} \equiv R$ in the non-compact case).

Lemma 4.8 : *If R is a boundary component of Θ , and $\hat{R} \subseteq \hat{\Theta}$ constructed as above, then \hat{R} is quasi-isometrically embedded in $\hat{\Theta}$.* \diamond

For applications, we will need another riemannian metric, ρ , on $\Lambda(\mathcal{W}, \mathcal{W})$, and its restriction to Θ . This can be taken to be equal to d on Θ and equal to 0 on each tube $T \in \mathcal{T}$. In other words we force each element of T to have diameter 0. As stated, we just

get a pseudometric, and it will be discontinuous at the toroidal boundary components. If we want, we can smooth it out in a small neighbourhood of these boundaries. The only important requirement is that each element of \mathcal{T} should have bounded diameter with respect to ρ .

We can list some geometric properties of Θ as follows:

(W1) Every point $x \in \Theta$ lies in a fibre $S(x)$ of uniformly bounded ρ -diameter.

We can simply take $S(x)$ to be the horizontal fibre as described above. In this way, $S(x)$, will vary continuously in x . (This will be used in Section 9.) Alternatively, we can push the fibre off each torus of \mathcal{T}_0 so as to give us a surface, $S(x)$, in Θ , while retaining a bound on its ρ -diameter in Θ . (Here we are referring to the extrinsic diameter in Θ , and not the induced path-metric in Θ , which may be arbitrarily large.) In this case, however, we can no longer assume that $S(x)$ varies continuously in x . (This alternative construction will be useful in Section 10.)

(W2) Each $x \in \Theta$ is contained in a loop $\gamma_x \subseteq \Theta$ of bounded d -length, and homotopic to a curve, $[\gamma_x] \in X(\Sigma)$. If x lies in a boundary component, R , of Θ , then we can take γ_x to be the horizontal curve in R containing x .

In fact, if $x \in D \in \mathcal{D}$, we take γ_x to be homotopic to one of the vertical boundary components of D . Thus, γ_x is homotopic to an annulus $\Omega_x \in \mathcal{W}$.

(W3) If $x, y \in \Theta$, with $d(x, y) \leq \eta$, then $d_{\mathcal{G}(\Sigma)}([\gamma_x], [\gamma_y])$ is bounded above.

Here we can take $\eta > 0$ to be the lower bound on injectivity radius, but any positive constant would do.

(W4) If $x, y \in \Theta$, then $\rho(x, y)$ is bounded above by a uniform linear function of $d_{\mathcal{G}(\Sigma)}([\gamma_x], [\gamma_y])$. ■

This is our geometric interpretation of the tautness condition (Theorem 4.1(3)) in the case where the base surface is Σ . Note that γ_x and γ_y are homotopic to Ω_x and Ω_y and a bounded d -distance from the corresponding tubes $T(\Omega_x)$ and $T(\Omega_y)$ (these tubes might lie in either \mathcal{T} or \mathcal{T}_1). By tautness, there is a ladder, $\Omega_x = \Omega_0, \Omega_1, \dots, \Omega_n = \Omega_y$ in \mathcal{W} , with n bounded above by a linear function of $d_{\mathcal{G}(\Sigma)}([\gamma_x], [\gamma_y])$ (in fact, $n \leq d_{\mathcal{G}(\Sigma)}([\gamma_x], [\gamma_y]) + c$, where c is the tautness constant). Two consecutive Ω_i and Ω_{i+1} meet a horizontal fibre which has bounded ρ -diameter. We see that $\rho(T(\Omega_i), T(\Omega_{i+1}))$ is bounded above, and so $\rho(x, y)$ is linearly bounded in terms of n and hence in terms of $d_{\mathcal{G}(\Sigma)}([\gamma_x], [\gamma_y])$ as claimed. (One needs to rephrase this slightly if Σ is 4HS or 1HT, but the argument is essentially the same — consecutive tubes are a bounded distance apart, since they meet a common building block. Here we use the metric on the modified curve graph $\mathcal{G}'(\Sigma)$.)

All these statements have analogues for the case of a band $B \subseteq \Lambda(\mathcal{W}, \mathcal{W})$. We are only really interested in the case where all vertical boundary components of B lie in \mathcal{W}_0 . In this case, $\partial_H B$ is the relative boundary of B in Θ . Let Φ be base surface of B .

Let \mathcal{T}_B^0 be the set of tubes in \mathcal{T} whose interiors meet B . We let d_B be the riemannian metric on B induced from d , and let ρ_B be the metric obtained from d_B by forcing each set $T \cap B$ for $T \in \mathcal{T}_B^0$ to have diameter 0. We can now perform the above constructions inside B . We get:

(W5) If $x \in B \cap \Theta$, then x lies in a fibre, $F(x) \subseteq B$ of B , of uniformly bounded ρ_B -diameter.

(W6) If $x \in B \cap \Theta$, then x lies in a loop $\gamma_x^B \subseteq B \cap \Theta$ of bounded d_B -length, with $[\gamma_x^B] \subseteq X(\Phi) \subseteq X(\Sigma)$. If $x \in B \cap \partial\Theta$ we can take γ_x^B to be a horizontal curve of that boundary component.

(W7) If $x, y \in B \cap \Theta$, and $d_B(x, y) \leq \eta$, then $d_{\mathcal{G}(\Phi)}([\gamma_x^B], [\gamma_y^B])$ is bounded above.

(W8) If $x, y \in B \cap \Theta$, the $\rho_B(x, y)$ is bounded above by a uniform linear function of $d_{\mathcal{G}(\Phi)}([\gamma_x^B], [\gamma_y^B])$.

Finally we need to be able to recognise when a point of Θ does not lie in a given maximal band. We can define a *maximal band* in $\Lambda(\mathcal{W}, \mathcal{W})$ to be the preimage of a maximal band in Ψ .

(W9) If $x \in \Theta \setminus B$, then x lies in a loop $\delta_x^B \subseteq B \cap \Theta$ of bounded d_B -length such that either $[\delta_x^B]$ is homotopic to a torus $T \in \mathcal{T}$ not lying entirely inside B , or else $[\delta_x^B]$ is not homotopic into Φ .

To see this, let $S(x)$ be the horizontal fibre through x , and let F be the component of $S(x) \cap \Theta$ containing x . This has bounded d -diameter. If $\pi_\Sigma F$ is not a subsurface of Φ , then we can choose $\delta_x^B \subseteq F$, not homotopic into Φ . If $\pi_\Sigma F \subseteq \Phi$, then by the maximality of B , at least one of the boundary components of F must be a torus $T \in \mathcal{T}$, and this is not contained in B . We can take $\delta_x^B \subseteq F$, freely homotopic to this boundary component.

5. Margulis tubes.

The constructions described in Section 4 are the basis of the model of the “thick part”. To complete the picture we will need some description of the “thin part”. There may be parabolic cusps, but the main thing we have to worry about is the existence of Margulis tubes.

Recall that *quasi-isometry* between two geodesic spaces, (X, d) and (X', d') is a map $f : X \rightarrow X'$, for which $c_1 > 0, c_2, c_3, c_4, c_5$ with $c_1 d(x, y) - c_2 \leq d'(f(x), f(y)) \leq c_3 d(x, y) + c_4$ for all $x, y \in X$, and with $X' \subseteq N(f(X), c_5)$. Here we shall make the following definition:

Definition : A *sesquilipschitz* map is a surjective lipschitz quasi-isometry.

In other words, in the definition of quasi-isometry we put $c_4 = c_5 = 0$.

Definition : A *universally sesquilipschitz* map between two spaces is a homotopy equivalence such that the lift to the universal covers is sesquilipschitz.

One can check that a universally sesquilipschitz map is indeed sesquilipschitz.

Throughout this section our results refer to implicitly assumed constants. We will take it as implied that the constants outputted are explicit functions of the constants inputted, though we will not bother to calculate these function explicitly. (Here they are all computable.) We shall use the adjective “uniform” if we want to stress this point.

We shall begin our discussion in dimension 1. It is well known that any quasi-isometry of the real line \mathbf{R} is a bounded distance from a bilipschitz homeomorphism. We note the following variation:

Lemma 5.1 : *Let $f : \mathbf{R} \longrightarrow \mathbf{R}$ is sesquilipschitz then there is a sesquilipschitz homotopy to a bilipschitz map.*

In other words there is a sesquilipschitz map $F : \mathbf{R} \times [0, 1] \longrightarrow \mathbf{R}$ with $f = [x \mapsto F(x, 0)]$ and with $g = [x \mapsto F(x, 1)]$ bilipschitz. Note that it follows that f is a bounded distance from g .

Proof : We can assume that f is end-preserving. We fix some sufficiently large, but bounded, constant, $k \geq 0$, so that $f(x+k) > f(x) + 1$ for all $x \in \mathbf{R}$. We set $g|_{k\mathbf{Z}} = f|_{k\mathbf{Z}}$, and interpolate linearly. We then take a linear homotopy between f and g . \diamond

A similar argument can be carried out equivariantly. We write $S(r) = \mathbf{R}/r\mathbf{Z}$ for the circle of length r . We obtain:

Lemma 5.2 : *Suppose that $r, s > 0$, and that $f : S(r) \longrightarrow S(s)$ is a universally sesquilipschitz map. Then there is a universally sesquilipschitz homotopy from f to a bilipschitz map from $S(r)$ to $S(s)$.* \diamond

In particular, the ratios s/r and r/s are bounded.

More will be said about 1-dimensional quasi-isometries in Section 9, but this will do for the moment. We move on to 2 dimensions.

Let Δ be a euclidean torus equipped with a preferred basis, (l, m_0) for the integral first homology. We refer to l as the *longitude* of Δ and to m_0 as the *standard meridian*. More generally a *meridian* will be a curve of the form $l + nm_0$ for some $n \in \mathbf{Z}$. In situations of interest to us the length of the longitude will be bounded both above and below, and so it is often convenient to normalise so that its length is 1. In this case the structure on Δ is determined by a complex *modulus* $\lambda \in \mathbf{C}$ with $\Im(\lambda) > 0$, so that $\Delta = \Delta(\lambda) = \mathbf{C}/\langle [z \mapsto z + 1], [z \mapsto z + \lambda] \rangle$, with $[z \mapsto z + \lambda]$ giving us the standard meridian. Note that the shortest meridian in Δ has length between $\Im(\lambda)$ and $\Im(\lambda) + \frac{1}{2}$.

We refer to a geodesic longitude as being *horizontal*: it is the projection of a line parallel to the real axis. These foliate Δ and we write $S(\Delta)$ for the leaf space obtained by collapsing each leaf to a point. It is a circle of length $\Im(\lambda)$.

In most cases of interest, the injectivity radius of $\Delta(\lambda)$ will be bounded below by some positive constant. One can see that this is equivalent to putting a lower bound on $\Im(\lambda)$. Moreover if there is an equivariant quasi-isometry between two such tori, a lower bound on the injectivity radius of one gives a lower bound for the other.

Let $\Delta = \Delta(\lambda)$ and $\Delta' = \Delta(\lambda')$.

Definition : A map $f : \Delta \longrightarrow \Delta'$ is *horizontally straight* if it sends each horizontal longitude of Δ isometrically to a horizontal longitude of Δ' .

In formulae, this means that, writing $\tilde{\Delta} = \mathbf{R}^2 = \tilde{\Delta}'$, we have $f(x, y) = (x + f_H(y), f_V(y))$,

where $f_H, f_V : \mathbf{R} \rightarrow \mathbf{R}$ satisfy $f_H(y + \mathfrak{S}(\lambda)) = f_H(y) + \mathfrak{R}(\lambda') - \mathfrak{R}(\lambda)$ and $f_V(y + \mathfrak{S}(\lambda)) = f_H(y) + \mathfrak{S}(\lambda')$.

Note that such a map induces a map, $S(f) : S(\Delta) \rightarrow S(\Delta')$ (lifting to $f_V : \mathbf{R} \rightarrow \mathbf{R}$). One can easily check that if f is lipschitz (respectively, sesquilipschitz, universally sesquilipschitz) then so is $S(f)$.

Lemma 5.3 : *Suppose $\mathfrak{S}(\lambda) \geq \epsilon > 0$. Suppose $f : \Delta \rightarrow \Delta'$ is a lipschitz map sending the longitude (homotopically) to the longitude. Then there is a lipschitz homotopy of f to a (lipschitz) horizontally straight map. Moreover if f is (universally) sesquilipschitz, we can take the homotopy to be (universally) sesquilipschitz. Here the constants only depend on ϵ and the initial (sesqui)lipschitz constants.*

Proof : Let m be a shortest meridian on Δ . The lower bound on $\mathfrak{S}(\lambda)$ means that there is a lower bound on its slope with respect to any horizontal longitude. We now define $g : \Delta \rightarrow \Delta'$ by taking $g|m = f|m$, and extending in the unique way to a horizontally straight map. We now take a linear homotopy between f and g . The above properties are easily verified. \diamond

Lemma 5.4 : *Suppose that $f : \Delta \rightarrow \Delta'$ is a universally sesquilipschitz horizontally straight map. Then there is a universally sesquilipschitz homotopy from f to a horizontally straight bilipschitz homeomorphism.*

Proof : By lemma 5.2 there is a universally sesquilipschitz homotopy F of $S(f)$ to a bilipschitz map. Lifting to \mathbf{R} gives us a homotopy $\tilde{F} : \mathbf{R} \times [0, 1] \rightarrow \mathbf{R}$ from $f_V : \mathbf{R} \rightarrow \mathbf{R}$ to a bilipschitz map h . Now define $\tilde{G} : \mathbf{R}^2 \times [0, 1] \rightarrow \mathbf{R}^2$ by $\tilde{G}(x, y, t) = (x + f_H(y), \tilde{F}(y, t))$. Projecting back down gives us a bilipschitz map, $g : \Delta \rightarrow \Delta$, with $g_H = f_H$ and $g_V = h$. \diamond

We have assumed that f lifts to a quasi-isometry of universal covers $\tilde{\Delta} \rightarrow \tilde{\Delta}'$. However, we only really require that it is a quasi-isometry for the covers corresponding to the longitudes of Δ and Δ' , which are bi-infinite cylinders. This is enough to show that f_V is a quasi-isometry. If we assume that f is lipschitz, then the same conclusion holds.

We remark that the existence of a k -bilipschitz map from $\Delta(\lambda)$ to $\Delta(\lambda')$ implies both that $k^{-1}\mathfrak{S}(\lambda) \leq \mathfrak{S}(\lambda') \leq k\mathfrak{S}(\lambda)$ and $k^{-1}|\lambda| \leq |\lambda'| \leq k|\lambda|$.

We need also to consider lipschitz maps to the circle. Suppose that $f : \Delta \rightarrow S(1)$ is k -lipschitz. Let m be a shortest meridian on Δ . It has length at most $\mathfrak{S}(\lambda) + \frac{1}{2}$, and so its image has length at most $k(\mathfrak{S}(\lambda) + \frac{1}{2})$. Thus the degree of $f|m$ is at most $k(\mathfrak{S}(\lambda) + \frac{1}{2})$ in absolute value.

Lemma 5.5 : *Suppose the torus $\Delta(\lambda)$ admits a k -lipschitz map to $S(1)$ which has degree 1 on the longitude and degree 0 on the standard meridian. Then $|\lambda| \leq (k + 1)(\mathfrak{S}(\lambda) + \frac{1}{2})$.*

Proof : Let l be the longitude, and m_0 and m be the standard and shortest meridians respectively. Thus, $m = m_0 + pl$ for some $p \in \mathbf{Z}$, and so $\deg(f|m) = \deg(f|m_0) +$

$p \deg(f|l) = p$, where f is the k -lipschitz map. By the above observation, $|p| \leq k(\mathfrak{S}(\lambda) + \frac{1}{2})$.
Now

$$\begin{aligned} |\lambda| &= \text{length}(m_0) \\ &\leq \text{length}(m) + |p| \\ &\leq (\mathfrak{S}(\lambda) + \frac{1}{2}) + k(\mathfrak{S}(\lambda) + \frac{1}{2}) \\ &= (k+1)(\mathfrak{S}(\lambda) + \frac{1}{2}). \end{aligned}$$

◇

Lemma 5.6 : *Given $c, k > 0$, there is some $h > 0$ such that if a map $f : \Delta(\lambda) \rightarrow S(1)$ is k -lipschitz and degree 1 on the longitude and degree 0 on the meridian, and if $\mathfrak{S}(\lambda) \leq c$, then there is a h -lipschitz homotopy of f to a k -lipschitz map, g , sending every geodesic standard meridian to a point.*

Proof : Let l be some horizontal longitude. There is a unique map $g : \Delta(\lambda) \rightarrow S(1)$ so that $g|l = f|l$ and sending every standard longitude to a point. This is also k -lipschitz. Clearly f and g are homotopic, and we take a linear homotopy between them. To bound its lipschitz constant, it is enough to note that every geodesic standard meridian of $\Delta(\lambda)$ gets mapped under f to curve of length at most $k|\lambda| \leq k(k+1)(c + \frac{1}{2})$ by Lemma 5.6. ◇

We now move into 3-dimensions to consider Margulis tubes. For the purposes of this section, a ‘‘Margulis tube’’ is just a particular kind on hyperbolic structure on the solid torus.

Let $r \geq 0$ and $R = 2\pi \sinh r$. Given $t \geq 0$, set $a_r(t) = \cosh(rt)/\cosh r$ and $t \geq 0$, set $b_r(t) = \sinh(rt)$. Define a riemannian metric on $\mathbf{R} \times S^1 \times (0, 1]$ by $ds^2 = a_r(t)^2 dx^2 + b_r(t)^2 dy^2 + r^2 dt^2$, where (x, y, t) are the local coordinates. Let N be the metric completion of this space.

What we have defined is just the r -neighbourhood, $N = N(\tilde{\alpha})$ of a bi-infinite geodesic, $\tilde{\alpha}$ in \mathbf{H}^3 . Its boundary, $\partial N = \mathbf{R} \times S^1 \times \{1\}$ is isometric to $\mathbf{R} \times S(R)$, by isometry that is the identity on the first co-ordinate.

A loxodromic isometry, g , with axis $\tilde{\alpha}$ acts by translating the x -coordinate and rotating the y -coordinate. We refer to the quotient, $T = N/\langle g \rangle$ as a *Margulis tube*. Thus, ∂T is a euclidean torus. The quotient, $\alpha = \tilde{\alpha}/\langle g \rangle$ is the *core* of T . We refer to $r = d(\alpha, \partial T)$ as the *depth* of T . We define the *standard meridian*, m_0 , of ∂T , to be homotopically trivial in T . It has length R . We deem another curve, l , in ∂T , homotopic to the core curve, α , to be the *longitude* of ∂T . (Its homotopy class gives us some additional structure to T .) Note that the length of the core curve is equal to $\text{area}(\partial T)/R \cosh r = 2\pi \text{area}(\partial T)/R\sqrt{R^2 + 4\pi^2}$.

The following seems well known:

Lemma 5.7 : *Given any euclidean torus Δ with preferred longitude and standard meridian, there is a unique Margulis tube, T , (up to isometry—) with $\partial T = \Delta$.*

Proof : The cover of Δ corresponding to the standard meridian is isometric to $\mathbf{R} \times S(R)$ for some $R > 0$. Set $r = \sinh^{-1}(R/2\pi)$ and construct N as above. The action of the longitude on $\partial N = \mathbf{R} \times S(R)$ extends to a loxodromic on N , and we take the quotient.

Uniqueness is easily established (see the remark after Lemma 5.8 below). \diamond

Lemma 5.8 : *Suppose T, T' are Margulis tubes and $f : \partial T \rightarrow \partial T'$ is a k -bilipschitz map sending the standard meridian of ∂T (homotopically) to the standard meridian of $\partial T'$. Then f extends to a k' -bilipschitz map, $f : T \rightarrow T'$, where k' depends only on k .*

Proof : Let r, s be the depths of T, T' respectively. The lengths of the standard meridians are $2\pi \sinh r$ and $2\pi \sinh s$. Their ratios are bounded by k . This also gives bounds on the ratios of r and s and of $\cosh r$ and $\cosh s$. For $t \in (0, 1]$, $a_r(t)/a_s(t)$ varies between $\cosh s/\cosh r$ and 1 and $b_r(t)/b_s(t)$ varies between r/s and $\sinh r/\sinh s$. We can thus define a bilipschitz map $[(x, t) \mapsto (f(x), t)] : \partial T \times (0, 1] \rightarrow \partial T' \times (0, 1]$ and extend over the completions. \diamond

We remark that in the case where $k = 1$, we can take $k' = 1$, giving the uniqueness part of Lemma 5.7.

The cases of interest to us the length of the longitude will be bounded above and below, and so, up to bilipschitz equivalence we can normalise so that it has length 1. (This is really just for notational convenience.) In this case, we can identify ∂T with $\Delta(\lambda)$ for some modulus $\lambda \in \mathbf{C}$. Note that $|\lambda| = R$ and that $\text{area } \partial T = \Im(\lambda)$. We see that the length of the core curve is $L(\lambda) = 2\pi \Im(\lambda)/|\lambda| \sqrt{|\lambda|^2 + 4\pi^2}$.

We earlier defined the leaf space, $S(\partial T)$ of ∂T by collapsing each geodesic longitude to a point. Its length is $\Im(\lambda)$. We can define another leaf space, $S_0(\partial T)$ by collapsing each geodesic standard meridian to a point. It has length $\text{area}(\partial T)/R = \Im(\lambda)/|\lambda|$. There is a natural linear homeomorphism from $S_0(\partial T)$ to the core curve, α , given by orthogonal projection. It contracts distances by a factor $L(\lambda)/(\Im(\lambda)/|\lambda|) = |\lambda|L(\lambda)/\Im(\lambda) = 2\pi/\sqrt{|\lambda|^2 + 4\pi^2}$. By precomposing the inverse of this projection with the projection of T to the core curve, we see that the projection $\partial T \rightarrow S_0(\Delta)$ extends to a $(\sqrt{|\lambda|^2 + 4\pi^2}/2\pi)$ -lipschitz map $T \rightarrow S_0(\Delta)$.

Lemma 5.9 : *Suppose that T is a Margulis tube where the longitude of ∂T has length 1. Suppose that ∂T has area at most c . Let $f : \partial T \rightarrow S(1)$ be a k -lipschitz map which has degree 1 on the longitude and degree 0 on the standard meridian. Then f extends to a k' -lipschitz map $f : T \rightarrow S(1)$, where k' depends only on k and c .*

Proof : Using Lemma 5.6, we can reduce to the case where f sends each geodesic meridian to a point — since the lipschitz homotopy given by Lemma 5.6 can be carried out in a uniformly small neighbourhood of ∂T in T .

We can thus assume we have a k -lipschitz map $f : \partial T \rightarrow S(1)$ which factors through the projection $\partial T \rightarrow S_0(\partial T)$. But the latter projection extends to a $(\sqrt{|\lambda|^2 + 4\pi^2}/2\pi)$ -lipschitz map $T \rightarrow S_0(\partial T)$. Composing this gives us a $k(\sqrt{|\lambda|^2 + 4\pi^2}/2\pi)$ -lipschitz map $T \rightarrow S(1)$. Finally, we note that, by Lemma 5.6, we have $|\lambda| \leq (k + 1)(\Im(\lambda) + \frac{1}{2}) \leq$

$(k + 1)(c + \frac{1}{2})$.

◇

6. Systems of convex sets.

The purpose of this section is to describe some constructions of lipschitz maps which we will apply in Section 7 to get a lipschitz map from the thick part of our model space into the thick part of our 3-manifold. Since the actual set-up in which we are interested is somewhat complicated to describe, we will present most of it in the fairly general setting of systems of convex sets, only adding assumptions as we need them. The main application we have in mind here is described by Lemma 6.8, whose explicit hypotheses are laid out before its statement. In practice all the convex sets we deal with will be either horoballs or uniform neighbourhoods of geodesics, which are lifts of closed geodesics or Margulis regions in our 3-manifold. The domain for our map will be a locally finite polyhedral complex, which in applications arises out of the combinatorial construction of the model space. We note that any two sensible path metrics on such a space will be bilipschitz equivalent, and so the actual choice doesn't much matter to us. We will describe a specific metric for definiteness. Much of the argument will apply in any dimension, though we restrict our attention to 3.

Let Π be a 3-dimensional simplicial complex with vertex set Π^0 . We write Π^i for the set of i -simplices. We assume Π^i to be locally finite away from Π^0 . We write $|\Pi|$ for its realisation. We are really interested in a *truncated realisation* of Π , denoted $R(\Pi)$, built out of truncated simplices. We can construct a truncated simplex by taking a regular euclidean simplex of side length 3, and removing a regular simplex of side length 1 about each vertex. The resulting polyhedron has all side lengths 1. In dimension 2, for example, we get a regular hexagon. Gluing these together we get a locally finite polyhedral complex, $R(\Pi)$, which we can view as a closed subset of $|\Pi|$. Associated to each $x \in \Pi^0$, we have a polyhedral subset of $D(x) \subseteq R(\Pi)$ — the boundary of a neighbourhood of x in $|\Pi|$. This is a 2-dimensional simplicial complex.

Given two convex subsets $P, Q \subseteq \mathbf{H}^3$ we write $\text{par}(P, Q) = \text{diam}(N(P, 1) \cap N(Q, 1))$. (This is 0 if the intersection is empty.) We view this as a convenient measure of the extent to which P and Q remain close (or “parallel”). An upper bound in $\text{par}(P, Q)$ means that they must diverge uniformly.

We start with a fairly simple construction that will be refined later. Suppose to each $x \in \Pi^0$ we associate a closed convex set $Q(x) \subseteq \mathbf{H}^3$. We assume:

(A1) If $x, y \in \Pi^0$ are distinct, then $Q(x) \cap Q(y)$ consists of at most one common boundary point, and $\text{par}(Q(x), Q(y))$ is bounded above.

(B1) If $xy \in \Pi^1$, then $d(Q(x), Q(y))$ is bounded above.

(We will later look into other hypothesis that imply (B1).)

Given $xy \in \Pi^1$, let $\beta(xy)$ be the shortest geodesic from $Q(x)$ to $Q(y)$. A key observation is, that is a simple exercise in hyperbolic geometry is the following:

Lemma 6.1 : *If $xyz \in \Pi^2$, then $\text{diam}(\beta(xy) \cup \beta(yz) \cup \beta(zx))$ is bounded.*

Proof : Consider the hexagonal path $\beta(xy) \cup \alpha(y) \cup \beta(yz) \cup \alpha(z) \cup \beta(zx) \cup \alpha(x)$ with $\alpha(x) \subseteq Q(x)$ etc. all geodesics. The lengths of the β paths are bounded. So, if the hexagon were very long, two of the α paths would have to run close together over a long distance, contradicting (A1). \diamond

Lemma 6.2 : *With the above hypotheses ((A1) and (B1)) there is a canonical uniformly bilipschitz map $\phi : R(\Pi) \rightarrow \mathbf{H}^3$ such that $\phi(D(x)) \subseteq Q(x)$. for all $x \in \Pi^0$, and such that if $a \in Q(x) \cap \phi(R(\Pi))$ then $d(a, \partial Q)$ is uniformly bounded above.*

Note that every point of $\phi(R(\Pi))$ lies a bounded distance from two distinct sets $Q(x)$, from which it follows that this image can only boundedly enter any such convex set. Explicitly we note:

(*) There is some constant $K \geq 0$ depending only on the constants of (A1) and (B1) such that for all $x \in \Pi^0$, and such that if $a \in Q(x) \cap \phi(R(\Pi))$ then $d(a, \partial Q) \leq K$.

Our construction of ϕ is as follows. For any $xy \in \Pi^1$, length $\beta(xy) = d(Q(x), Q(y))$ is bounded. Let $p(x, y) = \beta(xy) \cap Q(x) \in \partial Q(x)$ be the nearest point in $Q(x)$ to $Q(y)$. We map the corresponding edge of $R(\Pi)$ linearly to $\beta(xy)$. By this process, we will map the vertex set of each $D(x)$ into $\partial Q(x)$. We now extend linearly over $\phi(D(x))$. By convexity, $\phi(D(x)) \subseteq Q(x)$. Applying Lemma 6.1, we see that the images of simplices in $D(x)$ are bounded.

The *centre* of a finite diameter subset, $B \subseteq \mathbf{H}^3$, can be defined as the unique point $c \in \mathbf{H}^3$ such that $B \subseteq N(c, r)$ with r minimal. Given $xyz \in \Pi^2$, write $c(xyz)$ for the centre of $\beta(xy) \cup \beta(yz) \cup \beta(zx)$. Associated to xyz , we have a hexagonal 2-cell in $R(\Pi)$, and we have already defined ϕ on its boundary. We now extend over the interior by sending its centre to $c(xyz)$ and coning linearly over the boundary.

Similarly, given $xyzw \in \Pi^3$, we let $c(xyzw)$ be the centre of $\beta(xy) \cup \beta(yz) \cup \beta(zx) \cup \beta(xw) \cap \beta(yw) \cup \beta(zw)$. We have already defined ϕ on the boundary of the associated 3-cell of $R(\Pi)$ and now cone linearly over the centre $c(xyzw)$.

This gives us our lipschitz map ϕ , proving Lemma 6.2.

We want to refine the above construction to push ϕ off the interiors of convex sets. For this we need some additional assumptions.

Suppose that $A \subseteq \mathbf{H}^3$ is convex, and that $Q = N(A, t)$ for some $t \geq 0$. Given any $r \in (0, t)$ write $Q_r = N(A, t - r)$. Thus $Q_0 = Q$. We can define a projection $\pi : Q_r \rightarrow \partial Q$ so that each $a \in Q_r$ lies on the shortest geodesics from $\pi(a)$ to A . This projection is $(\sinh t / \sinh(t - r))$ -lipschitz.

We can refine Lemma 6.2 as follows. Suppose that to each $x \in \Pi^0$ we have associated some convex set, $A(x)$ and some $t(x) \geq 0$. Let $Q(x) = N(A(x), t(x))$. We suppose that the collection $(Q(x))_{x \in \Pi^0}$ satisfies the assumptions (A1) and (B2). Suppose that $t_0 \geq K + 1$, where K is the constant of (*) above, and suppose that $t_1 \geq t_0$.

Now suppose that we decompose Π^0 into two subsets, $\Pi_0^0 \sqcup \Pi_1^0$ satisfying:

For all $x \in \Pi_0^0$, $t(x) \leq t_0$, and

For all $x \in \Pi_1^0$, $t(x) \geq t_1$.

Lemma 6.3 : *With the above hypotheses, we can find a canonical uniformly lipschitz map, $\phi : R(\Pi) \longrightarrow \mathbf{H}^3$ such that*

- (1) if $x \in \Pi_1^0$, then $\phi(D(x)) \subseteq A(x)$,
- (2) if $x \in \Pi_0^0$, then $\phi(D(x)) \subseteq \partial Q(x)$, and
- (3) if $x \in \Pi_0^0$, then $Q(x) \cap \phi(R(\Pi)) \subseteq \partial Q(x)$.

Proof : We start with a map $\psi : R(\Pi) \longrightarrow \mathbf{H}^3$ as given by Lemma 6.3. If $x \in \Pi_0^0$, then $Q(x) \cap \phi(R(\Pi)) \subseteq Q_{t_0}(x)$. By composing with the outward projection $\pi : Q_{t_0}(x) \longrightarrow \partial Q(x)$ described above, we can push the image off the interior of $Q(x)$, while maintaining a control on the lipschitz constant.

If $x \in \Pi_1^0$, we have $\phi(D(x)) \subseteq Q(x) \subseteq N(A(x), t_1)$. We can now project $\phi(D(x))$ to $A(x)$ by nearest point projection, and extending by linear homotopy carried out in a uniformly small neighbourhood of $D(x)$ in $R(\Pi)$. In this way we can arrange that $\phi(D(x)) \subseteq A(x)$, again maintaining control over the lipschitz constant. \diamond

In applying these results, we will start from slightly different hypotheses. We suppose we have convex sets, $(Q(x))_x$, satisfying (A1), but we do not a-priori assume (B1). This we will need to deduce.

We begin by assuming:

(A2) If $xy \in \Pi^1$, then there is an involution $\sigma(xy)$ of \mathbf{H}^3 preserving both $Q(x)$ and $Q(y)$ setwise.

Unless these are both points (a situation of no interest to us) such an involution will be unique. Indeed if $Q(x) \cap Q(y) = \emptyset$ then it is rotation by π about the geodesic $\beta(xy)$.

Suppose now that $xyz \in \Pi^2$. Set $g(x, y, z) = \sigma(xz)\sigma(xy)$. Thus $g(x, y, z)Q(x) = Q(x)$ and $g(x, y, z)g(y, z, x)g(z, x, y) = 1$.

Lemma 6.4 : *Suppose $xyz \in \Pi^2$ and for all $a \in \partial Q(x)$, $d(a, g(x, y, z)a) \geq \epsilon > 0$. Then $d(\beta(xy), \beta(xz)) \geq \epsilon/2$.*

Proof : Recall that $p(x, y) = \beta(xy) \cap Q(x)$. It is clearly enough to show that $d(p(x, y), p(x, z)) \geq \epsilon/2$. But note that $g(x, y, z)p(x, y) = \sigma(xz)\sigma(xy)p(x, y) = \sigma(xz)p(x, y)$ and so we have $d(p(x, y), \sigma(xz)p(x, y)) \geq \epsilon$. Also $\sigma(xz)p(x, z) = p(x, z)$. It follows that $d(p(x, y), p(x, z)) \geq \epsilon/2$ as claimed. \diamond

Lemma 6.5 : *Suppose $xyz \in \Pi^2$ and length $\beta(xz)$ and length $\beta(yz)$ are bounded. Suppose also that there is some $w \in \Pi^0 \setminus \{z\}$ so that $Q(w) = \sigma(xy)Q(z)$. Then length $\beta(xy)$ is also bounded.*

Proof : If $\beta(xy)$ were very long, then it would have to remain close to the convex set $Q(z)$ over a long distance. But $\sigma(xy)$ is rotation by π about $\beta(xy)$ and so it would follow that

$Q(z)$ and $Q(w)$ would remain close over a long distance, in contradicting the assumption in (A1) that $\text{par}(Q(z), Q(w))$ is bounded. \diamond

Lemma 6.6 : *Suppose that $x_1x_2x_3 \in \Pi^2$. Suppose that for each i and all $a \in \partial Q(x_i)$, we have $d(a, g(x_i, x_{i+1}, x_{i+2})a) \geq \epsilon > 0$. Suppose that, for all i , there exists $y_i \in \Pi^0$ so that $\sigma(x_{i+1}, x_{i+2})Q(x_i) = Q(y_i)$. Then $\text{diam}(\beta(x_1, x_2) \cup \beta(x_2, x_3) \cup \beta(x_3, x_1))$ is bounded.*

Proof : First note that we must have $y_i \neq x_i$, for otherwise the three involutions $\sigma(x_1x_2)$, $\sigma(x_2x_3)$ and $\sigma(x_3x_1)$ would all be equal, and so $g(x_i, x_{i+1}, x_{i+2}) = 1$ contrary to our assumption on $\partial Q(x_i)$. Write $\beta_i = \beta(x_{i+1}, x_{i+2})$. By Lemma 6.4, $d(\beta_i, \beta_j) > \epsilon/2$ for all $i \neq j$. Recall that the β_i are the shortest paths between the convex sets $Q(x_j)$. Simple hyperbolic geometry now shows that at most of one of these β_i can be very long. But if it really were very long, we would contradict Lemma 6.5. Thus the three lengths are all bounded. We can now apply Lemma 6.1 to give the result. \diamond

We will be applying this in a situation where we have a group action. Suppose Γ acts on Π , and hence also on $R(\Pi)$, and that Γ also acts properly discontinuously on \mathbf{H}^3 . Let us suppose:

(A3) If $g \in \Gamma$ and $x \in \Pi$, $Q(gx) = gQ(x)$.

(A4) If $g \in \Gamma \setminus \{1\}$ fixes x , then for all $a \in \partial Q(x)$, $d(a, ga) \geq \epsilon$ for some fixed constant $\epsilon > 0$.

We also assume there is some Γ invariant subset, $\Pi_0^2 \subseteq \Pi^2$ such that:

(A5) If $xyz \in \Pi_0^2$, then the product of involutions, $g(x, y, z)$ described above lies in Γ (and so it fixes x)

(A6) Any 2-simplex of Π has at least two edges in simplices in Π_0^2 .

(A7) Every edge of Π lies in a 2-simplex.

Lemma 6.7 : *The assumptions (A1)–(A7) above imply (B1).*

Proof : Lemma 6.6 lets us that if $xyz \in \Pi_0^2$, then $\text{diam}(\beta(xy) \cup \beta(yz) \cup \beta(zx))$ is bounded above. Thus, by (A6), if $xyz \in \Pi^2$, at least two of $\beta(xy), \beta(yz), \beta(zx)$ have bounded length, and so by Lemma 6.5, they all do. By (A7), this accounts for all edges of Π , thereby verifying (B1). \diamond

We can therefore apply Lemma 6.3. Since the construction was canonical, the map we get is Γ -equivariant.

We now finally get to the specific application we have in mind.

To be clear about our hypotheses, we go back to the beginning. Let us suppose that Π be a simplicial complex, and let $R(\Pi)$ and $(D(x))_{x \in \Pi^0}$ be as constructed earlier. We assume that $R(\Pi)$ is locally finite. Let Γ be a group acting simplicially on Π , and so we get an induced isometric action on $R(\Pi)$. Given $x \in \Pi^0$, we write $\Gamma(x)$ for the stabiliser of x in Γ . Let Π_0^2 be a Γ -invariant subset of Π^2 . We suppose:

- (C1) Every edge of Π is contained in some 2-simplex.
- (C2) If $x \in \Pi^0$, then $\Gamma(x)$ is infinite cyclic.
- (C3) If $x, y \in \Pi^0$ are distinct, then $\Gamma(x) \cap \Gamma(y)$ is trivial.
- (C4) If $xyz \in \Pi_0^2$, we can choose generators, $g(x), g(y), g(z)$ respectively for $\Gamma(x), \Gamma(y), \Gamma(z)$ so that $g(x)g(y)g(z) = 1$.
- (C5) If $xyz \in \Pi^2$, then at least two of its edges, xy, yz, zx , are also edges of elements of Π_0^2 .

We now suppose that Γ also acts freely and properly discontinuously on hyperbolic 3-space, \mathbf{H}^3 . Writing $\Gamma(x) = \langle g(x) \rangle$ we let $l(x)$ for the infimum of translation distances of $g(x)$ on \mathbf{H}^3 . We assume:

- (C6) There is some $L \geq 0$ such that for all $x \in \Pi^0$ we have $l(x) \leq L$.

We note that this “translation bound” constant, L , is the only constant we are inputting into the proceedings. All other constants can be chosen dependent on that (though we will be free to exercise some choice).

Let $\Pi_P^0 = \{x \in \Pi^0 \mid l(x) = 0\}$. Thus if $x \in \Pi_P^0$, then $g(x)$ is parabolic with fixed point, $\pi(x) \in \partial\mathbf{H}^3$, If $x \in \Pi^0 \setminus \Pi_P^0$, then $g(x)$ is loxodromic and translates some axis, α .

Now fix some ϵ_0 less than the 3-dimensional Margulis constant, and sufficiently small in relation to L as we describe shortly. Suppose $x \in \Pi^0$ with $l(x) < \epsilon$. Let $P_0(x) \subseteq \mathbf{H}^3$ such that there exists a non-trivial $h \in \Gamma(x)$ with $d(x, hx) \leq \epsilon_0$. Thus, $P_0(x)$ is the *Margulis region* corresponding to $\Gamma(x)$. (We refer to the ϵ_0 -Margulis region if we need to specify the constant.) If $x \notin \Pi_P^0$ this is has the form $N(\alpha(x), r(x))$ for some $r(x) > 0$. We refer to $r(x)$ as the *depth* of $P_0(x)$. If $x \in \Pi_P^0$ this is a horoball centred at $\pi(x)$. We set $r(x) = \infty$. The Margulis lemma tells us that distinct Margulis regions are disjoint. Indeed we can assume the distance between them to be bounded below. (For the moment we are not excluding the possibility that there may be other points of \mathbf{H}^3 translated a very small distance by some element of $\Gamma \setminus \{1\}$ outside the regions we have described.) In the above situation ($l(x) < \epsilon_0$), we set $Q(x) = P_0(x)$. If $l(x) \geq \epsilon_0$, we set $Q(x) = \alpha(x)$.

The choice of ϵ_0 depends on the following standard fact of hyperbolic geometry. Given any $L > 0$, there is some $\epsilon(L) > 0$ such that if g, h are hyperbolic isometries generating a discrete group with $d(x, gx) \leq L$ and $d(x, hx) \leq \epsilon(L)$, the g and h generate an elementary group. In our case this will be cyclic. Thus, if we choose $\epsilon_0 \leq \epsilon(L)$, the no axis $\alpha(x)$ can enter any Margulis region $P_0(y)$. It follows that for all distinct $x, y \in \Pi^0$, $Q(x) \cap Q(y)$ is at most one point.

The following construction is most conveniently described now, even though its logical place in the argument comes a bit later. We fix some constant r_0 sufficiently large (to be specified later). Let Π_0^0 be the set of $x \in \Pi^0$ such that $l(x) < \epsilon_0$ and $r(x) \geq r_0$, and let $\Pi_1^0 = \Pi \setminus \Pi_0^0$. Note that $\Pi_P^0 \subseteq \Pi_0^0$. If $x \in \Pi_P^0$, let $A(x)$ be the horosphere about $\pi(x)$ so that $Q(x) = N(A(x), r_0)$. If $x \in \Pi_0^0 \setminus \Pi_P^0$, we let $A(x) = N(\alpha(x), r(x) - r_0)$, so that again $Q(x) = N(A(x), r_0)$. If $x \in \Pi_1^0$, then we set $A(x) = \alpha(x)$. In this case, we have $Q(x) = N(A(x), r)$ for some $r \leq r_0$.

Suppose now that $x, y \in \Pi^0$ are distinct. It follows that $\text{par}(Q(x), Q(y))$ is bounded

above. If these are both loxodromic axis, this is a standard fact following from the upper bound on translation lengths (C6) and the discreteness of Γ . Otherwise it is a standard fact about the geometry of Margulis regions: they cannot remain parallel over large distance, nor can they remain parallel to any loxodromic axis of bounded translation length. We have thus verified property (A1).

Another observation is that if $x, y \in \Pi^0$ are distinct, the $Q(x)$ and $Q(y)$ are both invariant under rotation by π about their common perpendicular $\beta(xy)$. (This has an obvious interpretation if $Q(x)$ and $Q(y)$ should happen be axes meeting at a single point.) This gives property (A2).

Now suppose that $xyz \in \Pi^2$. We can define hyperbolic isometries $g(x, y, z) = \sigma(xz)\sigma(xy)$ as before. We note that $g(x, y, z)Q(x) = Q(x)$ and that $g(x, y, z)g(y, z, x)g(z, x, y) = 1$. Suppose we also have $g(x)g(y)g(z) = 1$. We claim that $g(x) = g(x, y, z)$, $g(y) = g(y, z, x)$ and $g(z) = g(z, x, y)$. To see this, one first verifies that $g(x)\sigma(xy)$ is an involution that conjugates $g(x)$ to $g(x)^{-1}$ and $g(z)$ to $g(z)^{-1}$. Thus $g(x)\sigma(xy)$ must be equal to $\sigma(xz)$ and so $g(x) = \sigma(xz)\sigma(xy) = g(x, y, z)$. The other equalities follow by symmetry.

We are now in a set-up applicable to Lemma 6.7. Note that this makes reference only to the sets $Q(x)$, and so our construction of the sets $A(x)$ is irrelevant for the moment. In particular, Lemma 6.7 gives us an upper bound on $d(Q(x), Q(y))$ for all $xy \in \Pi^1$. This, ultimately depends only on L . We are now free to choose r_0 so that $r_0 - 1$ is greater than the constant given by Lemma 6.2. We are now in a position to apply Lemma 6.3 with $t_0 = t_1 = r_0$.

We finally note that for all $x \in \Pi_1^0$, we have $l(x) > \epsilon_1$ for some ϵ_1 depending only on ϵ_0 and r_0 , and thus ultimately only on r_0 .

Let us summarise what we have shown:

Lemma 6.8 : *Let Γ be a group acting on Π and on \mathbf{H}^3 in the manner described above, in particular satisfying (C1)–(C6). Then there are positive constants, $k, \epsilon_0, \epsilon_1$, depending only on the translation bound of property (C6), such that we can write Π^0 as an Γ -invariant disjoint union $\Pi^0 = \Pi_0^0 \sqcup \Pi_1^0$ such that there exists an equivariant k -lipschitz map, $\phi : R(\Pi) \longrightarrow \mathbf{H}^3$ satisfying:*

- (1) *If $x \in \Pi_1^0$, then the generator of $\Gamma(x)$ translates an axis $\alpha(x)$ a distance at least ϵ_1 , and $\phi(D(x)) \subseteq \alpha(x)$.*
- (2) *If $x \in \Pi_0^0$, the ϵ_0 -Margulis region, $P_0(x)$, corresponding to $\Gamma(x)$ is non-empty and $\phi(D(x)) \subseteq \partial P_0(x)$.*
- (3) *For all $x \in \Pi_0^0$, $P_0(x) \cap \phi(R(\Pi)) \subseteq \partial P_0(x)$.* ◇

This map projects to a map, $f : R(\Pi)/\Gamma \longrightarrow \mathbf{H}^3/\Gamma$. We write $\tilde{\Theta} = \mathbf{H}^3 \setminus \bigcup_{x \in \Pi_0^0} \text{int } P_0(x)$, and let $\Theta = \tilde{\Theta}/\Gamma$. Thus $f(R(\Pi)/\Gamma) \subseteq \Theta$.

We note:

Lemma 6.9 : *There is some $\epsilon_2 > 0$ depending only on L , such that if $f(R(\Pi)/\Gamma) = \Theta$, then the injectivity radius of Θ is at least ϵ_2 .*

Proof : In this case, every point of Θ is a bounded distance from some set of the form $\phi(D(x))/\Gamma(x)$ for some $x \in \Pi^0$. This is either a closed geodesic whose length is bounded below by ϵ_1 and above by L , or else the boundary of an ϵ_0 -Margulis region. This places an upper bound on the depth of any Margulis region contained in Θ and hence a lower bound on injectivity radius. \diamond

7. The model space.

In this section, we give a description of the model space for a doubly degenerate manifold, and show how the results of Section 6 can be used to construct a lipschitz map into such a hyperbolic 3-manifold.

Let $W = \bigcup \mathcal{W}$ be a complete annulus system in $\Psi = \Sigma \times \mathbf{R}$. Let $\Lambda = \Lambda(W)$ be the completion of $\Psi \setminus W$. In Section 4, we described the associated “brick decomposition”, $\mathcal{D} = \mathcal{D}(W)$ of Λ . Each element $B \in \mathcal{D}$ has the form $\Phi \times [0, 1]$, where Φ is a 3HS (“type 0”) or a 4HS or 1HT (“type 1”). Suppose B is of type 1. There is a curve $\gamma_+ \subseteq \partial_+ B$ that cuts $\partial_+ B$ into one or two 3HS components, each the lower boundary of an adjacent type 0 brick. We have a similar curve, $\gamma_- \subseteq \partial_- B$. By construction, the intersection number $\iota(\gamma_-, \gamma_+)$ is minimal (1 for a 1HT and 2 for a 4HS). Thus, if we forget about the marking (the map to Σ), then the local combinatorics of \mathcal{D} is bounded.

We want to put a path-metric on Λ using our combinatorial structure. Since we are only interested in the metric up to bilipschitz equivalence, it doesn’t much matter how we do this, but a fairly specific procedure is as follows. We fix the unique hyperbolic metric on the 3HS so that every boundary component has length 1. This will be our *standard* 3HS. For a type 0 brick, we just take a product with the unit interval. Suppose B is a type 1 brick. We put hyperbolic structures on $\partial_{\pm} B$ so that each component of $\partial_{\pm} B \setminus \gamma_{\pm}$ is a standard 3HS and so that there is no “twisting” (the topological symmetries give geometric symmetries). We now choose (once and for all) our favourite path between these structures in the space of pointwise smooth riemannian metrics for which the boundary components are always geodesic of length 1. This gives us a riemannian metric on $B = \Phi \times [0, 1]$. (It would be natural to do this in such a way that the topological symmetries of B give geometric symmetries, though this doesn’t really matter to us.)

We can now glue all the bricks back together to give us a riemannian metric on Λ . Each boundary component, $\Delta(\Omega)$ is a locally geodesic euclidean torus of the form $\Delta(\lambda)$ with respect to the longitude and standard meridian (in the notation of Section 5). Note that $\text{area}(\Delta(\lambda)) = \mathfrak{S}(\lambda)$ is the same as the “combinatorial length” of $\Delta(W)$ as defined in Section 4.

Recall that $\Lambda(\mathcal{W}, \mathcal{W}) = \Lambda(\mathcal{W}) \cup \bigcup_{\Omega \in \mathcal{W}} T(\Omega)$ is obtained by gluing in a solid torus, $T(\mathcal{W})$ to each $\Delta(\Omega)$ so that the standard meridian is trivial in $T(\Omega)$. The standard meridian was defined in such a way that $\Lambda(\mathcal{W}, \mathcal{W})$ gives us back Ψ up to homeomorphism. In particular, there is a projection map $\pi_{\Sigma} : \Lambda(\mathcal{W}, \mathcal{W}) \longrightarrow \Sigma$, well defined up to homotopy.

Now Lemma 5.7 gives $T(\Omega)$ the structure of a Margulis tube, and so we get a riemannian metric on all of $\Lambda(\mathcal{W}, \mathcal{W})$. (It is comforting to observe that by Lemma 5.8, if we chose a different metric on $\Lambda(\mathcal{W})$ with euclidean boundary and in the same bilipschitz class

we would get a bilipschitz equivalent metric on Ω , so the construction is quite “robust”. However, once we have chosen our model space, we don’t formally need to know this.)

Finally to arrive at our model space, $P = \text{int } \Sigma \times \mathbf{R}$, we glue in a Margulis cusp to each boundary component of $\Lambda(\mathcal{W}, \mathcal{W})$. Any such boundary component is a bi-infinite cylinder, $S(1) \times \mathbf{R}$, and the Margulis cusp is the quotient of a horoball by a \mathbf{Z} -action. We can regard $\Lambda(\mathcal{W}, \mathcal{W})$ as a subset of $P(\mathcal{W})$, which we shall denote by $\Psi(\mathcal{W})$ and refer to it as the *non-cuspidal* part of $P(\mathcal{W})$.

This is all we need to know about the model space to understand the statements of the main results of this section, notably Theorem 7.1. For proofs, however, we need more combinatorial constructions, in order to apply the results of Section 6. We go on to describe these next.

We will first need to cut up $\Lambda(\mathcal{W})$ into truncated simplices. This is done in a number of steps.

First, we replace our brick decomposition with a “block decomposition” — which is the same combinatorial structure as the “block decomposition” of Minsky [Mi4]. This is a fairly trivial adjustment. For each type 0 brick, $B \equiv \Phi \times [0, 1]$, we take the horizontal 3HS, $F_B = \Phi \times \{\frac{1}{2}\} \subseteq B$. The union of all these surfaces, $\bigcup_B F_B$ as B ranges over all type 0 bricks, cuts $\Lambda(\mathcal{W})$ into a collection of compact *blocks*, each of which is a type 1 brick with either two or four type 0 half bricks attached to it. Such a block is homeomorphic to $\Phi \times [0, 1]$. We write $A \cap \partial\Lambda = \partial_V A \sqcup \Upsilon_- \sqcup \Upsilon_+$, where $\partial_V A = \partial\Phi \times [0, 1]$ is the vertical boundary, and where $\Upsilon_- \subseteq \partial_- A = \Phi \times \{0\}$ and $\Upsilon_+ \subseteq \partial_+ A = \Phi \times \{1\}$. The annulus Υ_\pm retracts on to the curve γ_\pm of the type 1 brick. In particular, Υ_- and Υ_+ have minimal intersection number in Φ . We refer to the annuli Υ_- and Υ_+ as the *non-vertical annuli* of our block.

Before proceeding to the second step, we make the following observation regarding a 3HS, F , which we can take to have the standard hyperbolic structure. If $\alpha \subseteq \partial F$ is a boundary component, write $\sigma(\alpha)$ for the shortest geodesic from α to itself that separates the other two boundary components of F . If β is another boundary component, write $\sigma(\alpha, \beta)$ for the shortest geodesic from α to β . We can cut F into two right-angled hexagons in four different ways. We can cut it along $\sigma(\alpha, \beta) \cup \sigma(\beta, \gamma) \cup \sigma(\gamma, \alpha)$, or we can cut it along $\sigma(\alpha) \cup \sigma(\alpha, \beta) \cup \sigma(\alpha, \gamma)$ for any boundary curve α . We say these decompositions are of type D_0 or type D_α respectively.

The second step is to cut each block into truncated octahedra. The process is most conveniently described in reverse. Let O be a truncated octahedron — it has six square and eight hexagonal faces. We label the edges 1, 2, 3 so that the edges of each square face are alternately labelled 2 and 3, and the edges of each hexagonal face are either labelled alternately 1 and 2 or alternately 1 and 3. Thus all three labels appear at each vertex. Any two hexagons meet, if at all, in a 1-edge. There four 12-hexagons and four 13-hexagons arranged alternately.

To describe a 4HS block, we take two copies of O , and identify the corresponding pairs of 13-hexagons. This gives us a genus-3 handlebody, H . The square faces turn into a set of six disjoint annuli embedded in ∂H (each bounded by two curves labelled 2). Each component of the complement of these annuli in ∂H is a 3HS and is cut into two hexagons by three 1-arcs. This decomposition is of type D_0 . To identify H as a 4HS block, we select

four annuli which cut ∂H into two 4HS's, and deem them to be vertical. The other two annuli give us our non-vertical annuli. (They correspond to a pair of opposite squares in each copy of O .)

To describe a 1HT block, take a copy of O and partition the 13-hexagons into two pairs. Thus, the two hexagons in a pair meet a common square face. Now identify the two hexagons of a pair in such a way that their common adjacent square turns into an annulus. This gives us a genus-2 handlebody, H , and two annuli in ∂H . The other four square faces of O get strung together to form a third annulus, which we deem to be vertical. It separates ∂H into two 1HT's each containing a non-vertical annulus. This gives us a 1HT block. Each of the four 3HS components of the complement of these annuli is cut into two hexagons by three 1-arcs, as before. This time, these decompositions are of type D_α , where α is the boundary component in the vertical annulus.

We note that, in fact, the decomposition of a block as two octahedra in the manner described above is combinatorially canonical. We can thus cut each block of our decomposition up in this way.

The third step arises from the complication that the two decompositions of a horizontal 3HS into two hexagons (arising from the blocks on either side) might not match up. For example if a 4HS block meets a 1HT block along F , one decomposition will be of type D_0 and the other of type D_α . We can fix this by replacing F by a truncated simplex. Writing $\partial F = \alpha \cup \beta \cup \gamma$, we can think of one pair of opposite edges of this simplex as corresponding to $\sigma(\alpha)$ and $\sigma(\beta, \gamma)$. Another pair of opposite edges corresponding to $\sigma(\alpha, \beta)$ get identified, and a third pair, corresponding to $\sigma(\alpha, \gamma)$ also get identified. It is also possible to get two decomposition of type D_α and D_β arising from two 1HT blocks. In this case we replace F by two truncated simplices via an intermediate D_0 decomposition by applying the above construction.

This gives a polyhedral decomposition of Λ into truncated simplices and truncated octahedra. Since our discussion in Section 6 only considered truncated simplices, we should apply a fourth step. Each octahedron can be cut into four simplices by connecting two opposite vertices by an edge. This cuts a truncated octahedron into four truncated simplices. There are choices involved, but the manner in which we do it is not important.

To relate this to the discussion of Section 6, we need to pass to covers. Let $\Gamma = \pi_1(\Sigma)$, and let $\Lambda(\mathcal{W}, \mathcal{W}) = \Psi(\mathcal{W})$ for the non-cuspidal part of our model space. We have $\Lambda(\mathcal{W}) \subseteq \Psi(\mathcal{W})$, and we let R be the lift of this to the universal cover of $\Psi(\mathcal{W})$. Thus Γ acts on R with quotient $\Lambda(\mathcal{W})$. We can lift the polyhedral decomposition of $\Lambda(\mathcal{W})$ we just constructed to a polyhedral decomposition of R . This has the form $R = R(\Pi)$, where Π is the simplicial complex obtained by shrinking each boundary component of R to a point. These points become the vertices, Π^0 , of Π . The higher dimensional truncated simplices turn into simplices of Π . Thus, for each $x \in \Pi^0$, the complex $D(x)$ described in Section 6 is boundary component of $R(\Pi)$. We let Π_P^0 be the set of $x \in \Pi^0$ such that $D(x)$ is homeomorphic to \mathbf{R}^2 . In this case, $D(x)/\Gamma(x)$ is a bi-infinite cylinder, in fact a boundary component of $\Psi(\mathcal{W})$. If $x \in \Pi^0 \setminus \Pi_P^0$, then $D(x)$ is a bi-infinite cylinder, and $D(x)/\Gamma(x)$ is a torus of the form $\Delta(\Omega) = \partial T(\Omega)$ for some $\Omega \in \mathcal{W}$. In all cases, $\Gamma(x)$ is infinite cyclic (Property (C2)).

In Section 6, we defined a polyhedral metric on $R(\Pi)$, which also gives us a poly-

hedral metric on $R(\Pi)/\Gamma = \Lambda(\mathcal{W})$. Provided we carry out the subdivision of $\Lambda(\mathcal{W})$ in a geometrically sensible way, this will be bilipschitz equivalent to the model metric on $\Lambda(\mathcal{W})$ we described above. (Note that we are carrying out very explicit, locally bounded, combinatorial operations.)

We set Π_0^2 to be the set of 2-simplices that arose from hexagons in horizontal 3HS. In other words, these are 12-hexagons in the truncated octahedra constructed by the end of the second step, together with the hexagons introduced in the truncated simplices of the third step. (There will be other simplices arising from 13-hexagons in octahedra as well as those arising in the fourth step of the construction.) Note that every 2-simplex of Π has at least two edges in 2-simplices in Π_0^2 . This is property (C5). Clearly every edge lies in a 2-simplex, and so (C1) holds.

If $xyz \in \Pi_0^2$, then we can choose generators, $g(x), g(y), g(z)$, of $\Gamma(x), \Gamma(y), \Gamma(z)$ with $g(x)g(y)g(z) = 1$. This is just an observation about the boundary curves in the fundamental group of a 3HS. This is property (C4).

For property (C3), we need another assumption on \mathcal{W} , namely that no two annuli are parallel, i.e. if $\pi_\Sigma \Omega = \pi_\Sigma \Omega'$ then $\Omega = \Omega'$. In this case, distinct boundary components of $\Lambda(\mathcal{W})$ project to distinct curves in Σ (allowing peripheral curves for the cusp boundaries).

We have thus verified all the combinatorial hypotheses, (C1)–(C5), of Lemma 6.7. For the final hypothesis we need finally to introduce group actions on \mathbf{H}^3 .

Suppose that $M = \mathbf{H}^3/\Gamma$ is a complete hyperbolic 3-manifold, with a strictly type-preserving homotopy equivalence $\pi_\Sigma^M : M \rightarrow \Sigma$, to Σ . Now every curve $\alpha \in X(\Sigma)$ can be realised as a closed geodesic $\bar{\alpha}$ in M . We will abuse notation and write $\bar{\alpha} \subseteq M$, even if it is not simple. We write $l_M(\alpha)$ for the length of $\bar{\alpha}$. Given $r \geq 0$, we write $X(M, r) = \{\alpha \in X(\Sigma) \mid l_M(\alpha) \leq r\}$.

Suppose that \mathcal{W} is a complete annulus system such that no two annuli are parallel. Recall the notation $X(\mathcal{W}) = \{\pi_\Sigma \Omega \mid \Omega \in \mathcal{W}\}$. We shall make the following ‘‘a-priori bounds’’ assumption:

(APB) There is some constant $L \geq 0$, such that $X(\mathcal{W}) \subseteq X(M, L)$.

The constant, L , now gives us another constant, ϵ_0 , arising from Lemma 6.8. This is less than the Margulis constant. We write $\Psi(M)$ for the non-cuspidal part of M with respect to this constant, in other words M minus the ϵ_0 -cusps. By tameness [Bon] this is homeomorphic to $P = \Sigma \times \mathbf{R}$. Given $\alpha \in X(M, \epsilon_0)$ we write $T_0(\bar{\alpha})$ for the ϵ_0 -Margulis tube about α .

Let $P(\mathcal{W})$ be the model space constructed above, and let $\Psi(\mathcal{W})$ be its non-cuspidal part. The closures of components of $P(\mathcal{W}) \setminus \Psi(\mathcal{W})$ we shall refer to as *cusps*. Given $\Omega \in \mathcal{W}$, write $\bar{\Omega} = \overline{\pi_\Sigma(\Omega)}$ for the corresponding closed geodesic in M .

Theorem 7.1 : *Let \mathcal{W} and M be as above, in particular, satisfying (APB). Then there is a proper map $f : P(\mathcal{W}) \rightarrow M$ such that $\pi_\Sigma \circ f$ is homotopic to π_Σ^M with the following properties. Each cusp of $P(\mathcal{W})$ gets sent to a ϵ_0 -cusp of M , and $f(\Psi(\mathcal{W})) \subseteq \Psi(M)$. Moreover, we can write $\mathcal{W} = \mathcal{W}_0 \sqcup \mathcal{W}_1$ so that if $\Omega \in \mathcal{W}_1$, then $f(T(\Omega)) = \bar{\Omega}$ and $\bar{\Omega}$ has length at least $\epsilon_1 > 0$. If $\Omega \in \mathcal{W}_0$, then $\bar{\Omega}$ has length less than ϵ_0 and $T(\Omega) = f^{-1}T_0(\bar{\Omega})$.*

The map f is k -lipschitz on the complement of $\bigcup_{\Omega \in \mathcal{W}_0} \text{int } T(\Omega)$. Here, the constants, $\epsilon_0, \epsilon_1, k$ depend only on the constant L of the (APB) hypothesis.

Proof : We first pass the the covers corresponding to $\Gamma = \pi_1(\Sigma)$, and construct the polyhedral complex, $R(\Pi)$ as above. This satisfies (C1)–(C5) and (APB) gives us (C6). Thus, Lemma 6.8 gives us an equivariant map $\tilde{f} : R(\Pi) \longrightarrow \mathbf{H}^3$, which projects to a map $f : \Lambda(\mathcal{W}) \longrightarrow M$. This is lipschitz with respect to the polyhedral metric on $R(\Pi)/\Gamma = \Lambda(\mathcal{W})$, and hence also with respect to the model metric on $\Lambda(\mathcal{W})$.

Extending over a cusp of $\Psi(\mathcal{W})$ is a fairly trivial operation — essentially just coning over infinity. This does not change the lipschitz constant.

The Γ -invariant partition $\Pi^0 \setminus \Pi_P^0 = (\Pi_0^0 \setminus \Pi_P^0) \sqcup \Pi_1^0$ gives us a partition of \mathcal{W} as $\mathcal{W}_0 \sqcup \mathcal{W}_1$. If $\Omega \in \mathcal{W}_0$, then, by construction $f(\partial T(\Omega)) \subseteq \partial T_0(\bar{\Omega})$, and no other part of $\Lambda(\mathcal{W})$, or indeed the cusps, can enter $\text{int } T_0(\bar{\Omega})$. We can now extend f topologically over $T(\Omega)$. By slight adjustment in a neighbourhood of the boundary torus, we can assume that $f^{-1}(T_0(\bar{\Omega})) = T(\Omega)$. (This will not be affected by our remaining construction.)

Now suppose that $\Omega \in \mathcal{W}$. Now $\partial T(\Omega) = \Delta(\Omega)$ is a euclidean torus of the form $\Delta(\lambda)$ for some modulus λ (in the notation of Section 5). Moreover, $f|_{\partial T}$ is a lipschitz map to a circle, $\bar{\Omega} \subseteq M$, whose length is bounded above (by L), and below by the constant $\epsilon_1 > 0$ of Lemma 6.8. Thus, if we can place an upper bound on $\mathfrak{S}(\lambda)$ we can apply Lemma 5.9 to extend it to a lipschitz map of $T(\Omega)$ to $\bar{\Omega}$.

To bound $\mathfrak{S}(\lambda)$ we need again the assumption that no two annuli in Ω are parallel. From this it is a simple matter to construct a set of n curves in $\Lambda(\mathcal{W})$, all a bounded distance from $\Delta(\Omega)$ and of bounded length, which correspond to distinct elements of $X(\Sigma)$ and such that $\mathfrak{S}(\lambda)$ is bounded above by fixed (linear) function of n . For example we can take these curves to be boundary curves of type 0 brick meeting $\Delta(\Omega)$ and deleting repetitions.

Now the images of these curves in M also have bounded length, and are a bounded distance from $\bar{\Omega}$. Since $\bar{\Omega}$ has length at most L , these curves all lie in a subset of M of bounded diameter. Moreover, they are all homotopically distinct in Σ and hence in M . But a set of bounded diameter in any hyperbolic 3-manifold contains boundedly many distinct curves of bounded length (unless they are all multiples of a very short geodesics, which cannot arise here). Thus, n is bounded, and so is $\mathfrak{S}(\lambda)$ as required.

We still need to show that f is proper. First, to see that $f|_{\Lambda(\mathcal{W})}$ is is proper, we can use a variation on the above argument. Any bounded set of M can meet only finitely many toroidal boundaries, and hence only finitely many sets of the form $f(\Delta(\Omega))$ for $\Omega \in \mathcal{W}$. Since every point of $\Lambda(\mathcal{W})$ is a bounded distance from some $\Delta(\Omega)$ the properness of $f|_{\Lambda(\mathcal{W})}$ follows. The fact that f is proper on all of $P(\mathcal{W})$ now follows easily from the manner in which we have extended over tubes and cusps. \diamond

Now let $\Theta(M) = \Psi(M) \setminus \bigcup_{\Omega \in \mathcal{W}_0} \text{int } T_0(\bar{\Omega})$ and $\Theta(\Xi) = \Theta_M(P) = \Lambda(\mathcal{W}, \mathcal{W}_1) = \Psi(\mathcal{W}) \setminus \bigcup_{W \in \mathcal{W}_0} \text{int } T(\Omega)$. Thus, $f(\Theta(P)) \subseteq \Theta(M)$. Note that the definition of $\Theta(P)$ uses the partition of \mathcal{W} as $\mathcal{W}_0 \sqcup \mathcal{W}_1$ coming from Theorem 7.1, and so (at least a-priori) may depend on M .

Now the map $f : \Psi(P) \longrightarrow \Psi(M)$ is proper, and both spaces are homeomorphic to $\Psi = \Sigma \times \mathbf{R}$. It thus sends each end of $\Psi(P)$ to an end of $\Psi(M)$. We make the following

“end consistency” assumption:

(EC) Distinct ends of $\Psi(P)$ get sent to distinct ends of $\Psi(M)$.

Note that, in this case $f|\Psi(P)$ has degree 1 to $\Psi(M)$, and in particular is surjective. It also follows $f : P(\mathcal{W}) \rightarrow M$ has degree 1 and is surjective. In this case, the manifold M must be doubly degenerate.

Proposition 7.2 : *If (EC) is satisfied, then the injectivity radius of $\Theta(M)$ is bounded below by some constant ϵ_2 depending only in L .*

Proof : This is an immediate consequence of Lemma 6.9. ◇

Proposition 7.3 : *The map $f|\Theta(P) : \Theta(P) \rightarrow \Theta(M)$ is homotopic to a homeomorphism.*

Proof : By the result of Otal [O2], the set of Margulis tubes $T_0(\bar{\alpha})$ for $\Omega \in \mathcal{W}_0$ is unlinked in $\Theta(M)$. We can thus apply the results of Section 3, in particular, Proposition 3.1. ◇

Thus, applying Corollary 3.8, we get:

Lemma 7.4 : *Suppose F is a properly embedded π_1 -injective surface in $\Theta(P)$. Let U be any neighbourhood of $f(F)$ in $\Theta(M)$. Then there is a proper embedding $g : F \rightarrow U$ such that $f|F$ is homotopic to g in $\Theta(M)$ relative to ∂M .* ◇

We shall write $\mathcal{T}(P) = \{T(\Omega) \mid \Omega \in \mathcal{W}_0\}$ and write $\mathcal{T}(M) = \{f(T) \mid T \in \mathcal{T}(P)\}$ for the corresponding set of Margulis tubes in M . Thus $\Theta(P) = \Psi(P) \setminus \text{int} \bigcup \mathcal{T}(P)$ and $\Theta(M) = \Psi(M) \setminus \text{int} \bigcup \mathcal{T}(M)$.

8. Bounded geometry.

In this section, we describe some properties of “bounded geometry” manifolds. In Sections 9 and 10, these will be applied to the “thick parts”, $\Theta(P)$ and $\Theta(M)$ of our model space and hyperbolic 3-manifolds respectively. Much of the discussion is quite general.

Let Θ be a riemannian n -manifold (for us $n = 3$) with boundary $\partial\Theta$. Here, for simplicity, we shall assume that everything is smooth.

Write $B(\mathbf{R}^n) = \{\underline{x} \mid \|\underline{x}\| \leq 1\}$ for the unit ball in \mathbf{R}^n , and write $B^+(\mathbf{R}^n) = B(\mathbf{R}^n) \cap \{\underline{x} \mid x_n \geq 0\}$ for the unit half ball.

Definition : We say that Θ has *bounded geometry* if there is some $\mu > 0$ such that every $x \in \Theta$ has a neighbourhood $N \ni x$, with a smooth μ -bilipschitz homeomorphism to either $B(\mathbf{R}^n)$ or $B^+(\mathbf{R}^n)$ taking x to a point a distance at most $\frac{1}{2}$ from the origin.

(Up to modifying μ , we can equivalently replace $\frac{1}{2}$ by any constant between 0 and 1.)

One can draw a few immediate conclusions. The neighbourhood N contains and is contained in a ball of uniform positive radius about x . In particular, the injectivity radius, $\text{inj}(\Theta)$, is bounded below by some positive constant. We fix some positive constant, $\eta_0 < \frac{1}{2} \text{inj}(M)$ depending only on μ . If $x, y \in \Theta$ with $d(x, y) \leq \eta_0$, we write $[x, y]$ for the unique geodesic between them.

There are homeomorphisms, $V_{\pm} : [0, \infty) \rightarrow [0, \infty)$ such that for all $x \in \Theta$ and $r \geq 0$, we have $V_-(r) \leq \text{vol}(N(x, r)) \leq V_+(r)$.

Definition : Given $\epsilon > 0$, a subset $P \subseteq \Theta$ is ϵ -separated if $d(x, y) > \epsilon$ for all distinct $x, y \in P$.

We see that $|P| \leq \text{vol}(N(Q, \epsilon/2))/V_-(\epsilon/2)$. If $\text{diam}(P) \leq r$, then $|P| \leq V_+(r + \epsilon/2)/V_-(\epsilon/2)$.

Definition : If $P \subseteq Q \subseteq M$, we say that P is ϵ -dense in Q if $Q \subseteq N(P, \epsilon)$.

Definition : P is an ϵ -net in Q if it is $(\epsilon/2)$ -separated and ϵ -dense.

Note that any maximal $(\epsilon/2)$ -separated subset of Q is an ϵ -net in Q . We shall be taking $\epsilon < \eta_0$. The cardinality of any ϵ -net is thus bounded in terms of $\text{vol}(N(Q, \eta_0))$. We will use the following technical lemma in Section 10.

Lemma 8.1 : *Suppose Θ, Θ' are bounded geometry manifolds, and $f : \Theta \rightarrow \Theta'$ is λ -lipschitz. If $Q \subseteq \Theta$, then $\text{vol}(N(f(Q), \eta_0))$ is bounded above in terms of $\text{vol}(N(Q, \eta_0))$, λ and the bounded-geometry constants.*

(Here we are taking the same constant η_0 for Θ and Θ' .)

Proof : Let $P \subseteq Q$ be an η_0 -net. Thus $|Q|$ is bounded in terms of $\text{vol}(N(Q, \eta_0))$. Now $f(Q) \subseteq N(f(P), \lambda\eta_0)$, so $N(Q, \eta_0)$ lies inside $N(f(P), \lambda\eta_0 + \eta_0)$ whose volume is bounded above in terms of $|f(P)| \leq Q$ and λ, η_0 . \diamond

There are probably more efficient ways of dealing with this issue, but this observation will save us having to worry about taking η_0 -neighbourhoods in Θ' .

The following “nerve” construction will serve as a substitute for certain “geometric limit” arguments. (Indeed it is the basis of many precompactness results in bounded geometry, cf. [Gr2].)

Given $P \subseteq \Theta$ and $\epsilon > 0$, let $\Upsilon = \Upsilon_{\epsilon}(P)$ be the simplicial 2-complex with vertex set $V(\Upsilon) = P$ and with $A \subseteq P$ deemed to be a simplex in Υ if $\text{diam}(A) \leq 3\epsilon$. For us, P will be discrete, and so Υ will be locally finite. We write Υ^1 for its 1-skeleton. If $\epsilon < \eta_0$, then the inclusion of P into M extends to a map $\theta : \Upsilon \rightarrow \Theta$. This can be taken to send each edge of Υ linearly to a geodesic segment. We then extend over each 2-simplex by coning over a vertex. (The latter construction may entail putting some order on the vertices, and so may not be canonical.) We easily verify that $\theta(\Upsilon) \subseteq N(P, 3\epsilon)$.

Definition : We say that two paths α and β in Θ are η -close if we can parametrise them so that $d(\alpha(t), \beta(t)) \leq \eta$ for all parameter values, t .

Note that if α and β have the same endpoints, and $\eta \leq \eta_0$, then this implies that α and β are homotopic relative to their endpoints.

Suppose now that $Q \subseteq \Theta$ and that P is ϵ -dense in Q , with $3\epsilon \leq \eta_0$. If α is a path in Q with endpoints, in P , then we can find a path $\bar{\alpha}$ in Υ^1 with the same endpoints of combinatorial length at most $3((\text{length}(\alpha)/\epsilon) + 1)$ so that $\theta \circ \bar{\alpha}$ is 3ϵ -close to α . In particular, if Q is connected, then so is Υ , and the image of $\pi_1(Q)$ in $\pi_1(M)$ is contained in $\theta_*(\pi_1(\Upsilon))$.

If $\pi_1(Q)$ injects into $\pi_1(M)$, it would be nice to say that $\pi_1(Q)$ were isomorphic to $\pi_1(\Upsilon)$, but this is complicated by the fact that Q may have wriggly boundary. To help us cope with this problem, we make the following definition:

Definition : We say that Q is r -convex if given any $x, y \in Q$ with $d(x, y) \leq \eta_0$, there is some arc α in Q from x to y so that $\alpha \cup [x, y]$ bounds a disc of diameter at most r in Θ .

Note that an immediate consequence is that if $P \subseteq Q$ is ϵ -dense for some $\epsilon \leq \eta_0$, then if β is any path in Υ^1 then $\theta \circ \beta$ can be homotoped relative to its endpoints into Q . In particular, given our previous observation, we see that the image of $\pi_1(Q)$ in $\pi_1(M)$ must equal $\theta_*(\pi_1(\Upsilon))$. The problem remains that $\pi_1(\Upsilon)$ may have lots non-trivial loops “near the boundary”.

Suppose then that Q is r -convex, and let $Q' = N(Q, r)$. Let $P \subseteq Q$ be an ϵ -net in P and extend to an ϵ -net $P' \subseteq Q'$. We thus have an inclusion of Υ in Υ' . Write $\Gamma(\Upsilon, \Upsilon')$ for the image of $\pi_1(\Upsilon)$ in $\pi_1(\Upsilon')$. Note that θ induces a natural map of $\Gamma(\Upsilon, \Upsilon')$ into $\pi_1(M)$.

Lemma 8.2 : *Suppose that Q is r -convex and suppose that $\pi_1(Q)$ injects into $\pi_1(M)$. Let Υ and Υ' be as constructed above. Then the natural map of $\Gamma(\Upsilon, \Upsilon')$ into $\pi_1(M)$ is injective, and its image equals the image of $\pi_1(Q)$. (In particular, $\Gamma(\Upsilon, \Upsilon')$ is isomorphic to $\pi_1(Q)$.)*

Proof : We have already observed that the image of $\pi_1(\Upsilon)$ and hence of $\Gamma(\Upsilon, \Upsilon')$ in $\pi_1(M)$ equals the image of $\pi_1(Q)$. We thus need to show that the map of $\Gamma(\Upsilon, \Upsilon')$ into $\pi_1(M)$ is injective. If β is any closed curve in Υ^1 . then $\theta \circ \beta$ consists of a sequence of geodesics arcs of length at most $3\epsilon \leq \eta_0$ connecting points of $P \subseteq Q$. By r -convexity, $\theta \circ \beta$ can be homotoped to Q inside $Q' = N(Q, r)$. If $\theta \circ \beta$ is trivial in $\pi_1(M)$, then it the homotoped curve is also trivial in $\pi_1(Q)$. It thus follows that $\theta \circ \beta$ bounds a disc in Q' . We can now pull back this disc to Υ' showing that β is trivial in $\pi_1(\Upsilon')$, and hence in $\Gamma(\Upsilon, \Upsilon')$ as required. \diamond

As an application, we have the following lemma to be used in Section 9. For the purposes of this lemma, we can define a “band” in a 3-manifold, Θ , to be a closed subset, $B \subseteq \Theta$, isomorphic to $\Sigma \times [0, 1]$, where Σ is compact surface such that $B \cap \partial\Theta = \partial_V B$, where $\partial_V B = \partial\Sigma \times [0, 1]$ is the “vertical boundary”. Note that the relative boundary of B in Θ is the “horizontal boundary” $\Sigma \times \{0, 1\}$.

Lemma 8.3 : *Let Θ be a bounded geometry 3-manifold, with $\eta_0 \leq \frac{1}{2} \text{inj}(\Theta)$ as before. Suppose that $B \subseteq \Theta$ is a band with $\pi_1(B)$ injecting into $\pi_1(\Theta)$ and that B is r -convex. Suppose there is a constant $s \geq 0$ such that each component of $\partial_V B$ is homotopic to a curve of length at most s . Suppose that α, β are curves in B of length at most t for some other constant $t \geq 0$, and such that α and β are homotopic to curves $[\alpha], [\beta]$ in $X(\Sigma)$. Then the intersection number, $\iota([\alpha], [\beta])$, of α and β in Σ is bounded above in terms of $r, s, t, \text{diam}(B)$ the constant of bounded geometry (including η_0), and the complexity, $\kappa(\Sigma)$, of the surface Σ .*

Note that the intersection number is independent of the choice of homeomorphism of B with $\Sigma \times [0, 1]$. The bound also places a bound on the distance, $d_X([\alpha], [\beta])$, between $[\alpha]$ and $[\beta]$ in the curve graph (which is what we are really interested in).

The proof relies on the observation that the intersection number of two curves is a function of the pair of conjugacy classes in $\pi_1(\Sigma)$ representing their free homotopy class. (We need not explicitly describe what this function is, though of course this is in principle computable.) In the case where $\pi_1(\Sigma)$ has boundary, we need also to take into account the peripheral structure — the set of conjugacy classes of boundary curves.

In practice the “short” peripheral curves in B will just be core curves of the corresponding annuli.

Proof : Fix some $\epsilon \leq \eta_0/6$. Let $P \subseteq B$ be an ϵ -net and extend to an ϵ -net P' of $Q' = N(Q, r)$, and construct $\Upsilon = \Upsilon_\epsilon(P)$ and $\Upsilon' = \Upsilon_\epsilon(P')$ as above. Note that the diameter of Q' is bounded, and so $|V(\Upsilon')| = |P'|$ is bounded. By Lemma 8.2, there is a natural isomorphism of $\pi_1(Q) \cong \pi_1(\Sigma)$ with $\Gamma(\Upsilon, \Upsilon')$. Note that α and β correspond to curves, $\bar{\alpha}$ and $\bar{\beta}$ of bounded length in the 1-skeleton of Υ . If $\Omega_1, \dots, \Omega_n$ is the (possibly empty) set of vertical boundary components, then each Ω_i is homotopic to a curve, γ_i in B of bounded length, and thus corresponds to some bounded length curve, $\bar{\gamma}_i$, in the 1-skeleton of Υ . We see that there are boundedly many combinatorial possibilities for $\Upsilon, \Upsilon', \bar{\alpha}, \bar{\beta}, (\gamma_i)_i$. Among all such possibilities for which $\Gamma(\Upsilon, \Upsilon')$ is isomorphic to $\pi_1(\Sigma)$ with the γ_i peripheral, there is a maximal intersection number of $[\alpha]$ and $[\beta]$ which will serve as our bound. \diamond

We remark that this argument does not give us a computable bound, since it involves sifting out those pairs, Υ, Υ' for which $\Gamma(\Upsilon, \Upsilon')$ is a surface group, and this is not algorithmically testable. In principle, the above argument could be translated into a “geometric limit” argument.

Here is another application of this construction, to be used in Section 10 (see Proposition 10.11).

Lemma 8.4 : *Suppose that $f : \Theta \rightarrow \Theta'$ is a surjective lipschitz homotopy equivalence between two bounded geometry manifolds Θ and Θ' . Suppose there is some positive $\epsilon < \eta_0$ such that if $x, y \in M$ with $d'(f(x), f(y)) \leq \epsilon$. then there is a path α from x to y in Θ with $\text{diam}(\alpha)$ bounded such that $f(\alpha) \cup [f(x), f(y)]$ bounds a disc of bounded diameter in Θ' . Then f is universally sesquilipschitz.*

In other words, the lift, $\tilde{f} : \tilde{\Theta} \rightarrow \tilde{\Theta}'$ is a quasi-isometry. The constants of quasi-isometry depend only on the constants of the hypotheses (though, again, we do not show this dependence to be computable). Of course, the fact that f itself is sesquilipschitz (i.e. a quasi-isometry) is an immediate consequence of the hypotheses.

Note that, since f is a homotopy equivalence, the conclusion means that α must lie in a particular homotopy class relative to its endpoints x and y . We refer to this as the “right homotopy class”. The hypotheses tell us that, in particular, that there is path of bounded diameter in the right class, and we need to find one of bounded length.

Proof : For the purposes of the proof (rescaling the metric on Θ or Θ' if necessary, and modifying the bounded geometry constants) we can assume, for notational convenience, that f is 1-lipschitz. We can also take the same constant η_0 for both Θ and Θ' . Fix some $\epsilon \leq \eta_0/6$.

Let R be the bound on the length of α in Θ and let R' be the bound on the length of the disc in Θ' . We assume $R \leq R'$. Let $Q = N(x, r) \subseteq \Theta$ and let $Q' = N(f(x), R + 6\epsilon) \subseteq \Theta'$. Let P_0 and P'_0 be ϵ -nets in Q and Q' respectively. Let $P = P_0 \cup \{x, y\}$, and let $P' = P'_0 \cup f(P) \subseteq Q'$. Thus, P and P' are ϵ -dense in Q and Q' respectively. Note that since the diameters of Q and Q' are bounded, $|P_0|$ and $|P'_0|$ and hence $|P|$ and $|P'|$ are bounded.

Let $\Upsilon = \Upsilon_\epsilon(P)$ and $\Upsilon' = \Upsilon_\epsilon(P')$, and write $\theta : \Upsilon \rightarrow \Theta$ and $\theta' : \Upsilon' \rightarrow \Theta'$ for the corresponding maps as constructed earlier. Note that the map $f : V(\Upsilon) \rightarrow V(\Upsilon')$ extends to a simplicial map $g : \Upsilon \rightarrow \Upsilon'$. We see that $f \circ \theta$ and $\theta' \circ g$ are 3ϵ -close on the 1-skeleton of Υ .

Let α be the path connecting x to y as given by the hypotheses. Now $\alpha \subseteq Q$, so there is a path $\bar{\alpha}$ from x to y in Υ such that $\theta \circ \bar{\alpha}$ is 3ϵ -close to α . Thus $f \circ \alpha$ is 3ϵ -close to $f \circ \theta \circ \bar{\alpha}$ and hence 6ϵ -close to $\theta' \circ g \circ \bar{\alpha}$. Now $f(\alpha) \cup [f(x), f(y)]$ bounds a disc $D \subseteq N(f(x), R')$, and so $\theta' \circ g \circ \bar{\alpha} \cup [f(x), f(y)]$ bound a disc in $Q' = N(f(x), R' + 6\epsilon)$. This pulls back to a disc in Υ' bounding $(g \circ \bar{\alpha}) \cup e$, where e is the edge connecting $f(x)$ to $f(y)$ in Υ' .

In summary, we have two simplicial 2-complexes, Υ, Υ' with $|V(\Upsilon)|$ and $|V(\Upsilon')|$ bounded, a simplicial map $g : \Upsilon \rightarrow \Upsilon'$ and vertices $x, y \in V(\Upsilon)$, with the property that $f(x)$ and $f(y)$ are connected by some edge e in Υ' , and such x and y are connected by some path whose image under f together with e bounds a disc in Υ' . Now there are boundedly many combinatorial possibilities for $\Upsilon, \Upsilon', g, x, y$. For each such $\Upsilon, \Upsilon', g, x, y$, we choose some path, say β so that $(g \circ \beta) \cup e$ bounds a disc. Since there are only finitely many cases, there is some upper bound for the length of any such β depending only on the bounds on $|V(\Upsilon)|$ and $|V(\Upsilon')|$. Let $\gamma = \theta \circ \beta$. This has bounded length in Θ , and since $(g \circ \beta) \cup e$ bounds a disc in Υ , we see that $(\theta' \circ g \circ \beta) \cup [f(x), f(y)]$ bounds a disc in Θ' (of bounded diameter). Now $\theta' \circ g \circ \beta$ is 3ϵ -close to $f \circ \theta \circ \beta = f \circ \gamma$, so $(f \circ \gamma) \cup [f(x), f(y)]$ bounds a disc in Θ' .

In other words, we can find a path from x to y of bounded length in the right homotopy class. The rest of the argument is now fairly standard.

Reinterpreting in terms of universal covers, we have a lift $\tilde{f} : \tilde{\Theta} \rightarrow \tilde{\Theta}'$. If $a, b \in \tilde{\Theta}$ with $d(\tilde{f}(a), \tilde{f}(b)) \leq \epsilon$, then $d(a, b)$ is bounded by some constant, say k . Given any $a, b \in \Theta$, we can connect a to b by a geodesic σ in Θ' from $\tilde{f}(a)$ to $\tilde{f}(b)$. Choose points

$\tilde{f}(a) = c_0, c_1, \dots, c_n = \tilde{f}(b)$ along σ so that $d'(c_i, c_{i+1}) \leq \epsilon$ for all i , and so that $n \leq d'(f(a), \tilde{f}(b))/\epsilon + 1$. Since f is assumed surjective, \tilde{f} is also surjective, so we can find points $a = a_0, a_1, \dots, a_n = b$ in Θ with $\tilde{f}(a_i) = c_i$. Thus, $d(a_i, a_{i+1}) \leq k$ for all i , so that $d(a, b)$ is bounded above by a linear function of $d'(f(a), \tilde{f}(b))$.

Since f is lipschitz, so is \tilde{f} , and so \tilde{f} is a quasi-isometry as required. \diamond

Before finishing this section, we make the observation that if a group Γ acts freely properly discontinuously on a riemannian manifold Θ , then Θ/Γ has bounded geometry if and only if Θ has bounded geometry and the orbits of Γ are uniformly separated sets.

9. Lower bounds.

We have constructed, in Section 7, a lipschitz map between a model space and our 3-manifold. In this section, we begin the project of showing that there is also a linear lower bound on distortion of distances. We will do this in a series of steps. First we shall restrict our attention to the modified metrics, ρ and ρ' as defined at the end of Section 4. We will need a result about 1-dimensional quasi-isometries. In what follows we will write $[t, u] = [u, t]$ for the interval between $t, u \in \mathbf{R}$, regardless of the order of t and u .

Suppose that $I \subseteq \mathbf{R}$ is an interval and that $\sigma : I \rightarrow \mathbf{R}$ is a continuous map. We shall say that σ is *quasi-isometric* if it is a quasi-isometry to its range, $\sigma(I)$. Writing $I = [\partial_- I, \partial_+ I]$, it necessarily follows that $\sigma(I)$ lies in a bounded neighbourhood of $[\sigma(\partial_- I), \sigma(\partial_+ I)]$. We shall allow the possibility that $\partial_- I = -\infty$ and $\partial_+ I = \infty$. In this case, σ is self quasi-isometry of \mathbf{R} .

We list the following properties of a continuous map $\sigma : I \rightarrow \mathbf{R}$ which together will imply that it is quasi-isometric. Let Q be a closed subset of I . We suppose:

(Q0) $(\forall k)(\exists K_0(k))$ if $t, u \in I$ with $|t - u| \leq k$ then $|\sigma(t) - \sigma(u)| \leq K_0(k)$.

(Q1) $(\forall k)(\exists K_1(k))$ if $t, u \in Q$ and $|\sigma(t) - \sigma(u)| \leq k$ then $|t - u| \leq K_1(k)$

(Q2) $(\forall k)(\exists K_2(k))$ if $t, u \in I$ and $\text{diam}(\sigma[t, u]) \leq k$, then $|t - u| \leq K_2(k)$.

(Q3) $(\exists k_3)(\forall k)(\exists K_3(k))$ if $t, u \in I$ and $N([\sigma(t), \sigma(u)], k_3) \cap \sigma(Q) = \emptyset$ and $|\sigma(t) - \sigma(u)| \leq k$ then $|t - u| \leq K_3(k)$.

(Q4) $(\exists k_4)$ if $t, u \in I$, $[t, u] \cap Q = \emptyset$, then $[\sigma(t), \sigma(u)] \cap \sigma(Q) \subseteq N(\{\sigma(t), \sigma(u)\}, k_4)$.

We can paraphrase the above conditions informally as follows. (Q0) gives an upper bound on distortion, and (Q1) gives a lower bound restricted to Q . (Q2) tells us that no long interval can get sent onto a short interval. (Q3) gives a lower bound on distortion, so long as we stay away from $\sigma(Q)$. Finally (Q4) tells us that intervals in the complement of Q can not fold too deeply over Q . Note that we always apply (Q4) to a subinterval $[t', u'] \subseteq [t, u]$ with $\sigma([t', u']) = [\sigma(t'), \sigma(u')] = \sigma([t, u])$.

Lemma 9.1 : *Let $\sigma : I \longrightarrow \mathbf{R}$ be a continuous map, with $Q \subseteq I$ satisfying (Q0)–(Q4) above. Then σ is a quasi-isometry, and the constants of quasi-isometry depend only on the constants of the hypotheses.*

Proof : Let $Q' = \sigma(Q)$ and let $P' = \{t \in \mathbf{R} \mid N(t, 2k_3) \cap Q = \emptyset\}$, and let $P = \sigma^{-1}P'$. Then $P \subseteq I$ is closed. We claim that each component of $I \setminus (P \cup Q)$ has bounded length.

Note that each component of $\mathbf{R} \cap (P' \cup Q')$ has length at most k_3 . Suppose that $J \subseteq I$ is an interval with $J \cap (P \cup Q) = \emptyset$. Now $\sigma(J) \cap P' = \emptyset$, and property (Q4) bounds the extent to which $\sigma(J)$ can cross Q' . In fact, we get $\text{diam}(\sigma(J)) \leq 2k_4 \cup 2k_3$. It now follows by (Q2) that $\text{diam}(J)$ is bounded (by $K_2(2k_4 + 2k_3)$). This proves the claim.

Now fix some $k_0 < k_3$, and suppose that $t < u \in I$ with $|\sigma(t) - \sigma(u)| \leq k_0$. We claim that $u - t$ is bounded. If $t \in P$, then $\sigma(u) \in N(P', k_0) \subseteq N(P', k_3)$ so $N([\sigma(t), \sigma(u)], k_3) \cap Q' = \emptyset$, and $u - t$ is bounded by (Q3). Similarly if $u \in P$. We can thus assume that $t, u \notin P$.

Now if $[t, u] \cap (P \cup Q) = \emptyset$, then $u - t$ bounded by our earlier claim. If not, let t_0 and u_0 be, respectively, the minimal and maximal points of $[t, u] \cap (P \cup Q)$. Again, $u - u_0$ and $t_0 - t$ are bounded. It in turn follows that $|\sigma(t_0) - \sigma(u_0)|$ is bounded (by (Q0)). If $t_0, u_0 \in Q$, then $u_0 - t_0$ is bounded by (Q1), and so $u - t$ is bounded. If $t_0 \in P$, then let t_1 be the maximal point of $[t, u] \cap \sigma^{-1}\sigma(t_0)$. Since $\sigma(t_0) = \sigma(t_1) \in P'$, $t_1 - t_0$ is bounded by (Q3), and so it is enough to consider the interval $[t_1, u_0]$. Similarly, if $u_0 \in P$ let u_1 be the minimal point of $[t, u] \cap \sigma^{-1}\sigma(u_0)$, we see that $u_0 - u_1$ is bounded. But if $\sigma|_{[t, u]}$ enters any component of P' it must eventually leave by the same point. Thus the above observations allow us to reduce to the case where $(t, u) \cap P = \emptyset$, and so again we get $u - t$ bounded, as before.

This proves the second claim. The fact that σ is quasi-isometric is now standard. \diamond

We shall be applying this to spaces quasi-isometric to intervals, and we will need some general observations concerning such spaces.

Suppose that Ψ is a locally compact locally connected space with two ends, deemed “positive” and “negative”. (In practice $\Psi \cong \Sigma \times \mathbf{R}$.) By an *end-separating* set Q , we mean a compact connected subset which separates the two ends of Ψ . We write $C_+(Q)$, $C_-(Q)$ for the components of $\Psi \setminus Q$ containing the positive and negative ends of Ψ respectively. We write $C_0(Q) = \Psi \setminus (C_+(Q) \cup C_-(Q))$. If Q' is another end separating set, we write $Q < Q'$ to mean that $Q \subseteq C_-(Q')$. One can verify that this is equivalent to stating that $Q' \subseteq C_+(Q)$, and that $<$ is a total order on any locally finite pairwise disjoint collection of end-separating sets. If $Q < Q'$ we write $[Q, Q'] = [Q', Q]$ for the closure of $C_+(Q) \cap C_-(Q')$. This is the compact region of Ψ between Q and Q' . We also note that if $P \subseteq Q$ is end-separating, then $C_-(Q) \subseteq C_-(P)$ and $C_+(Q) \subseteq C_+(P)$.

Suppose now that Ψ has a complete path-metric, ρ , and suppose that every point $x \in \Psi$ lies in some end-separating set with diameter bounded by some constant, l , say. (Here we measure diameter with respect to ρ , not with respect to the intrinsic path metric.) We choose such a set $Q(x)$. Note that if $y \in C_0(Q(x))$, then $Q(y) \cap Q(x) \neq \emptyset$, and so $\text{diam } C_0(Q(x)) \leq 3l$.

Let $\pi : \mathbf{R} \setminus \Psi$ be a bi-infinite end-respecting geodesic, i.e. $\pi|_{(-\infty, 0]}$ goes out the negative end, and $\pi|_{[0, \infty)}$ goes out the positive end of Ψ . Clearly any end-separating set

must meet $\pi(\mathbf{R})$ and so $\Psi = N(\pi(\mathbf{R}), l)$. In particular, π is a quasi-isometry. We can see in fact, that any two geodesics in Ψ with the same endpoints must remain a bounded distance (in fact l) apart. The same remains true if the respective endpoints are bounded distance apart. If $x, y \in \Psi$ with $Q(x) \cap Q(y) = \emptyset$, then $[Q(x), Q(y)]$ is a bounded Hausdorff distance from any geodesic from x to y . If $t, u \in \mathbf{R}$ with $u > t + 2l$, then $Q(\pi(t)) \cap Q(\pi(u)) = \emptyset$, and $[Q(\pi(t)), Q(\pi(u))]$ is a bounded Hausdorff distance from $\pi([t, u])$.

We shall be applying this in the case where $\Psi = \Sigma \times \mathbf{R}$, and every point of Π lies in the image of a homotopy equivalence, ψ , of Σ into Ψ of bounded diameter in Ψ . (Here all homotopy equivalences are assumed to be relative to the boundaries $\partial\Sigma$ and $\partial\Psi$. By [FHS] we can find an embedded surface, Z , in an arbitrarily small neighbourhood of $\psi(\Sigma)$. This is a ‘‘fibre’’ of Ψ in the sense discussed in Section 3. If we have two such ψ and ψ' , with the $\psi(\Sigma) \cap \psi'(\Sigma) = \emptyset$ and Z, Z' are nearby fibres, then $[Z, Z']$ is a band in Ψ (with base surface Σ) and $[Z, Z']$ is a bounded Hausdorff distance from $[\psi(\Sigma), \psi'(\Sigma)]$.

Suppose that (Ψ, ρ) and (Ψ', ρ') are two such product spaces, and $f : \Psi \rightarrow \Psi'$ is a proper lipschitz end-preserving map. Let $\pi : \mathbf{R} \rightarrow \Psi$ and $\pi' : \mathbf{R} \rightarrow \Psi'$. We can find a map $\sigma : \mathbf{R} \rightarrow \mathbf{R}$ such that for all $t \in \mathbf{R}$, $\rho'(\pi'(\sigma(t)), f(\pi(t)))$ is uniformly bounded.

Let us now focus on the case of interest where $f : \Psi(P) \rightarrow \Psi(M)$ is the map between the non-cuspidal parts of our model space P and hyperbolic 3-manifold M . These have riemannian metrics, d and d' respectively. Let $\mathcal{T}(P)$ and $\mathcal{T}(M)$ be the sets of Margulis tubes in $\Psi(P)$ and $\Psi(M)$ respectively, and let $\Theta(P) = \Psi(P) \setminus \text{int} \bigcup \mathcal{T}(P)$ and $\Theta(M) = \Psi(M) \setminus \text{int} \bigcup \mathcal{T}(M)$ be the respective thick parts. Now f maps $\Theta(P)$ onto $\Theta(M)$ and is lipschitz with respect to the metrics d and d' . We can define reduced riemannian (pseudo)metrics ρ and ρ' on $\Psi(P)$ and $\Psi(M)$ respectively, agreeing with d or d' on the thick parts, and equal to zero on each Margulis tube. The map $f : (\Psi(P), \rho) \rightarrow (\Psi(M), \rho')$ is then lipschitz. In what follows, all distances and diameters etc, refer to the metrics ρ or ρ' unless otherwise specified. As observed in Section 4, we can foliate $\Psi(P)$ with fibres of bounded diameter. We write $S(x)$ for the fibre containing x , so that $S(x)$ varies continuously in the Hausdorff topology. We are thus in the situation described above with $\Psi = \Psi(P)$ and $\Psi' = \Psi(M)$. We have geodesics π and π' and a map $\sigma : \mathbf{R} \rightarrow \mathbf{R}$.

Lemma 9.2 : *The map σ arising as above (from the map $f : \Psi(P) \rightarrow \Psi(M)$) is a self-quasi-isometry of \mathbf{R} .*

To prove this, we shall apply Lemma 9.1. We set $Q = \{t \in \mathbf{R} \mid S(\pi(t)) \cap \bigcup \mathcal{T}(P) \neq \emptyset\}$. Thus Q is a closed subset of \mathbf{R} . Note that property (Q0) of Lemma 9.1, is an immediate consequence of the construction and the fact that f is lipschitz. We now set about verifying properties (Q1)–(Q4).

(Q1): This argument is based on a similar construction due to Bromberg, which is discussed in [Brocb]. (We are using images of fibres under f , in place of interpolations of pleated surfaces used by Bromberg.) Indeed the proof of (Q1) given here could be shortened by quoting Bromberg’s result (namely that two Margulis tubes a bounded apart in M are a bounded distance apart in the curve complex). However, we will need some variants of this constructions later.

Suppose $t, u \in Q$ with $|\sigma(t) - \sigma(u)| \leq k$. By the definition of Q , $S(\pi(t))$ and $S(\pi(u))$ meet Margulis tubes $T_0, T_1 \in \mathcal{T}(P)$ respectively. Let $T'_i = f(T_i) \in \mathcal{T}(M)$. We want to show that $\rho(T_0, T_1)$ is bounded, since it then follows that $\rho(\pi(t), \pi(u))$ is bounded, and so, since π is geodesic, $|t - u|$ is bounded, as required. Clearly, we can assume that $T_0 \neq T_1$, or there is nothing to prove.

Now $f(S(\pi(t)))$ meets T'_0 and has bounded diameter. Also $f(\pi(t)) \in f(S(\pi(t)))$ is a bounded distance from $\pi'(\sigma(t))$ (by definition of σ) and so $\rho'(\pi'(\sigma(t)), T'_0)$ is bounded. Similarly, $\rho'(\pi'(\sigma(u)), T'_1)$ is bounded. Also $\rho'(\sigma(t), \sigma(u)) = |\sigma(t) - \sigma(u)| \leq k$, and so $\rho(T'_0, T'_1)$ is bounded (in terms of k). In other words we can connect T'_0 to T'_1 by a path β of bounded ρ' -length. Indeed we can assume that $\beta \cap \Theta(M)$ consists of a bounded number of paths β_1, \dots, β_n of bounded d -length connecting different Margulis tubes.

Since $f : \Psi(P) \rightarrow \Psi(M)$ is a homotopy equivalence and $f^{-1}(T'_0) = T_0$ and $f^{-1}(T'_1) = T_1$, there is a path $\alpha \subseteq f^{-1}\beta$ connecting T_0 to T_1 in $\Psi(P)$. For each $x \in \alpha$, there is a loop γ_x in $\Psi(P)$ based at x of bounded d -length, with the properties (W2)–(W4) described in Section 4. In particular, γ_x is freely homotopic to a curve, $[\gamma_x] \in X(\Sigma)$ (via the natural homotopy equivalence). We can also take γ_x to lie either in $\Theta(P)$ or else inside a Margulis tube and hence freely homotopic to the core of that tube. If x, y are sufficiently close then $d_X([\gamma_x], [\gamma_y])$ is bounded: less than $2r \geq 0$ say. (“Sufficiently” can be taken to imply some uniform positive constant, but uniformity is not needed here.) Let $Y = \{[\gamma_x] \mid x \in \alpha\} \subseteq X(\Sigma)$. Since α is connected, by the above observation, the r -neighbourhood of Y in $\mathcal{G}(\Sigma)$ is connected. We claim that $|Y|$ is bounded. This will place a bound on the diameter of Y in $X(\Sigma)$.

To see that $|Y|$ is bounded, note that if $x \in \alpha$, then $f(\gamma_x)$ is a loop in $\Psi(M)$ based at $f(x) \in \beta$. Either $f(x)$ lies in one of the segments β_i or else it is freely homotopic to the core of one of (the bounded number of) Margulis tubes passed through by β . But now for each β_i , there can be only boundedly many possibilities for the free homotopy classes of β_i . This follows from the standard fact of hyperbolic 3-manifolds that there are a bounded number of (based) homotopy classes of curves of bounded length based at any point of the 3-manifold (at least if we rule out multiples of a very short curve, which cannot happen here). Alternatively, this is a general statement about bounded geometry manifolds, given that we are working in the thick part. We see that there is bound on the number of possible $[\gamma_x]$. This bounds $|Y|$ as required.

But now the core curves of T_0 and T_1 lie in Y and so we have bounded the distance between these curves in $\mathcal{G}(\Sigma)$. By the tautness assumption (see (W4) Section 4), it follows that $\rho(T_0, T_1)$ is bounded as required, and so (Q1) follows.

(Q2): This is essentially a variation on the same argument. Suppose $u, t \in \mathbf{R}$ with $\text{diam}(\sigma([s, t]))$ bounded. Let $R \subseteq \Psi$ be the region between $S(\pi(t))$ and $S(\pi(u))$. Again, since $f : \Psi(P) \rightarrow \Psi(M)$ is a homotopy equivalence, there is a path $\alpha \subseteq R \cap f^{-1}(\pi'(\mathbf{R}))$ connecting $S(\pi(t))$ to $S(\pi(u))$. Now α will be a bounded Hausdorff distance from any geodesic connecting also connecting $S(\pi(t))$ to $S(\pi(u))$ in $\Psi(P)$, in particular, the geodesic $\pi([t, u])$. Thus, $f(\alpha)$ is a bounded Hausdorff distance from $f(\pi([t, u]))$, which, by the definition of σ , is in turn a bounded Hausdorff distance from $\pi'(\sigma([t, u]))$ in $\Psi(M)$. By assumption, $\text{diam}(\sigma([t, u]))$ is bounded. It therefore follows that $\text{diam}(f(\alpha))$ is bounded. We see that

$f(\alpha) \cap \Theta(M)$ consists of a bounded number of segments β_1, \dots, β_n , each of bounded d' -length, and we proceed exactly as in (Q1) to show that $|t - u|$ is bounded.

(Q3): Let us suppose that $t, u \in \mathbf{R}$ and that $N([\sigma(t), \sigma(u)], k_3) \cap \sigma(Q) = \emptyset$ for some constant k_3 to be determined shortly, and suppose that $|\sigma(t) - \sigma(u)| \leq k$. Let $x = \sigma(t)$ and $y = \sigma(u)$. We want to bound $|t - u|$, which is the same as bounding $\rho(x, y)$. Since $t, u \notin Q$, we have $S(x), S(y) \subseteq \Theta(P)$. In particular, $x, y \in \Theta(P)$ and so the paths γ_x and γ_y lie in $\Theta(P)$. Since $|\sigma(t) - \sigma(u)|$ is bounded, $\rho'(f(x), f(y))$ is bounded. Also $f(\gamma_x)$ and $f(\gamma_y)$ have bounded d' -length. We want to apply Lemma 8.3. This means finding a band $B \subseteq \Theta(M) \subseteq \Psi(M)$, which base surface σ , of bounded diameter, containing γ_x and γ_y , and which is r -convex for some uniform $r > 0$. For the latter, it is enough to find another band $A \subseteq \Theta(M)$ containing a uniform neighbourhood of B .

Let $h_1 > h_0 > 0$ be sufficiently large constants, to be determined shortly. Let $t' = \sigma(t)$ and $u' = \sigma(u)$. We can suppose that $t' \leq u'$. Let $t_0 = t' - h_0$, $t_1 = h_1$, $u_0 = u' + h_0$, $u_1 = u' + h_1$. Thus $t_1 < t_0 < t' < u' < u_0 < u_1$. Let $i = 0, 1$, let Z_i^-, Z_i^+ be a fibres in a small neighbourhood of $f(S(\pi(t_i)))$ and $f(S(\pi(u_i)))$ respectively. Now if h_0 and $h_1 - h_0$ are sufficiently large, then we have $Z_1^- < Z_0^- < Z_0^+ < Z_1^+$. Let $B = [Z_0^-, Z_0^+]$ and $A = [Z_1^-, Z_1^+]$, so that $B \subseteq A$. Note that B and A are a bounded Hausdorff distance from $\pi'([t_0, u_0])$ and $\pi'([t_1, u_1])$ respectively, and that $f(x)$ and $f(y)$ are a bounded distance from $\pi'(t')$ and $\pi'(u')$ respectively. Thus, again by choosing h_0 and $h_1 - h_0$ sufficiently large, we can assume that $f(\gamma_x), f(\gamma_y) \subseteq B$ and that $N(B, r) \subseteq A$, for some uniform $r > 0$ (as usual, with respect to the metric ρ').

We claim that, provided $k_3 - h_1$ is sufficiently large, then $A \subseteq \Theta(M)$. Suppose that $T \cap A \neq \emptyset$ for some $T \in \mathcal{T}(M)$. Let $s \in I$ be such that the fibre $S(\pi(s))$ meets the corresponding Margulis tube in P . By definition, $s \in Q$. Thus $f(S(\pi(s))) \cap T \neq \emptyset$. Now $\pi'(\sigma(s))$ is a bounded distance from $f(\pi(s))$ which is a bounded distance from T and hence from $\pi'([t_1, u_1])$. Thus $\sigma(s)$ is a bounded distance from $[t_1, u_1]$. But $\sigma(s) \in \sigma(Q)$, so we get a contradiction by taking $k_3 - h_1$ large enough. This shows that $A \subseteq \Theta(M)$ as claimed.

But now γ_x, γ_y, B satisfy the hypotheses of Lemma 8.3, which means that $d_{\mathcal{G}(\Sigma)}([\gamma_x], [\gamma_y])$ is bounded. By (W4) of Section 4, it follows that $\rho(x, y)$ is bounded, thereby giving a bound on $|t - u|$ as required. ■

(Q4): Suppose $t, u \in \mathbf{R}$ with $[t, u] \cap Q = \emptyset$. Let $x = \pi(t)$ and $y = \pi(u)$. Let $R = [S(x), S(y)]$ be the band bounded by $S(x)$ and $S(y)$. Since $[t, u] = \emptyset$, we have $R \subseteq \Theta(P)$. If $f(S(x)) \cap f(S(y)) \neq \emptyset$, then $\rho'(f(x), f(y))$ is bounded, so $\rho'(\pi'(\sigma(t)), \pi'(\sigma(u)))$ and hence $|\sigma(t) - \sigma(u)|$ is bounded, and there is nothing to prove. If not, let $R' = [f(S(x)), f(S(y))]$ be the compact region between $f(S(x))$ and $f(S(y))$. Thus R' is a bounded Hausdorff distance from the geodesic segment $\pi'[\sigma(t), \sigma(u)]$. Now, for homological reasons, $f|f^{-1}R'$ must have degree 1, and so $R' \subseteq f(R)$. Since $f(\Theta(P)) = \Theta(M)$, we see that $R' \subseteq \Theta(M)$. Suppose now that $v \in [\sigma(t), \sigma(u)] \cap \sigma(Q)$. Let $v = \sigma(s)$ for $s \in Q$. Thus $S(\pi(s))$ meets some Margulis tube $T \in \mathcal{T}(P)$. Now $f(S(\pi(s)))$ has bounded diameter, and is a bounded distance from $\pi'(v)$ and from $f(T)$. But $f(T) \cap R' = \emptyset$. It follows that $\pi'(v)$ must be a bounded distance from an endpoint of the segment $\pi'([\sigma(t), \sigma(u)])$, and so v is a bounded distance from either $\sigma(t)$ or $\sigma(u)$ as required.

We have verified properties (Q0)–(Q4) for the map, $\sigma : \mathbf{R} \longrightarrow \mathbf{R}$, and so it is a quasi-isometry, proving Lemma 9.2.

We remark that an immediate consequence is that the map $f : (\Psi(P), \rho) \longrightarrow (\Psi(M), \rho')$ is a quasi-isometry. ■

We need a version of this result for bands. We can express this by passing to appropriate covers.

Let $B \subseteq \Psi(P)$ be a band with base surface Φ . Let $\mathcal{T}_B(P) = \{T \in \mathcal{T}(P) \mid T \cap \partial_V B \neq \emptyset\}$. Thus $\partial_V B$ consists of a set of annuli in the boundaries of elements of $\mathcal{T}_B(P)$ and components of $\partial\Psi(P)$. Let $\Xi_B(P) = \Psi(P) \setminus \text{int} \bigcup \mathcal{T}_B(P)$, and let $\Psi_B(P)$ be the cover of $\Xi_B(P)$ corresponding to B . Thus B lifts to a compact subset of $\Psi_B(P)$, which we also denote by B . We note that B cuts $\Psi_B(P)$ into two non-compact components bounded by $\partial_- B$ and $\partial_+ B$ respectively. The inclusion of B into $\Psi_B(P)$ is a homotopy equivalence. Indeed, if we remove those boundary components of $\Psi_B(P)$ that do not meet B , then the result is homeomorphic to $\Phi \times \mathbf{R}$.

We can perform the same construction in $\Psi(M)$. We let $\mathcal{T}_B(M) = \{f(T) \mid T \in \mathcal{T}_B(P)\}$, and $\Xi_B(M) = \Psi(M) \setminus \text{int} \bigcup \mathcal{T}_B(M)$, and by Lemma 3.7, $f : \Xi_B(P) \longrightarrow \Xi_B(M)$ is a homotopy equivalence. We let $\Psi_B(M)$ be the cover of $\Xi_B(M)$ corresponding to $\Psi_B(P)$, so that f lifts to a homotopy equivalence $\tilde{f} : \Psi_B(P) \longrightarrow \Psi_B(M)$. Indeed, $\Psi_B(M)$ is homeomorphic to $\Phi \times \mathbf{R}$ after removing certain boundary components, so we are in the same topological situation as before. We write $g : B \longrightarrow \Psi_B(M)$ for its restriction to B .

We shall assume that B has positive height. Let $\mathcal{T}_B^0(P) = \{T \in \Theta(P) \setminus \mathcal{T}_B(P) \mid T \cap B \neq \emptyset\}$. If $T \in \mathcal{T}_B^0(P)$, then either $T \subseteq B$, or it $T \cap B$ is a half torus bounded by an annulus in $\partial_H B \cap T$. We denote the lifted riemannian metric d by d_B , and write ρ_B for the modified metric obtained by setting the metric equal to zero on each $T \in \mathcal{T}_B^0(P)$. Every point $x \in B$ lies in a fibre $F(x) \subseteq B$, of bounded ρ_B -diameter. If $x \in \partial_\pm B$, we can take $F(x) = \partial_\pm B$. We can assume that such fibres foliate B . We similarly define a metric ρ'_B on $\Psi_B(M)$. In what follows, all distances are measured with respect to ρ_B or ρ'_B , unless otherwise specified. Note that $g : (B, \rho_B) \longrightarrow (\Psi_B(M), \rho'_B)$ is lipschitz.

Let $\pi : [a, b] \longrightarrow B \subseteq \Psi_B(P)$ be a shortest geodesic from $\partial_- B$ to $\partial_+ B$. Each fibre, $F(x)$, of B meets $\pi([a, b])$ and we see that B lies in a uniform neighbourhood of $\pi([a, b])$. Let $\pi' : \mathbf{R} \longrightarrow \Psi_B(M)$ be a bi-infinite geodesic, with $\pi'|(-\infty, 0]$ and $\pi'|[0, \infty)$ going out a negative and positive end of $\Psi_B(M)$ respectively. If $x \in B$, then $g(F(x))$ intersects $\pi'(\mathbf{R})$. This enables us to define a continuous map $\sigma_B : [a, b] \longrightarrow \mathbf{R}$ such that $\rho'_B(\pi'(\sigma_B(t)), g(\pi(t)))$ is uniformly bounded.

Lemma 9.3 : *The map $\sigma_B : [a, b] \longrightarrow \mathbf{R}$ is uniformly quasi-isometric.*

Proof : We define $Q = \{t \in [a, b] \mid F(\pi(t)) \cap \bigcup \mathcal{T}_B^0(P)\} \neq \emptyset$. Property (Q0) is an immediate consequence of the fact that g is lipschitz. We need to verify (Q1)–(Q4). The argument is essentially the same as before. There are a few subtleties we should comment on.

(Q1): Here we use the loops $\gamma_x^B \subseteq B$ instead of γ_x . This time property (W8) of Section 4 tells us that $\rho_B(x, y)$ is bounded above in terms of the distance between $[\gamma_x^B]$ and $[\gamma_y^B]$ in the curve graph $\mathcal{G}(\Phi)$ or the modified curve graph $\mathcal{G}'(\Phi)$ if $\kappa(\Phi) = 1$.

There is a slight complication in that $g^{-1}\beta \subseteq B$ might not connect T_0 to T_1 . We may therefore need to allow the path α to have up to three components, possibly connecting T_0 or T_1 to $\partial_H B$, and maybe also $\partial_- B$ with $\partial_+ B$. But this makes no essential difference to the argument, since if x, y both lie in $\partial_- B$ or both in $\partial_+ B$, then $[\gamma_x^B]$ and $[\gamma_y^B]$ are equal or adjacent in the curve graph. It follows that a uniform r -neighbourhood of $Y \subseteq X(\Phi)$ is connected and the argument proceeds as before.

(Q2): As before.

(Q3): Here we apply [FHS] as before to find embedded surfaces in $\Psi_B(M)$ close to the surfaces $g(F(x))$. These surfaces will be fibres in $\Psi_B(M)$ and any two disjoint fibres bound a band.

There is, however, an added complication in that in order to find our surfaces Z_i^\pm , we will need that t_1 and u_1 lie in $\sigma([a, b])$. We therefore need that t and u are not too close to the boundary of $\sigma([a, b])$. We are saved by property (Q2) which we have already proven.

Let $\sigma([a, b]) = [a', b']$, and suppose that $t, u \in I$ with $\rho'_B(\sigma(t), \sigma(u)) \leq k$ and that $N([\sigma(t), \sigma(u)], 2k_3) \cap \sigma(Q) = \emptyset$. Suppose that $\sigma(t) \in N(a', k_3)$, say. If $\sigma([t, u]) \subseteq N(a', k_3)$, then $|t - u|$ is bounded using (Q2). If not, let $s \in [t, u]$ be the first time that $\sigma|[t, u]$ leaves $N(a', k_3)$. Now $\sigma([t, s]) \subseteq N(a', k_3)$, and so by (Q2) $|s - t|$ is bounded. We can now replace t by s and continue the argument. We can do the same for the other end u . We are then reduced to the case where $t, u \in [a' + k_3, b' - k_3]$, and the argument proceeds as before. This time, the constant $2k_3$ becomes our “new” k_3 .

(Q4): As before. ◇

This shows that $\sigma_B : [a, b] \rightarrow \mathbf{R}$ is quasi-isometric, and it follows that $\sigma_B([a, b])$ lies in a bounded neighbourhood of $[\sigma_B(a), \sigma_B(b)]$. The constants depend only on the various constants inputted. To simplify notation in what follows, we will assume that $\sigma(a) < \sigma(b)$. If not, we could interpret everything by reversing the order on the range. (It will turn out, retrospectively, that if $b - a$ is sufficiently large then this is necessarily the case, though we won't formally need to worry about this. The issue of vertical orientation of bands will eventually be taken care of automatically by the topology of the situation.)

If $b - a$ is sufficiently large, then we can find embedded fibres close to $f(\partial_\pm B)$ in $\Psi_B(M)$ which will bound a band. We would like to find an embedded band B' in $\Psi(M)$. It would be enough to show that the projection to $\Xi_B(M)$ is injective far enough away from $\partial_H B'$. For this, we use need the following lemma. For the statement, we can interpret the term “band” A to be a 3-manifold, Ψ , homeomorphic to $\Phi \times [0, 1]$ with $\partial_V A \equiv \partial\Phi \times [0, 1]$ and $\partial_H A = \partial_- A \sqcup \partial_+ A = \Phi \times \{0, 1\}$. As usual, a “subband” is a subset bounded by disjoint fibres. We shall assume that A carries a metric ρ . This need not be a path metric. (For our application it will be the restriction of an ambient path-metric.)

Lemma 9.4 : *Suppose that (A, ρ) is a band and that each point of A lies in a fibre of (extrinsic) diameter at most k . Let Ξ be a complete non-compact riemannian manifold, and suppose that $\theta : A \rightarrow \Xi$ is a locally isometric map with $\partial_V A = \theta^{-1}(\partial\Xi)$ (and so θ is 1-lipschitz). Suppose that any fibre of A is homotopic to an embedded surface in Ξ . Then $\theta|_{B \setminus N(\partial_H A, 13k)}$ is injective.*

Proof : First, we claim that if $x \in A$ with $\rho(x, \partial_H A) \geq 13k$, then there is some sub-band, $A' \subseteq A$ (with the same base surface), containing x , with $\rho'(\theta(x), \theta(\partial_H A')) \geq 2k$. To see this, we can argue as follows.

Let C_{\pm} be the set of $p \in A$ such that there is an arc $\tau_{\pm}(p)$ from p to $\partial_{\pm} A$ such that $\theta(\tau_{\pm}(p))$ is geodesic in Ξ . Clearly $\partial_{\pm} A \subseteq C_{\pm}$ (set $\tau_{\pm}(p) = \{p\}$). Also $A = C_- \cup C_+$, since if $p \in A$, we can find a geodesic ray σ in Ξ based at p in Ξ (since Ξ is non-compact). Some component of $\theta^{-1}(\sigma)$ must connect p to $\partial_H A$ in A , and so $p \in C_- \cup C_+$. But now A is connected, and C_- and C_+ are closed, so there must be some $p \in C_- \cap C_+$. In other words, there are arcs, τ_{\pm} from p to points $a_{\pm} \in \partial_{\pm} A$ with $\theta(\tau_{\pm})$ geodesic in Ξ . Since $\tau_+ \cup \tau_-$ connects $\partial_- A$ to $\partial_+ A$ we have $A = N(\tau_- \cup \tau_+, k)$. Suppose now that $x \in A$ with $\rho(x, \partial_H A) \geq 13k$. Let $y \in \tau_- \cup \tau_+$ with $\rho(x, y) \leq k$, so $\rho'(\theta(x), \theta(y)) \leq k$. We can assume that $y \in \tau_+$. Now, $\rho'(\theta(x), \theta(y)) \geq 12k$, and so, since $\theta(\tau_+)$ is geodesic, $\rho'(\theta(y), \theta(a_+)) \geq 12k$. If $\rho'(\theta(y), \theta(p)) \leq 4k$, let F be a fibre of diameter at most k through p . Note that $\rho'(\theta(x), \theta(F)) \geq 4k - 2k = 2k$, so we can set A' to be the band between F and $\partial_+ A$. On the other hand, suppose $\rho'(\theta(y), \theta(p)) \leq 4k$. The total length of τ_- together the segment of τ_+ from p to y is at least $\rho(y, a_+) \geq 12k$. Since $\theta(\tau_-)$ and $\theta(\tau_+)$ are both geodesic, it follows that $\rho'(\theta(y), \theta(a_-)) \geq 12k - 8k = 4k$. This time, we can take $A' = A$. This proves the claim.

Now if $x \in A$ with $\rho(x, \partial_H A) \geq 13k$, we claim there is a fibre $Z = Z(x)$ through x , with $\theta|Z$ injective. To see this, start with any fibre, F , through x . By hypothesis, $\theta(F)$ is homotopic to an embedded surface in Ξ , and so by [FHS] we can find such a surface S , in an arbitrarily small neighbourhood of $\theta(F)$. We shall assume that $\theta(x) \in S$. Since $\text{diam}(S) < 2k$, we have $S \cap \theta(\partial_H A') = \emptyset$. Now $\theta|A'$ is a local homeomorphism away from $\partial_H A'$, and so S lifts to an embedded surface, Z in $A' \subseteq A$ with $x \in Z$. Since the inclusion of Z into A is π_1 -injective, it follows that F must be a fibre of A . Thus, the map $\theta|Z \rightarrow S$ is a homeomorphism. In particular, $\theta|Z$ is injective as required.

Note also that, in the above construction, if $\theta(x) = \theta(y)$ we could take the same surface S for both, and we get fibres $Z(x)$ and $Z(y)$ with $\theta|Z(x)$ and $\theta|Z(y)$ both homeomorphisms to S . If $x \neq y$, then we must have $Z(x) \cap Z(y) = \emptyset$.

Suppose finally for contradiction, that $x, y \in A$ with $x \neq y$, $\rho(x, \partial_H A) \geq 13k$, $\rho(y, \partial_H A) \geq 13k$ and $\theta(x) = \theta(y)$. Construct fibres $Z(x)$ and $Z(y)$ as above. Let $C = [Z(x), Z(y)]$ be the band between $Z(x)$ and $Z(y)$. We construct a closed manifold R by gluing together $\partial_- C = Z(x)$ and $\partial_+ C = Z(y)$ via the homeomorphism $(\theta|Z(y))^{-1} \circ (\theta|Z(x))$. Now θ induces a map from R to Ξ with is a local homeomorphism away from $Z(x) \equiv Z(y)$, and hence, by orientation considerations, a local homeomorphism everywhere. It must therefore be a covering space, giving the contradiction that Ξ is compact. \diamond

We now apply this in the situation of interest. We return to the set-up of Lemma 9.3. We will need to assume that $b - a$ is sufficiently large. As described earlier, we will assume, for notational convenience that $\sigma(a) < \sigma(b)$. The first step is to note that $\partial_- B = F(\pi(a))$ and so $g(\partial_- B)$ is a bounded distance from $\pi'(\sigma(a))$. Similarly, $\partial_+ B = F(\pi(b))$ is a bounded distance from $\pi'(\sigma(b))$. Thus if $b - a$ and hence $\sigma(b) - \sigma(a)$ is sufficiently large, $g(\partial_- B) \cap g(\partial_+ B) = \emptyset$. We can find disjoint fibres, Z_{\pm} in a small neighbourhood of $g(\partial_{\pm} B)$, and let A be the band between Z_- and Z_+ . Let $\theta : A \rightarrow \Xi_B(M)$ be the inclusion of A into $\Psi_B(M)$ composed with the covering map $\Psi_B(M) \rightarrow \Xi_B(M)$. Now

if l_0 , and $b - a - 2l_0$ are sufficiently large, then by Lemma 9.3, we can arrange that $\sigma(a) < \sigma(a + l_0) < \sigma(b - l_0) < \sigma(b)$. Moreover, $g(F(\pi(a + l_0)))$ is bounded diameter and a bounded distance from $\pi(\sigma(a + l_0))$ and we can find a fibre, Z'_- , close to $g(F(\pi(a + l_0)))$. We similarly find Z'_+ close to $g(F(\pi(b - l_0)))$. If l_0 and $b - a - 2l_0$ are sufficiently large, then we will have $Z_- < Z'_- < Z'_+ < Z_+$. Moreover, by Lemma 9.4, we can assume that $\theta|_{B'}$ is injective, where $B' = [Z'_-, Z'_+]$. Now, if $l_1 > 0$ is sufficiently large, we can assume a uniform neighbourhood $\sigma_B([a + l_0 + l_1, b - l_0 - l_1])$ lies inside $[\sigma_B(a + l_0), \sigma_B(b - l_0)]$. If this neighbourhood is large enough, then $g(B) \subseteq B'$, where $B_0 = [F(a + l_0 + l_1), F(b - l_0 - l_1)]$. Note that the depth of B_0 in B (measured in the metric ρ_B) is equal to $l_0 + l_1$ up to an additive constant. Since B' embeds in $\Psi_B(M)$, we can project the whole picture to $\Xi_B(P) \subseteq \Psi(P)$ and $\Xi_B(M) \subseteq \Psi(M)$.

In summary, we have shown:

Lemma 9.5 : *There is some $l > 0$ such that if $B \subseteq \Psi(P)$ is a band, and B_0 is a sub-band of depth at least l , then there is a band $B' \subseteq \Psi(M)$ with the same base surface such that $f(B_0) \subseteq B'$. \diamond*

For the moment, we can just interpret this to mean that $B' \cong \Phi \times [0, 1]$ with $\partial_V B' = B' \cap \mathcal{T}_B(M)$. (In Section 10, we will insist in addition that the horizontal boundaries of bands should lie in the thick part.) In fact, from our construction of B' means that we see that every point of B' lies inside some fibre of bounded (extrinsic) diameter. By [FHS] we can take such fibres to be embedded (but we do not claim that such fibres foliate B'). We can also refine Lemma 9.5 in various ways. Note, in particular, that if $r \geq 0$, then by choosing $l = l(r)$ sufficiently large, we can assume that an r -neighbourhood of $f(B_0)$ (with respect to ρ') lies inside B' .

At this point, it is still conceivable that f might also send a point far away from B into B' . To rule this out, we need to bring (W9) of Section 4 into play.

Suppose now that B is a maximal band. There is some $r \geq 0$ such that if $x \in \Psi(P) \setminus B$, there there is some loop, $\delta_x^B \ni x$, of d -length at most r and such that either δ_x^B is freely homotopic into a Margulis tube T not meeting B , or else $[\delta_x^B]$ is not freely homotopic into the base surface, Φ , in Σ . We can assume that B' does not meet any of the images of Margulis tubes of the above type, and so there is a bound on how deeply the image $f(\delta_x^B)$ can enter into B' . Using the refinement of Lemma 9.5 mentioned above, we see:

Lemma 9.6 : *There is some $l' > 0$ such that if $B \subseteq \Psi(P)$ is a maximal band and if $B_0 \subseteq B$ is a parallel subband of depth at least l' in B (with respect to the metric ρ) then there is a band, $A \subseteq \Psi(M)$, with the same base surface such that $f(B_0) \subseteq A$ and $f(\Psi(P) \setminus B) \cap A = \emptyset$. \diamond*

10. Controlling the map on thick parts.

In this section, we shall show that the map $f : \Theta(P) \rightarrow \Theta(M)$ as defined in Section 7 is universally sesquiplisclitz (Proposition 10.11). To this end, we will take the results of Section 9, and get ourselves into a position to apply Lemma 8.4.

Before we begin, we need a few observations about surfaces in 3-manifolds. As usual, so simplify the discussion, our surfaces and 3-manifolds will be assumed orientable. Let Θ be an irreducible 3-manifold with boundary $\partial\Theta$. By a *proper surface*, in Θ we mean a properly embedded π_1 -injective surface, $S \subseteq \Theta$, with $\partial S = S \cap \partial\Theta$. Our orientability assumption means that S is 2-sided, i.e. extends to an embedding of $S \times [-1, 1]$ in Θ . We say that two compact proper surfaces, S and S' are *parallel* if they are disjoint and homotopic in Θ relative to $\partial\Theta$. By Waldhausen's cobordism theorem, $S \cup S'$ bounds a product region $S \times [0, 1] \subseteq \Theta$ (meeting $\partial\Theta$ in $\partial S \times [0, 1]$).

Suppose that \mathcal{S} is a locally finite collection of pairwise disjoint proper surfaces. We construct the dual graph $\sigma = \sigma(\mathcal{S})$ with edge set $E(\sigma) = \mathcal{S}$ and vertex set the set of components of $\Theta \setminus \bigcup \mathcal{S}$, by deeming an edge to be incident on a vertex if the corresponding surface is a boundary component of the corresponding complementary region. (We allow loops and multiple edges.)

Let $\tilde{\Theta}$ be the universal cover of Θ . The set \mathcal{S} lifts to a set, $\tilde{\mathcal{S}}$, of proper surfaces in $\tilde{\Theta}$. We write $\tau(\mathcal{S}) = \sigma(\tilde{\mathcal{S}})$. In this case, $\tau(\mathcal{S})$ is a tree, and $G = \pi_1(\Theta)$ acts on $\tilde{\mathcal{S}}$ with quotient $\sigma(\mathcal{S})$. Note also that if $\mathcal{S}_0 \subseteq \mathcal{S}$ is a maximal collection of parallel compact surfaces in \mathcal{S} , then the edges corresponding to \mathcal{S}_0 form an arc in $\sigma(\mathcal{S})$, and lifts to a collection of arcs in $\tau(\mathcal{S})$. We can view $\tau(\mathcal{S})$ as obtained by subdividing edges of the tree formed by removing all but one of the surfaces of \mathcal{S}_0 .

The above is a standard construction of Bass-Serre theory — the collection \mathcal{S} determines a splitting of G as a graph, $\sigma(\mathcal{S})$, of groups, and $\tau(\mathcal{S})$ is the associated Bass-Serre tree.

Lemma 10.1 : *Suppose that Θ and Θ' are 3-manifolds and that $f : \Theta \rightarrow \Theta'$ is a proper homotopy equivalence. Suppose that \mathcal{S} and \mathcal{S}' are locally finite collections of pairwise disjoint non-parallel compact proper non-sphere surfaces in Θ and Θ' respectively. Suppose that $\bar{f} : \mathcal{S} \rightarrow \mathcal{S}'$ is a bijection so that $\bar{f}(S)$ is homotopic to $f(S)$ for all $S \in \mathcal{S}$. Then there is a graph-of-groups isomorphism $g : \sigma(\mathcal{S}) \rightarrow \sigma(\mathcal{S}')$, which agrees with f on $E(\sigma(\mathcal{S})) = \mathcal{S} \rightarrow E(\sigma(\mathcal{S}')) = \mathcal{S}'$.*

Here we can interpret a “graph-of-groups isomorphism” to mean a G -equivariant isomorphism, $\tilde{g} : \tilde{\mathcal{S}} \rightarrow \tilde{\mathcal{S}'}$, where $G = \pi_1(\Theta) \cong \pi_1(\Theta')$, which projects to a graph isomorphism $g : \sigma(\mathcal{S}) \rightarrow \sigma(\mathcal{S}')$.

To prove this we need the following observation. Suppose that Θ_0 is a codimension-0 submanifold of Θ , with relative boundary $\partial_R \Theta_0$ in Θ . Suppose that the $\partial_R \Theta_0$ is compact and incompressible and has no sphere components, and that $f|_{\partial_R \Theta_0}$ is homotopic to an embedding $h : \partial_R \Theta_0 \rightarrow \Theta'$. Then, $h(\partial \Theta_0)$ bounds a codimension-0 submanifold, $\Theta'_0 \subseteq \Theta'$, with $\pi_1(\Theta_0) = f_*(\pi_1(\Theta_0))$. This is essentially a homology argument. Note that we can extend h to a proper map $h : \Theta_0 \rightarrow \Theta'$ which is homotopic to $f|_{\Theta_0}$. We can now let Θ'_0 be the set of points onto which h maps with degree 1. Since the relative boundaries of Θ_0 and Θ'_0 are incompressible we can identify $\pi_1(\Theta_0) \subseteq \pi_1(\Theta)$ and $\pi_1(\Theta'_0) \subseteq \pi_1(\Theta')$. We can now lift this picture to the covers of Θ and Θ' corresponding to $\pi_1(\Theta_0)$ and $f_*(\pi_1(\Theta_0))$. By degree arguments we see that Θ'_0 lifts, and that the cover retracts onto this lift. The observation now follows easily.

Proof of Lemma 10.1 : Let $\tau = \tau(\mathcal{S})$ and $\tau' = \tau(\mathcal{S}')$, and $G = \pi_1(\Theta) = \pi_1(\Theta')$ as above. We already have a bijection $\tilde{g} : E(\tau) \equiv \tilde{\mathcal{S}} \longrightarrow E(\tau') \equiv \tilde{\mathcal{S}'}$, and we need to check that this preserves adjacency.

Let v be a vertex of $\sigma(\mathcal{S})$ with group $H \subseteq G$. This corresponds to a component, Θ_v of $\Theta \setminus \bigcup \mathcal{S}$ with $\pi_1(\Theta_v) \equiv H$. By the above observation, there is a submanifold, $\Theta'_v \subseteq \Theta'$ with $\pi_1(\Theta'_v) = H$, bounded by the subsurfaces $\bar{f}(S)$ where S runs through the boundary components of Θ_v . (It is possible that these manifolds may have pairs boundary components that are parallel on the outside, and hence identified to a single surface of our collection.) We claim that Θ'_v contains no other surface of \mathcal{S} . For suppose that $S \in \mathcal{S}$ were such that $\bar{f}(S)$ lies in the interior of Θ'_v . We see that $\pi_1(S) \subseteq H$. We can lift Θ_0 and S to the cover the cover $\hat{\Theta} = \tilde{\Theta}/H$ corresponding to H . Note that $\hat{\Theta}$ retracts onto Θ_0 . Now S lies outside Θ_0 but is homotopic into it. Since it is a properly embedded compact π_1 -injective surface, one deduces that S is homotopic to a boundary component of Θ_0 . Back down in Θ we see that S is parallel to one of the boundary components of Θ_0 contradicting our hypothesis that no two elements of \mathcal{S} are parallel.

This shows that the interiors manifolds Θ'_v are all disjoint in Θ' . We see that (at least on the level of fundamental groups) the combinatorial situation in Θ' is identical to that in Θ , and so we get a graph-of-groups isomorphism. \diamond

Suppose that $f : \Theta \longrightarrow \Theta'$ is a proper homotopy equivalence. Suppose that \mathcal{S} is a finite disjoint collection of compact proper surfaces, and that $f(S) \cap f(S') = \emptyset$ for all $S, S' \in \mathcal{S}$. Suppose that for all $S \in \mathcal{S}$, $f(S)$ is homotopic (relative to $\partial\Theta'$) to an embedding in Θ' . By [FHS], we can find an embedded surface, $Z(S)$, homotopic to $f(S)$ in an arbitrarily small neighbourhood of $f(S)$. Let $\mathcal{S}' = \{Z(S) \mid S \in \mathcal{S}\}$. These surfaces are all disjoint.

If no two surfaces in \mathcal{S} are parallel, we are in a position to apply Lemma 10.1. This breaks down if we allow parallel surfaces since f might fold up the associated product region in Θ . However, we can take care of this using the following criterion:

Lemma 10.2 : *Suppose that \mathcal{S} is a finite collection of disjoint compact proper non-sphere surfaces in Θ , and that $f : \Theta \longrightarrow \Theta'$ is a homotopy equivalence, such that the surfaces $f(S)$ for $S \in \mathcal{S}$ are all disjoint, and each is homotopic to an embedded surface in Θ' . Suppose that whenever $S, S' \in \Theta$ are parallel there is a properly embedded ray (i.e. semi-infinite path) $\alpha \subseteq \Theta$ based in S , such that $f(\alpha) \cap f(S') = \emptyset$. Let $\mathcal{S}' = \{Z(S) \mid S \in \mathcal{S}\}$ be constructed as above. Then there is a graph-of-groups isomorphism from $\sigma(\mathcal{S})$ to $\sigma(\mathcal{S}')$ sending an edge corresponding to $S \in \mathcal{S}$ to the edge corresponding to $Z(S) \in \mathcal{S}'$.*

Proof : Let \mathcal{S}_0 be a transversal to the parallel relation, and let $\mathcal{S}'_0 = \{Z(S) \mid S \in \mathcal{S}_0\}$. By Lemma 10.1, there is a graph-of-groups isomorphism $\sigma(\mathcal{S}_0) \longrightarrow \sigma(\mathcal{S}'_0)$. Now $\sigma(\mathcal{S})$ is obtained from $\sigma(\mathcal{S}_0)$ by subdividing edges with valence-2 vertices. We similarly get from $\sigma(\mathcal{S}'_0)$ to $\sigma(\mathcal{S})$. One easily checks that the above criterion ensures that the orders of these subdivision is the same on both sides. \diamond

We now move on the the case of interest, namely where $\Theta = \Theta(P)$ and $\Theta' = \Theta(M)$, and $f : \Theta(P) \longrightarrow \Theta(M)$ is the map constructed in Section 7. Recall that we have riemannian path metrics, d and d' on $\Theta(P)$ and $\Theta(M)$ respectively, and that ρ and ρ' are the metrics obtained by shrinking each margulis tube to diameter 0. The map f

is uniformly lipschitz from $(\Theta(P), d)$ to $(\Theta(M), d')$, and hence also from $(\Theta(P), \rho)$ to $(\Theta(M), \rho')$. We write $G = \pi_1(\Theta(P)) \equiv \pi_1(\Theta(M))$. We shall say that a subset of $\Theta(P)$ (or $\Theta(M)$) is k -small if its diameter in the metric ρ (or ρ') is at most k .

Given $x \in \Theta(P)$, we can find a uniformly small fibre $S(x) \subseteq \Theta(P)$. This can be achieved by taking a horizontal fibre, and then pushing it slightly off any Margulis tube. This may significantly increase its d -diameter, but only increases the ρ -diameter by an arbitrarily small amount. (Note that this is slightly different from the notion used in Section 9. Here we are assuming that $S(x)$ lies in the thick part. This greatly simplifies the description of various topological operations. The cost is that we can no longer assume that $S(x)$ varies continuously in x , but that will not matter to us in this section.) By Lemma 7.4, we can find a proper surface $S'(x) \subseteq \Theta(M)$ in an arbitrarily small neighbourhood of $f(S(x))$, and homotopic to $f(S(x))$ in $\Theta(M)$. Note that $S'(x)$ is a fibre in the product space $\Psi(M)$. It is also uniformly small in $\Theta(M)$.

We now use the following construction of expanding bands in $\Theta(P)$. We fix a constant, h_0 , to be defined shortly. Given $x \in \Theta(P)$, set $R_x[0] = S(x)$, and $R'_x[0] = S'(x)$. Let $\pi : \mathbf{R} \rightarrow \Theta(P)$ be a bi-infinite geodesics respecting the ends of $\Theta(P)$. Since π must cross $R_x[0]$, we can assume that $\pi(0) \in R_x[0]$. Given $n \in \mathbf{Z}$, let $S_n = S(\pi(nh_0))$ and $S'_n = S'(\pi(nh_0))$. If h_0 is large enough, the surfaces S_n will all be disjoint and occur in the correct order in $\Psi(P)$. Let $R_x[n] = \Theta(P) \cap [S_{-n}, S_n]$, in other words, the compact region of $\Theta(P)$ bounded by S_{-n} and S_n . This gives an increasing sequence, $R_x[0] \subseteq R_x[1] \subseteq R_x[2] \subseteq \dots$ of bands that eventually exhaust $\Theta(P)$. Now applying Lemma 9.2, again if h_0 is large enough, the surfaces S'_n are all disjoint and occur in the correct order in $\Theta(M)$, and we similarly construct bands $R'_x[n] = \Theta(M) \cap [S'_{-n}, S'_n]$. We write $CR_x[n]$ and $CR'_x[n]$ for the closures of $\Theta(P) \setminus R_x[n]$ and $\Theta(M) \setminus R'_x[n]$ respectively. We can assume that:

Lemma 10.3 : *For all $x \in \Theta(M)$ and an all $n \in \mathbf{N}$ we have*

- (1) $f(R_x[n]) \cap CR'_x[n+1] = \emptyset$,
- (2) $f(CR_x[n+1]) \cap R'_x[n] = \emptyset$.

Proof : This is a simple consequence of the discussion of end-separating sets in Section 9, and the fact that σ is a quasi-isometry (Lemma 9.2). \diamond

Note in particular, that $f(R_x[n]) \cap f(R_x[n+2]) = \emptyset$. It will also be convenient to fix some $\eta > 0$ smaller than the injectivity radii of $\Theta(P)$ and $\Theta(M)$, and we can refine Lemma 10.3 slightly to say that $d'(f(R_x[n]), CR'_x[n+1]) \geq \eta$ and $d'(f(CR_x[n+1]), R'_x[n]) \geq \eta$. We also may as well assume that $d(f(x), f(R_x[1])) \geq \eta$.

We can also make a stronger statement concerning the nesting of the regions $R_x[n]$. See Lemma 10.6 below.

We have a similar process of shrinking bands. Let B be a maximal band in $\Psi(P)$, with base surface Φ . Let ρ_B be the metric on B described in Section 9. (Recall that this is essentially obtained by taking the path metric induced by d , and then shrinking each Margulis tubes whose interior meets B to have diameter 0. We may need to modify the metric near $\partial_H B$ to take account of the fact that we cannot easily control the local geometry of the surfaces $\partial_{\pm} B$. It is formally defined by passing to the appropriate covering space.) In this case, we use “small” to refer to diameter with respect to the metric ρ_B .

Now each $x \in B$ lies in some uniformly small surface $F(x)$ in $B \cap \Theta(P)$ that is a fibre for B . Applying Lemma 7.4 again, we can find a surface $F'(x)$ in an arbitrarily small neighbourhood of $f(F(x))$ in $\Theta(M)$, and homotopic to $f(F(x))$ in $\Theta(M)$.

Now let $\pi_B : [a^-, a^+] \rightarrow B \cap \Theta(P)$ be a shortest geodesic from $\partial_- B$ to $\partial_+ B$ in $B \cap \Theta(P)$, with respect to the metric ρ_B . We write $h(B) = a^+ - a^-$ for the length of this geodesic. We fix h_0 and h_1 as described below, and set $h(n) = 2h_1 + (2n + 1)h_0$. Suppose $h(B) \geq h(m)$. For each $n = 0, 1, \dots, m$, let $F_{n,\pm} = F(\pi_B(a^\pm \pm (h_1 + nh_0)))$, and let $F'_{n,\pm} = F'(\pi_B(a^\pm \pm (h_1 + nh_0)))$. If h_0 is big enough then the surfaces $F_{0,-}, F_{1,-}, \dots, F_{m,-}, F_{m,+}, \dots, F_{1,+}, F_{0,+}$ are all disjoint and occur in this order in B . Let $B[n] = B_\Phi[n] = [F_{n,-}, F_{n,+}] \cap \Theta(M)$ be the compact region of $\Theta(P)$ bounded by $F_{n,-}$ and $F_{n,+}$. Thus $B[n] \subseteq B$. Thus, $B[m] \subseteq \dots \subseteq B[1] \subseteq B[0] \subseteq B \cap \Theta(P)$. By applying Lemma 9.6, if h_1 is big enough we can assume that $f(B[0])$ lies inside a band A in $\Theta(M)$, and that $f(\Theta(P) \setminus B)$ does not enter A . We are thus effectively reduced to considering the map of $f|B[0]$ into A . Applying Lemma 9.3, we can assume that the surfaces $F'_{0,-}, F'_{1,-}, \dots, F'_{m,-}, F'_{m,+}, \dots, F'_{1,+}, F'_{0,+}$ are disjoint and occur in this order in B . We set $B'[m] = B'_\Phi[n] = [F_{n,-}, F_{n,+}] \cap \Theta(M)$, and so $B'[m] \subseteq \dots \subseteq B'[1] \subseteq B'[0] \subseteq A \cap \Theta(M)$. We write $CB[n]$ and $CB'[n]$ for the closures of $\Theta(P) \setminus B[n]$ and $\Theta(M) \setminus B'[n]$ respectively. Again, if h_0 and h_1 are large enough, we have:

Lemma 10.4 : *For each maximal band, and for all n , we have:*

- (1) $f(B_x[n]) \cap CB'_x[n+1] = \emptyset$,
- (2) $f(CB_x[n+1]) \cap B'_x[n] = \emptyset$.

◇

The above bands are defined provided $h(n+1) \leq h(B)$. If $h(n) > h(B)$, we can set $B[n] = \emptyset$. This is consistent with Lemma 10.4.

We can also assume that the bands $B[0]$ lies inside a 1-collared band, $B_0 \subseteq B$. This means that the collection of bands $B[0]$ that we construct will have a nesting property (see Lemma 10.6(1)).

Finally, we can carry out the expanding band construction within a band. Suppose that B is a maximal band with base surface Φ . Suppose that $x \in B_\Phi[m]$. We set $R_{\Phi,x}[0]$ to be a uniformly small fibre containing x , which we can assume lies in $B_\Phi[n]$. As before, we construct increasing sequences of bands $R_{\Phi,x}[0] \subseteq R_{\Phi,x}[1] \subseteq \dots \subseteq R_{\Phi,x}[n]$ in $B \cap \Theta(P)$, and $R'_{\Phi,x}[0] \subseteq R'_{\Phi,x}[1] \subseteq \dots \subseteq R'_{\Phi,x}[n]$ in $A \cap \Theta(M)$. We can assume that $R_{\Phi,x}[n] \subseteq B_\Phi[m-n]$. As before, applying Lemma 9.3, we get:

Lemma 10.5 : *For each maximal band, and for all n , we have:*

- (1) $f(R_{\Phi,x}[n]) \cap CR_{\Phi,x}[n+1] = \emptyset$,
- (2) $f(CR_{\Phi,x}[n+1]) \cap R_{\Phi,x}[n] = \emptyset$.

◇

Let \mathcal{F} be the set of subsurfaces of Σ . Given $\Phi \in \mathcal{F}$, let $B_\Phi \subseteq \Psi(P)$ be the (possibly empty) maximal band with base surface Φ . Given $n \in \mathbf{N}$, let $\mathcal{F}[n] = \{\Phi \in \mathcal{F} \mid B_\Phi \neq \emptyset, h(B_\Phi) \geq h(n)\}$. Let $\mathcal{B}[n] = \{B_\Phi[n] \mid \Phi \in \mathcal{F}[n]\}$, and let $\mathcal{B}'[n] = \{B'_\Phi[n] \mid \Phi \in \mathcal{F}[n]\}$.

Definition : We refer to elements of $\mathcal{B}[n]$ and $\mathcal{B}'[n]$ as *level n bands* in $\Theta(P)$ and $\Theta(M)$ respectively.

Given $\phi \in \mathcal{F}$, we write $\mathcal{F}_\Phi \subseteq \mathcal{F}$ for the set of proper subsurfaces of Φ . Let $\mathcal{F}_\Phi[n] = \mathcal{F}_\Phi \cap \mathcal{F}[n]$, $\mathcal{B}_\Phi[n] = \{B_{\Phi'}[n] \mid \Phi' \in \mathcal{F}_\Phi\}$ and $\mathcal{B}'_\Phi[n] = \{B'_{\Phi'}[n] \mid \Phi' \in \mathcal{F}_\Phi\}$.

If we choose h_0 and h_1 large enough we have the following:

Lemma 10.6 :

- (1) If $A, B \in \mathcal{B}[0]$ are distinct, and $A \cap B \neq \emptyset$, the base surfaces, $\pi_\Sigma A$ and $\pi_\Sigma B$ are nested (one is proper subset of the other).
- (2) Suppose that $x \in \Theta(P)$ and $A \in \mathcal{B}[0]$. If $A \cap R_x[n] \neq \emptyset$ then $A \cap CR_x[n+1] = \emptyset$.
- (3) Suppose $\Phi \in \mathcal{F}[n+1]$, $x \in B_\Phi[n+1]$ and $A \in \mathcal{B}_\Phi[0]$. If $A \cap R_{\Phi,x}[n] \neq \emptyset$ then $A \cap CR_{\Phi,x}[n+1] = \emptyset$.
- (4) Suppose $\Phi \in \mathcal{F}[n+1]$ and $A \in \mathcal{B}_\Phi[0]$. If $A \cap CB_\Phi[n] = \emptyset$ then $A \cap CB_\Phi[n+1] = \emptyset$. \diamond

Lemma 10.7 : The same statement in $\Theta(M)$, with $R'_x[n]$ replacing $R_x[n]$ and $B'[n]$ replacing $B[n]$ etc. \diamond

We also have:

Lemma 10.8 : Suppose that $p, q \in \mathbf{N}$.

- (1) If $x \in \Theta(P)$, the volume of $R'_x[p] \setminus \bigcup \mathcal{B}[q]$ is bounded above in terms of p and q .
- (2) If $\Phi \in \mathcal{F}[p]$ and $x \in B_\Phi[p]$, then the volume $R_{\Phi,x}[p] \setminus \bigcup \mathcal{B}_\Phi[q]$ is bounded above in terms of p and q .

Proof : The riemannian notions of distance and volume (with respect to ρ) are linearly bounded in terms of the combinatorial notions used in Section 4. This is therefore a direct corollary of Lemma 4.5. \diamond

Lemma 10.9 : Suppose that $p, q \in \mathbf{N}$.

- (1) If $x \in \Theta(P)$, the volume of $R'_x[p] \setminus \bigcup \mathcal{B}'[q]$ is bounded above in terms of p and q .
- (2) If $\Phi \in \mathcal{F}[p]$ and $x \in B_\Phi[p]$, then the volume $R'_{\Phi,x}[p] \setminus \bigcup \mathcal{B}'_\Phi[q]$ is bounded above in terms of p and q .

Proof :

- (1) By Lemma 10.3(1) and 10.4(2) and the fact that f is surjective, we have

$$R'_x[p] \setminus \bigcup \mathcal{B}'[q] \subseteq f(R'_x[p+1] \setminus \bigcup \mathcal{B}'[q+1]).$$

Since f is uniformly lipschitz, the volume of the right hand side is bounded by Lemma 10.8(1), and the result follows.

(2) By Lemma 10.5(1) and 10.4(2) we have:

$$R'_{\Phi,x}[p] \setminus \bigcup \mathcal{B}'_{\Phi}[q] \subseteq R'_{\Phi,x}[p+1] \setminus \bigcup \mathcal{B}'_{\Phi}[q+1]$$

and the result follows by Lemma 10.8(2). \diamond

In fact, using Lemma 8.1, we see that we can also bound the volume of an η -neighbourhood of these sets in terms of η .

We can now set about verifying the hypotheses of Lemma 8.4. Fix some constant η less than the injectivity radius of $\Theta(M)$.

Proposition 10.10 : *Suppose $x, y \in \Theta(P)$ and $d'(f(x), f(y)) \leq \eta$. Then there is a path, α , in $\Theta(P)$, of bounded diameter with respect to the metric d , such that $f(\alpha) \cup [f(x), f(y)]$ bounds a disc in $\Theta(M)$ of bounded diameter with respect to the metric d' .*

The conclusion of Proposition 10.10 determines a homotopy class of path from x to y in $\Theta(P)$, which we shall refer to as the *right homotopy class*.

The basic strategy is to start with any path α from x to y in the right homotopy class. (Such a path exists, since f is a homotopy equivalence from $\Theta(P)$ to $\Theta(M)$.) We first push this into a region of bounded depth about x , and then push it off all bands of a given bounded depth. Lemma 10.8 then gives a bound on the diameter of such a path in $(\Theta(P), d)$. By our choice of α , $f(\alpha) \cup [f(x), f(y)]$ bounds a disc in $\Theta(M)$. We now push this disc into a region of bounded depth, and then off bands of bounded depth. Lemma 10.9 then bounds the diameter of this disc in $(\Theta(M), d')$. In practice we will only need to construct bands up to depth 10. (The proof Lemma 10.9 takes us up to depth 11.)

Let us first deal with the case where $x, y \notin \bigcup \mathcal{B}[6]$. We connect x to y by a path, α , in the right homotopy class. We write $R[n] = R_x[n]$. Now $d'(f(x), f(CR[1])) \geq \eta$ and so $y \in R[1]$. We first claim that we can push α into $R[2]$.

To see this we pass to the universal covers, $\tilde{f} : \tilde{\Theta}(P) \rightarrow \tilde{\Theta}(M)$. Let $\mathcal{S}[i] = \{\partial_- R[i], \partial_+ R[i]\}$ and $\mathcal{S}'[i] = \{\partial_- R'[i], \partial_+ R'[i]\}$. Let $\mathcal{S} = \mathcal{S}[1] \cup \mathcal{S}[2]$ and $\mathcal{S}' = \mathcal{S}'[1] \cup \mathcal{S}'[2]$. Write $\tau = \tau(\mathcal{S})$ and $\tau' = \tau(\mathcal{S}')$ for the corresponding Bass-Serre trees. Now $R[1] \subseteq R[2] \subseteq R[3]$, and so applying Lemma 10.3, given any two distinct surfaces, $S, S' \in \mathcal{S}$, it is easy to construct a ray β from S to infinity such that $f(\beta) \cap f(S') = \emptyset$. (We only really need to do this if S and S' are parallel in $\Theta(P)$.) By Lemma 10.2, it then follows that there is an equivariant isomorphism from τ to τ' .

Let $\tilde{\alpha}$ be a lift of the path α to $\tilde{\Theta}(P)$ connecting \tilde{x} to \tilde{y} . A lift, δ , of the short geodesic $[f(x), f(y)]$ connects $\tilde{f}(\tilde{x})$ to $\tilde{f}(\tilde{y})$ in $\tilde{\Theta}(M)$. Now if $\alpha \subseteq R[1]$, there is nothing to prove. If not, let $z, w \in \alpha$ be the first and last intersection of α with $\partial_H R[1]$ (the relative boundary of $R[1]$ in $\Theta(P)$). We have $z \in S_0 \in \mathcal{S}[1]$ and $w \in S_1 \in \mathcal{S}[1]$. Let β, γ be the subpaths of α from x to z and y to w respectively, and let $\tilde{z}, \tilde{w}, \tilde{\beta}, \tilde{\gamma}, \tilde{S}_0$ and \tilde{S}_1 be the lifts to $\tilde{\Theta}(P)$. Now $\beta, \gamma \subseteq R[1]$ and so, by Lemma 10.3, $f(\beta), f(\gamma) \subseteq R'[2]$. Thus $\tilde{f}(\tilde{\beta}) \cup \delta \cup \tilde{f}(\tilde{\gamma})$ is a path in $\tilde{\Theta}(M)$ connecting $\tilde{f}(\tilde{z})$ to $\tilde{f}(\tilde{w})$ and not meeting $\bigcup \tilde{\mathcal{S}}'[2]$. Thus $\tilde{f}(\tilde{S}_0)$ and $\tilde{f}(\tilde{S}_1)$ are not separated by any element of $\tilde{\mathcal{S}}'[2]$. Since the corresponding surfaces \tilde{S}'_0 and \tilde{S}'_1 lie in arbitrarily small neighbourhoods of $\tilde{f}(\tilde{S}_0)$ and $\tilde{f}(\tilde{S}_1)$, these are not separated by any element of $\tilde{\mathcal{S}}'[2]$. Since we have an isomorphism of Bass-Serre trees, the corresponding

surfaces in $\Theta(P)$ have the same separation properties, and so \tilde{S}_0 and \tilde{S}_1 are not separated by any element of $\mathcal{S}[2]$. We can thus connect \tilde{z} to \tilde{w} by a path in $\tilde{\Theta}(P)$ not meeting $\bigcup \mathcal{S}[2]$. Together with the paths $\tilde{\beta}$ and $\tilde{\gamma}$, this gives a path from \tilde{x} to \tilde{y} . Projecting back down to $\Theta(P)$, this gives a path from x to y in $R[2]$ in the right homotopy class, as claimed.

In fact, we can refine the above observation slightly. Note that every time $\tilde{\alpha}$ crosses some component of $\mathcal{S}[2]$ it must eventually cross back again, and so we can replace the intervening path by a path in an arbitrarily small neighbourhood of this component. Projecting to $\Theta(P)$, we see that we can find a new path α in the right homotopy class in $R[2]$ and in an arbitrarily small neighbourhood of our original α union $\partial_H R[2]$. We refer to this operation as “pushing α into $R[2]$ ”.

Our next job is to push α off every level 7 band. Suppose that $B[7] \in \mathcal{B}[7]$. By our initial assumption, $x, y \notin B[6]$. We can now apply the above argument, with $CB[6]$ playing the role of $R[2]$ and $CB[7]$ playing the role of $R[3]$ to push α off $B[7]$. In other words, we replace α by another path in the right homotopy class in $CB[7]$, and in a small neighbourhood of our previous α union $\partial_H B[7]$. Our new path might now leave $R[2]$, however, since the pushing operations took place inside $B[6]$ and so certainly inside $B[0]$, Lemma 10.6(2) ensures that the resulting path lies inside $R[3]$.

We want to perform this construction for all level 7 bands, however there is a risk that the various “pushing” operations may interfere with each other. We therefore proceed by (reverse) induction on the complexity of the bands. By Lemma 10.6(1), any two level 0 bands of the same complexity are disjoint, and therefore the pushing operations on such bands can be performed simultaneously (or more precisely, in any order). We thus start with the level 7 bands of complexity $\kappa(\Sigma) - 1$, and then move onto those of complexity $\kappa(\Sigma) - 2$ and continue all the way down to bands of complexity 1 (observing that there are no 3HS bands). The pushing operations of a given complexity may affect those already performed at a higher complexity, but Lemma 10.6 parts (1) and (4) ensure that we will never enter a level 8 band. Again, Lemma 10.6(2) ensures we remain inside $R[3]$. We thus end up with a path $\alpha \subseteq R[3] \setminus \bigcup \mathcal{B}[8]$ in the right homotopy class.

Now by Lemmas 10.3 and 10.4, $f(\alpha) \subseteq R'[4] \setminus \bigcup \mathcal{B}'[9]$. Since α lies in the right homotopy class, $f(\alpha) \cup [f(x), f(y)] \subseteq R'[4] \setminus \bigcup \mathcal{B}'[9]$ bounds (the continuous image of) a disc D in $\Theta(M)$. Now the boundaries, $\partial_{\pm} R'[4]$ are incompressible in $\Theta(M)$, and so we can push D into $R'[4]$, so that the resulting disc lies in a small neighbourhood of our original disc union $\partial_H R'[4]$.

Next, we push D off all level 10 bands, by reverse induction on complexity as before. For this we only need to observe that the boundaries of bands are incompressible. By Lemma 10.7, we end up with a disc D lying in $R'[5] \setminus \bigcup \mathcal{B}'[10]$.

In summary, we have found $\alpha \subseteq R[3] \setminus \bigcup \mathcal{B}[8]$ such that $f(\alpha) \cup [f(x), f(y)]$ bounds a disc in $R'[5] \setminus \bigcup \mathcal{B}'[10]$. Using Lemma 10.8, we see that the diameter of α in $(\Theta(P), d)$ is bounded. Using Lemma 10.9, and the subsequent remark about the η -neighbourhood, we see that the diameter of the disc is bounded in $(\Theta(M), d')$. This proves Proposition 10.10 in this case.

We now move on the case where x or y lies in some level 6 band. Among all bands in $\mathcal{B}[6]$ that meet $\{x, y\}$ we choose one, say $B[6]$, of minimal complexity. We can assume that $x \in B[6]$. Let $\Phi = \pi_{\Sigma} B[6]$ be the base surface. By the minimal complexity assumption,

we see that $x, y \notin \bigcup \mathcal{B}_\Phi[6]$. Let $R_\Phi[n] = R_{\Phi,x}[n]$. Since $x \in B[6]$, we get that $R[5]$ exists and lies inside $B[0]$.

We can now carry out the above construction, with $R_\Phi[n]$ replacing $R[n]$, and with $\mathcal{B}_\Phi[n]$ replacing $\mathcal{B}[n]$. In this way, we get a path $\alpha \subseteq R_\Phi[3] \setminus \bigcup \mathcal{B}_\Phi[8]$ such that $f(\alpha) \cup [f(x), f(y)]$ bounds a disc in $R'_\Phi[5] \setminus \bigcup \mathcal{B}'_\Phi[10]$. By Lemmas 10.8 and 10.9 again, we see that these have bounded diameter in $(\Theta(P), d)$ and $(\Theta(M), d')$ respectively.

This proves Proposition 10.10.

Finally, putting Proposition 10.10 together with Lemma 8.4, we get:

Proposition 10.11 : *The map $f : (\Theta(P), d) \longrightarrow (\Theta(M), d')$ (constructed as in Section 7) is uniformly universally sesquilipschitz. \diamond*

11. The doubly degenerate case.

In this section, we gather our constructions together to show that two doubly degenerate hyperbolic 3-manifolds with the same pairs of end invariants are isometric (Theorem 11.12).

First we need to clarify what we mean by “end invariant”. This discussion makes sense for any complete hyperbolic 3-manifold, M . Let $\Psi(M)$ be the non-cuspidal part of M , i.e. we fix some Margulis constant, and remove the interiors of Margulis cusps. (See Section 12, for more details).

We say that an end, e , of $\Psi(M)$ is *tame* if there is a compact surface, Σ , such that a neighbourhood of the end is homeomorphic to $\Sigma \times [0, \infty)$. In this case, we write $\partial_H e = \Sigma \times \{0\}$ and $\partial_V e = \partial\Sigma \times [0, \infty)$ for the *horizontal* and *vertical* boundaries respectively. The horizontal boundary is thus the relative boundary in $\Psi(M)$. We say that e is *incompressible* if its inclusion in M is π_1 -injective.

Let e be an incompressible tame end of $\Psi(M)$. By [Bon], such an end is either geometrically finite or (simply) degenerate. For our purposes, we can define an end to be *degenerate* if it contains an infinite sequence of distinct curves of bounded length, each homotopic to a simple closed curve in $\partial_H e$. Note that such a sequence must necessarily tend out the end. We fix a homotopy class of identification of Σ with $\partial_H e$, and write $\mathcal{G} = \mathcal{G}(\Sigma)$. We have the following:

Proposition 11.1 : *Suppose that e is a degenerate end. Then there is a unique $a(e) \in \partial\mathcal{G}$ with the following property. Suppose that $(\gamma_i)_i$ is a sequence of curves in $X(\Sigma)$ which have representatives of bounded length in e , then $\gamma_i \rightarrow a(e)$ in $\mathcal{G} \cup \partial\mathcal{G}$. \diamond*

We should make some remarks about this definition. The existence of such a sequence of curves is given by our definition of degenerate end. In fact (as we see below) one could choose the length bound to depend only on $\kappa(\Sigma)$, though the statement holds true for any bound. We can also take our representative in e to be closed geodesics. This is a now standard argument (cf. [Bon]). Suppose we start with some set of bounded length representatives, γ'_i in e . Let $\bar{\gamma}_i$ be the closed geodesic representative in M . (Only finitely many might be parabolic so we can exclude these if we want.) There is a bounded area

homotopy between γ'_i and $\bar{\gamma}_i$ in M . If the $\bar{\gamma}_i$ did not go out the end, then we would get a subsequence of curves in these homotopy classes all of bounded length in a compact subset of e , which is impossible.

To relate this the more usual notion of end invariant in terms of laminations, we use the result of Klarreich [Kl] (see also [Ham1]) which identifies $\partial\mathcal{G}$ with the set of “arational laminations” on Σ :

Proposition 11.2 : *There is a bijection, a , from the set of arational laminations on Σ to the boundary $\partial\mathcal{G}(\Sigma)$ such that sequence of curves $(\gamma_i)_i$ in $X(\Sigma)$ converges on an arational lamination λ , if and only if, $\gamma_i \rightarrow a(\lambda)$ in $\mathcal{G} \cup \partial\mathcal{G}$. \diamond*

We shall not define the terms “arational lamination” or “converges on an arational lamination” here, since only the logic of the argument is relevant.

The work of Bonahon proves an analogous statement to Lemma 11.1, where “representatives of bounded length” is replaced by “closed geodesic representatives tending out the end”, and “ $\gamma_i \rightarrow a(e)$ ” is replaced by “ γ_i converges on a lamination $\lambda(e)$ ” (such a lamination is necessarily arational). We can arrange that some sequence curves have bounded length, since any closed geodesic that is simple in Σ can be extended to a “pleated surface”, which, for us, is a uniformly lipschitz homotopy equivalence of a finite are hyperbolic surface into the end. Such a surface must contain a simple closed curve of (uniformly) bounded length, and we take its image in the pleated surface. By our earlier argument, we can then take these curves to be geodesics in e . Bonahon’s statement implies that these also converge on the same arational lamination $\lambda(e)$. Indeed any set of curves of bounded length will converge on $\lambda(e)$. We can now set $a(e) = a(\lambda(e))$ and Proposition 11.1 follows via Proposition 11.2.

We now return to our doubly degenerate manifold, M . This means that $\Psi(M)$ is homeomorphic to $\Sigma \times \mathbf{R}$, and that both ends are degenerate. We write e_- and e_+ for the positive and negative ends of $\Psi(M)$ respectively. This gives us two end invariants, $a(e_-), a(e_+) \in \partial\mathcal{G}$.

Lemma 11.3 : $a(e_-) \neq a(e_+)$.

Proof : This follows from the fact [Bon] that $\lambda(e_-) \neq \lambda(e_+)$. \diamond

One could also deduce Proposition 11.3 by a slightly convoluted argument using the “a-priori bounds” estimate (see Proposition 11.6). If the end invariants were equal, we could take sequence of curves of bounded length in each end and connect them by tight geodesics in \mathcal{G} . This would then give rise a sequence of curves also tending to the end invariant, while remaining in a compact set in $\Psi(M)$. We shall not give the details of this argument here.

We are now in a position to construct our model space for M .

To begin, Theorem 4.1 associates to the pair $a(e_-), a(e_+)$ a complete annulus system. $W = \bigcup \mathcal{W} \subseteq \Psi = \Sigma \times \mathbf{R}$. (There might, of course, be many annulus systems satisfying the conditions of Theorem 4.1. We can arbitrarily choose one of them.) The construction of Section 7, now gives us a riemannian manifold $(\Psi(P), d)$ — first open up each annulus

into a torus, and then glue in a Margulis tube. Thus, $\Psi(P)$ is also homomorphic to $\Sigma \times \mathbf{R}$, and there is a natural (topological) proper homotopy equivalence of $\Psi(P)$ with $\Psi(M)$. Each component of $\partial\Psi(P)$ is a bi-infinite cylinder isometric to $S^1 \times \mathbf{R}$, in the induced path metric. We now construct P by gluing in a standard \mathbf{Z} -cusp to each such boundary component (a quotient of a horoball in \mathbf{H}^3 by a \mathbf{Z} -action). Thus P is a complete riemannian manifold with empty boundary.

To relate the geometry of P to the geometry of M , the following ‘‘a-priori bounds’’ estimate is key. Given $\alpha \in \Sigma$, we write $\bar{\alpha}$ for the closed geodesic in M in this homotopy class, and write $l_M(\alpha)$ for its length.

Theorem 11.4 : *Suppose that $\Phi \subseteq \Sigma$ is a subsurface. Suppose that $\alpha, \beta, \gamma \in X(\Phi)$ and that γ lies in a tight geodesic from α to β in $\mathcal{G}(\Phi)$ (or in $\mathcal{G}'(\Phi)$ if $\kappa(\Phi) = 1$). Then $l_M(\gamma)$ is bounded above in terms of $\kappa(\Sigma)$, $l_M(\alpha)$, $l_M(\beta)$ and $\max\{l_M(\delta)\}$ as δ ranges over the boundary components $\partial^\Sigma\Phi = \partial\Phi \setminus \partial\Sigma$. \diamond*

Here we are allowing the possibility that $\Phi = \Sigma$. We do not require that M is doubly degenerate for this result, only that $\Psi(M)$ is homeomorphic to $\Sigma \times \mathbf{R}$, possibly with a number of cusps corresponding to accidental parabolics removed. If α is homotopic to such a parabolic, we set $l_M(\alpha) = 0$. We need not worry about this in this section, though it cannot be avoided in general — see Section 12.

This (or a very similar) a-priori bounds estimate was proven by Minsky [Mi4]. Another proof is given in [Bow2].

By induction, we see that if $\alpha, \beta, \gamma \in X(\Sigma)$ and $\gamma \in Y^\infty(\{\alpha, \beta\})$ as defined in Section 2, then $l_M(\gamma)$ is bounded above in terms of $\kappa(\Sigma)$, $l_M(\alpha)$ and $l_M(\beta)$.

Now in Section 4, we constructed the annulus system \mathcal{W} out of a sequence of sets of the form $Y(\alpha_i, \beta_i)$ as $\alpha_i \rightarrow a$ and $\beta_i \rightarrow b$. By the defining property of end invariants, we can choose α_i and β_i so that $l_M(\alpha_i)$ and $l_M(\beta_i)$ remain bounded, and so by Theorem 11.4, all the curves of we construct will have bounded length in M . However, the choice of these sequences might depend on M . To see that all the curves have bounded length for any choice of sequences, we use the following variation on Theorem 11.4 also proven in [Bow2].

Theorem 11.5 : *Given any $r \geq 0$ there is some $r' \geq 0$ such that if $\alpha, \beta, \gamma \in X(\Sigma)$ and γ lies on a tight geodesic from α to β and $d(\alpha, \gamma) \geq r'$ and $d(\beta, \gamma) \geq r'$, then $l_M(\gamma)$ is bounded above in terms of $\kappa(\Sigma)$, $\min\{l_M(\delta) \mid \delta \in N(\alpha, r)\}$ and $\min\{l_M(\epsilon) \mid \epsilon \in N(\gamma, r)\}$. \diamond*

Putting these together with Proposition 11.1 and the fact that $\mathcal{G}(\Sigma)$ is uniformly hyperbolic, we obtain:

Proposition 11.6 : *There is some constant $L \geq 0$ such that $\Omega \in \mathcal{W}$, $l_M(\bar{\Omega}) \leq L$, where Ω is the closed geodesic in M in the homotopy class of Ω . \diamond*

This is precisely the hypothesis (APB) of Section 7 that allowed us to construct the map $f : \Psi(P) \rightarrow \Psi(M)$. In particular, Proposition 7.1 gives us a partition, $\mathcal{W} = \mathcal{W}_0 \cup \mathcal{W}_1$ such that $f : \Theta(P) \rightarrow \Theta(M)$ is a proper lipschitz homotopy equivalence of the thick parts. (We remark that the partition of \mathcal{W} and hence the definition of $\Theta(P)$ might depend on

M , but that does not affect the logic of the argument.)

Lemma 11.7 : *The map f sends the positive (negative) end of $\Psi(P)$ to the positive (negative) end of $\Psi(M)$.*

Proof : Let Ω_i be a sequence of annuli tending out the positive end of $\Psi(P)$. Now f sends $\partial T(\Omega_i)$ either to the associated geodesic $\bar{\Omega}_i$ in $\Psi(M)$, or else to the boundary of the Margulis tube about $\bar{\Omega}_i$. In any case, since f is proper, the sequence $\bar{\Omega}_i$ must go out an end, e , of $\Psi(M)$. By construction of \mathcal{W} , the homotopy classes of Ω_i tend to $a(e_+)$ in $\mathcal{G} \cup \partial \mathcal{G}$ and so by the definition of end invariant (Proposition 11.1), we see that $a(e) = a(e_+)$, and so, by Lemma 11.3, $e = e_+$, as required. \diamond

This proves the end consistency assumption (EC) of Section 7, and so $f(\Psi(P)) = \Psi(M)$.

This was all that was needed to get us to Proposition 10.11, and so we see that the map $f : \Theta(P) \rightarrow \Theta(M)$ is uniformly universally sesquilipschitz.

For the moment, f is only defined topologically on each of the Margulis tubes in $\mathcal{T}(P)$. If $T \in \mathcal{T}(P)$, then we have a lipschitz map $f : \partial T \rightarrow \partial T'$.

Lemma 11.8 : *If $T \in \mathcal{T}(P)$, then f extends to a uniformly universally sesquilipschitz map, $f : T \rightarrow T'$.*

In other words, the extension, f is uniformly lipschitz and its lift to the universal covers, $\tilde{f} : \tilde{T} \rightarrow \tilde{T}'$ is a quasi-isometry.

Proof : Let $\tilde{\Theta}(P)$ and $\tilde{\Theta}(M)$ be the universal covers of $\Theta(P)$ and $\Theta(M)$, and let $\hat{\Theta}(P) = \tilde{\Theta}(P)/H$ and $\hat{\Theta}(M) = \tilde{\Theta}(M)/H$ be the covers corresponding to the subgroup H of $G = \pi_1(\Theta(P)) \cong \pi_1(\Theta(M))$ generated by the longitude of T . We can identify $\partial \tilde{T}$ and $\partial \tilde{T}'$ with boundary components of $\hat{\Theta}(P)$ and $\hat{\Theta}(M)$ respectively. In the induced path metrics, they are euclidean cylinders whose longitudes have length uniformly bounded above and below.

By Proposition 10.11, the map $\tilde{f} : \tilde{\Theta}(P) \rightarrow \tilde{\Theta}(M)$ is a lipschitz quasi-isometry, and so therefore is its projection, $\hat{f} : \hat{\Theta}(P) \rightarrow \hat{\Theta}(M)$. By Lemma 4.8, $\partial \tilde{T}$ is quasi-isometrically embedded in $\hat{\Theta}(P)$, and so we can conclude that $\hat{f}|_{\partial \tilde{T}}$ is a quasi-isometry from $\partial \tilde{T}$ to $\partial \tilde{T}'$ in the induced euclidean path metrics.

We are therefore in the situation described by Lemma 5.4 and the subsequent remark. In particular, there is a universally sesquilipschitz homotopy from $f|_{\partial T} : \partial T \rightarrow \partial T'$ to a bilipschitz homeomorphism $g : \partial T \rightarrow \partial T'$. By Lemma 5.8, such a map g extends to a bilipschitz homeomorphism $g : T \rightarrow T'$.

Now we can carry out the sesquilipschitz homotopy between $f|_{\partial T}$ and $g|_{\partial T}$ in a uniformly small neighbourhood of ∂T in T , and then use g to extend over T . This way, we extend f to a universally sesquilipschitz map $f : T \rightarrow T'$.

Performing this for each tube $T \in \mathcal{T}(P)$ we get a lipschitz map $f : \Psi(P) \rightarrow \Psi(M)$. One can show this to be universally sesquilipschitz (cf. Proposition 11.10), but we are really interested in a further extension of f to the whole model space P . For this we still need to deal with the cusps.

Let R be a cusp of P , i.e. the closure of a component of $P \setminus \Psi(P)$. We have a corresponding cusp in R' in M , the closure of a component of $M \setminus \Psi(M)$. We have a proper lipschitz map $f|_{\partial R} : \partial R \rightarrow \partial R'$, between bi-infinite euclidean cylinders.

Lemma 11.9 : *The map then $f|_{\partial R}$ extends to a uniformly universally sesquilipschitz map, $f : R \rightarrow R'$.*

Proof : The argument is similar to that for Lemma 11.8. Using Lemma 4.8 as before, we see that $f|_{\partial R}$ is a uniform quasi-isometry to $\partial R'$. (In this case, we lift R rather than the universal cover of R .) Note that ∂R and $\partial R'$ are both uniformly quasi-isometric to the real line, under horizontal projection. We can apply Lemma 5.1 directly to see that there is a bounded homotopy to a bilipschitz homeomorphism of the real line. Thus (as in Lemma 5.4) we get a universally sesquilipschitz homotopy from $f|_{\partial R}$ to a bilipschitz homeomorphism $g : \partial R \rightarrow \partial R'$. (We use that fact that a lipschitz map from \mathbf{R}^2 to \mathbf{R}^2 which sends horizontal lines to horizontal lines and is bilipschitz in the vertical direction must, in fact, be bilipschitz.) The extension of g over R is now trivial — just send rays isometrically to rays. \diamond

Performing this for each cusp we get a proper lipschitz homotopy equivalence, $f : P \rightarrow M$.

Proposition 11.10 : *The map $f : P \rightarrow M$ is uniformly universally sesquilipschitz.*

Proof : In other words, we claim that the lift $\tilde{f} : \tilde{P} \rightarrow \tilde{M}$ is a uniform quasi-isometry. Since \tilde{f} is surjective, it is enough to put an upper bound on $d(x, y)$ whenever $d'(\tilde{f}(x), \tilde{f}(y)) < \eta$ for some fixed $\eta > 0$. But \tilde{P} and \tilde{M} are equivariantly decomposed into pieces, namely the lifts of thick parts, Margulis tubes, and cusps. We have shown that \tilde{f} respects this decomposition and that \tilde{f} restricted to each of the pieces is a uniform quasi-isometry. Moreover, we can assume that any two distinct pieces are distance at least η apart in \tilde{M} . The result now follows. \diamond

We can summarise what we have shown as follows:

Theorem 11.11 : *Given two distinct $a, b \in \partial \mathcal{G}$, we can construct a complete riemannian manifold, P , homeomorphic to $\text{int}(\Sigma) \times \mathbf{R}$, such that if M is a doubly degenerate hyperbolic 3-manifold with base surface Σ and end invariants a, b , then there is a uniformly universally sesquilipschitz map from P to M .* \diamond

Here “uniform” means that the constants depend only on $\kappa(\Sigma)$. (I don’t know if this dependence is computable.)

As a consequence we have:

Theorem 11.12 : *Suppose M, M' are doubly degenerate hyperbolic 3-manifolds with the same base surface and end invariants. Then M and M' are isometric.*

Proof : The argument is now standard. We can use the same model space P for both M and M' . The universally sesquiplipschitz maps, $P \rightarrow M$ and $P \rightarrow M'$ give us an equivariant quasi-isometry between \tilde{M} and \tilde{M}' , both isometric to \mathbf{H}^3 . This extends to an equivariant quasiconformal map $\partial\mathbf{H}^3 \rightarrow \partial\mathbf{H}^3$. The result of Sullivan [Su] now tells us that this is fact conformal. Thus the two actions on \mathbf{H}^3 are conjugate by an isometry of \mathbf{H}^3 , and so M is isometric to M' . \diamond

This is, of course, the “ending lamination conjecture” for the case of doubly degenerate manifolds, as originally proven in [Mi4, BrocCM1].

12. The indecomposable case.

In this section we describe how earlier arguments can be adapted to construct a model space for an indecomposable (orientable) complete hyperbolic 3-manifold. With some further modifications, the decomposable case can also be dealt with in this way. We begin with a topological discussion.

Definition : A 3-manifold Ψ with boundary, $\partial\Psi$, is *topologically finite* if we can embed Ψ in a compact 3-manifold, $\bar{\Psi}$, with boundary, $\partial\bar{\Psi}$, so that $\partial\Psi$ is a subsurface of $\partial\bar{\Psi}$, and $\bar{\Psi} = \Psi \cup \partial\bar{\Psi}$.

In other words, we can compactify Ψ by adjoining the “ideal” boundary $\partial_I\Psi = \partial\bar{\Psi} \setminus \partial\Psi$. In fact, the topology of $(\bar{\Psi}, \Psi)$ is determined by Ψ , though here we can regard $\bar{\Psi}$ as part of the structure associated to Ψ .

Suppose that Ψ is topologically finite.

Definition : By a *core* of Ψ we will mean a compact submanifold, $\Psi_0 \subseteq \Psi$, such that $\Psi \setminus \Psi_0$ is homeomorphic to $\partial_I\Psi \times \mathbf{R}$.

We write $\partial_H\Psi_0$ for the relative, or *horizontal* boundary of Ψ_0 in Ψ , and $\partial_V\Psi_0 = \Psi_0 \cap \partial\Psi$ for the *vertical* boundary. Thus $\partial\Psi_0 = \partial_H\Psi_0 \cup \partial_V\Psi_0$. The ends of Ψ are in bijective correspondence with the components of $\partial_H\Psi_0$, and are just products.

Definition : An end $e \cong \Sigma \times [0, \infty)$ is *incompressible* in Ψ if its inclusion into Ψ is π_1 -injective.

Definition : We say that Ψ is *indecomposable* if all its ends are incompressible.

Via Dehn’s lemma, this is equivalent to saying that there is no disc in $\bar{\Psi}$ whose boundary lies in $\partial_I\Psi$.

Now suppose that M is a complete orientable hyperbolic 3-manifold, and that $\pi_1(M)$ is finitely generated. We shall also assume that M is not elementary, i.e. that $\pi_1(M)$ is not abelian. Let $\Psi(M)$ be the non-cuspidal part of M . The tameness theorem [Ag, CalG] says that $\Psi(M)$ is topologically finite. Here we will be assuming that $\Psi(M)$ is indecomposable.

Under the equivalent indecomposability assumption on $\pi_1(M)$, tameness was proven by Bonahon [Bon].

Let Ψ_0 be a core for $\Psi(M)$. We note that each component of $\partial\Psi(M)$ is either a bi-infinite euclidean cylinder, meeting Ψ_0 in a compact annular component of $\partial_V\Psi_0$, or else is a torus and a component of $\partial_V\Psi_0$. These components bound \mathbf{Z} -cusps and $\mathbf{Z} \oplus \mathbf{Z}$ -cusps respectively in M . We also note that no component of $\partial_H\Psi_0$ is a disc or annulus.

We write $\mathcal{E}(M)$ for the set of ends of $\Psi(M)$. We write $\mathcal{E}_D(M) \subseteq \mathcal{E}(M)$ for the set of degenerate ends (as defined in Section 11), and write $\mathcal{E}_G(M) = \mathcal{E}(M) \setminus \mathcal{E}_D(M)$. The geometrical interpretation of tameness says that each end of $\mathcal{E}_G(M)$ is geometrically finite. The geometry of geometrically finite ends have long been well-understood. Note that any end whose base surface is a 3HS must be geometrically finite.

Let $C = C(M)$ be the convex core of M , and let $C(r)$ be its r -neighbourhood. Thus if $r > 0$, $\partial C(r)$ is a C^1 -submanifold of M . If we choose r sufficiently large, then $\Psi(M) \setminus \text{int } C(r)$ will consist of a disjoint union of geometrically finite ends. We can thus fix some $r > 0$ and assume that each $e \in \mathcal{E}_G(M)$ has this form. Let $\partial_H e$ be the relative boundary of e in $\Psi(M)$, let S be the component of $\partial C(r)$ containing $\partial_H e$, and let E be the component of $M \setminus \text{int } C(r)$ with boundary $\partial E = S$. Thus $e = E \cap \Psi(M)$ and $\partial_H e = S \cap \Psi(M)$. We write $\Psi(E) = E \cap \Psi(M) = e$, $\partial_H \Psi(E) = \partial_H e$ and $\partial_V \Psi(E) = \Psi(E) \cap \partial \Psi(M)$. Each component of $\partial_V \Psi(E)$ is a euclidean half-cylinder and an end of a component of $\partial \Psi(M)$.

By Ahlfors's finiteness theorem, each geometrically finite end has associated to it a Riemann surface of finite type, which can be thought of as geometrically finite end invariant. For the moment, however we will not be using this structure.

We now describe the main aim of this section. Suppose that Ψ is a topologically finite 3-manifold, with a decomposition of the of ends, $\mathcal{E} = \mathcal{E}_F \sqcup \mathcal{E}_D$, so that no base surface of an end is a disc, annulus, sphere or torus, and no base surface in \mathcal{E}_F is a 3HS. Suppose that to each $e \in \mathcal{E}_D$ there is associated an element, $a(e)$, in the boundary of the corresponding curve graph. We will associate to Ψ , $(a(e))_{e \in \mathcal{E}_D}$, a "model" manifold P . This will be a complete reimannian manifold together with a submanifold, $\Psi(P)$, homeomorphic to Ψ , such that each component of $P \setminus \Psi(P)$ is either a "standard" \mathbf{Z} -cusp or "standard" $\mathbf{Z} \oplus \mathbf{Z}$ -cusp. (We could also include into P information about geometrically finite end invariants, but we will not worry about that for the moment.)

We will show:

Theorem 12.1 : *Let M be a tame indecomposable hyperbolic 3-manifold with non-cuspidal part $\Psi(M)$. Let P be the model manifold referred to above, constructed from $\Psi(M)$, the partition of its ends, $\mathcal{E}(M) = \mathcal{E}_G(M) \sqcup \mathcal{E}_D(M)$ into geometrically finite and degenerate, and the collection $(a(e))_{e \in \mathcal{E}_D(M)}$ of degenerate end invariants. Then there is a universally sesquipschitz map from P into M .*

Note that here we are no longer claiming that the constants of our sesquipschitz map are uniform. They might depend on the geometry as well as the topology of M . We suspect that some uniform statement could be made, but one would need to take into account the geometrically finite end invariants when constructing the model space. In any case, this would considerably complicate the construction.

The doubly degenerate case of Section 11, is where $\Psi(M) \cong \Sigma \times \mathbf{R}$ and $\mathcal{E}_G(M) = \emptyset$. In this case, we can take the constants to depend only on $\kappa(\Sigma)$.

The basic idea is to construct a model for each end of $\Psi(M)$, and then put any riemannian metric on the core Ψ_0 . The only requirement of the latter is that it should match up with the metric we have already on the boundaries of the model ends.

Let Σ be a compact surface. We want to associate to Σ a geometrically finite model P_Σ . Here we just give a very crude model that only depends on the topological type of Σ . A more sophisticated model, which takes into account an end invariant (Riemann surface) is described in [Mi4].

Let us fix any finite-area hyperbolic structure on $\text{int } \Sigma$. This is given by the quotient, \mathbf{H}^2/H , of a properly discontinuous action of $H = \pi_1(\Sigma)$ on \mathbf{H}^2 . We embed \mathbf{H}^2 as a totally geodesic subspace of \mathbf{H}^3 and extend the action to \mathbf{H}^3 . Let P_Σ be the quotient of one of the half-spaces bounded by \mathbf{H}^2 . We write $\Psi(P_\Sigma)$ for the non-cuspidal part of P_Σ . Thus each component of $\partial\Psi(P_\Sigma)$ is a euclidean half-cylinder, which (at least for notational convenience) we can assume to be isometric to $S^1 \times [0, \infty)$.

Note that P_Σ has a product structure as $\partial P_\Sigma \times [0, \infty)$, where the first co-ordinate of $x \in P_\Sigma$ is the nearest point to x in ∂P_Σ , and the second co-ordinate, $t = t(x)$, is the distance of x from ∂P_Σ . Thus P_Σ is a warped riemannian product where the distances in the horizontal (constant t) direction are expanded by a factor of $\cosh t$.

A geometrically finite end of M has qualitatively similar geometry. This is well understood. We only give an outline here.

Let $e \in \mathcal{E}_G(M)$, and let $E, \Psi(E)$, etc. be as defined above. Note that ∂E is component of $\partial C(r)$. Now E has a product structure $E \cong \partial E \times [0, \infty)$ defined exactly as with P_Σ . In this case, the horizontal expansion at distance t from ∂E need not be constant, but will be bounded between two constants, namely, $k_-(t) = \cosh(t+r)/\cosh(r)$ and $k_+(t) = \sinh(t+r)/\sinh(r)$. We note that the ratios of both $k_-(t)$ and $k_+(t)$ with $\cosh t$ are bounded above and below (in terms of r).

Now ∂E meets each \mathbf{Z} -cusp in a constant curvature cusp and so we can find a bilipschitz homeomorphism $g : \partial P_\Sigma \rightarrow \partial E$. We can now extend, using the product structures, to a homeomorphism, $g : P_\Sigma \rightarrow E$, which, by the above observations will also be bilipschitz. Unfortunately, this need not send $\Psi(P_\Sigma)$ to $\Psi(E)$, though it is not hard to modify it so that it does. One way to describe this is as follows.

We fix some positive $k < 1$ as described below, and choose g so that for each $s \leq 1$ it sends any horocycle of length s in ∂P_Σ to a horocycle of length ks in ∂E . Now given any $t \geq 0$, the level t surfaces in P_Σ and E , meet the \mathbf{Z} -cusps in cusps of constant curvature determined by t and r . Under the above construction, g will send a horocycle of length $s \leq 1$ in such a surface in P_Σ to a horocycle in the corresponding surface in E , whose length is bounded above and below by fixed multiples of ks . By choosing k sufficiently small, we can assume that this length is always less than the Margulis constant. Thus, g sends each \mathbf{Z} -cusp in P_Σ to the corresponding \mathbf{Z} -cusp in E . We can now modify g by post-composing with projection of such a cusp in E to its boundary, using nearest point projection in the level surfaces in E . This projection will have bounded expansion on $g(\Psi(P_\Sigma))$. The resulting map $f : \Psi(P_\Sigma) \rightarrow \Psi(E)$ is universally sesquilipschitz.

We have shown:

Lemma 12.2 : *If $e \in \Psi(E)$ is a geometrically finite end of $\Psi(M)$ with base surface Σ , then there is a universally sesquilateral map $f : \Psi(P_\Sigma) \longrightarrow \Psi(E)$.* \diamond

We next want to construct models for simply degenerate ends. We can use the following variation on Theorem 4.1. The proof is essentially the same, indeed a more direct application of Lemma 4.3. Given a compact surface Σ , write $\Psi_+ = \Sigma \times [0, \infty)$ and $\partial_H \Psi_+ = \Sigma \times \{0\}$.

Lemma 12.3 : *Given a complete multicurve, α , and some $a \in \partial\mathcal{G}(\Sigma)$ we can find a complete annulus system $W = \bigcup \mathcal{W} \subseteq \Psi$ with $\pi_\Sigma(W \cap \partial_H \Psi_+) = \alpha$, and satisfying the conditions (1)–(4) of Theorem 4.1.* \diamond

We need to interpret condition (1) which said that $X(\mathcal{W}) \subseteq \bar{Y}^\infty(Y)$. Here we can take Y to be the limit of the sets $Y^\infty(X(\alpha) \cup \{\beta_i\})$ where $\beta_i \in X(\Sigma)$ is some sequence converging to a .

Now let $\mathcal{W}_I = \{\Omega \in \mathcal{W} \mid \Omega \cap \partial_H \Psi_+ = \emptyset\}$ and $\mathcal{W}_\partial = \mathcal{W} \setminus \mathcal{W}_I$. Let $\Lambda(\mathcal{W})$ be the space obtained by opening out each annulus of \mathcal{W} as before. We have a natural map, $p : \Lambda(\mathcal{W}) \longrightarrow \Psi_+$. Each $\Omega \in \mathcal{W}_I$ gives us a solid torus, $\Delta(\Omega)$, and each $\Omega \in \mathcal{W}_\partial$ gives us an annulus, $A(\Omega)$ with boundary $p^{-1}(\Omega \cap \partial_H \Psi_+)$.

Now let $\Psi(P_a) = \Lambda(\mathcal{W}, \mathcal{W}_I)$ be the space obtained by gluing in a solid torus, $T(\Omega)$, to each $\Delta(\Omega)$ for $\Omega \in \mathcal{W}_I$. (We won't need to define a space P_a , but will write $\Psi(P_a)$ for the sake of maintaining consistent notation.) We write $\partial_H \Psi(P_a) = p^{-1}(\partial_H \Psi_+) \cup \bigcup_{\Omega \in \mathcal{W}_\partial} A(\Omega)$. In other words, it consists of all the (3HS) components of $\partial_H \Psi_+ \setminus \alpha$ connected by annuli $A(\Omega)$, so as to recover Σ up to homeomorphism. In fact, $(\Psi(P_a), \partial_H \Psi(P_a)) \cong (\Sigma \times [0, \infty), \Sigma \times \{0\})$.

We can now put a riemannian metric, d , on $\Psi(P_a)$, exactly as we did with $\Psi(P)$, by giving each $T(\Omega)$ the structure of a Margulis tube. It also has a pseudometric, ρ , obtained by deeming each $T(\Omega)$ to have diameter 0. Near the boundary, $\partial_H \Psi(P_a)$, these metrics may be a bit of a mess, but in a neighbourhood of the end of $\Psi(P_a)$ they will have all of the properties, (W1)–(W9) laid out in Section 4.

We are now in a position to describe the model space, P . The only data we need is the topology of $\Psi(M)$, the partition of its ends as $\mathcal{E}(M) = \mathcal{E}_G(M) \sqcup \mathcal{E}_D(M)$, and the assignment of degenerate end invariants, $(a(e))_{e \in \mathcal{E}_D(M)}$.

Let $\Psi_0(P)$ be a homeomorphic copy of the core, $\Psi_0(M)$, of $\Psi(M)$. We have a decomposition of its boundary into the horizontal and vertical parts, $\partial \Psi_0(P) = \partial_H \Psi_0(P) \cup \partial_V \Psi_0(P)$. For each $e \in \mathcal{E}_G(M)$, we take a copy $\Psi(P_e) = \Psi(P_{\Sigma(e)})$ of the geometrically finite model, for the base surface $\Sigma(e)$, and glue $\partial_H \Psi(P_e)$ to the corresponding component of $\partial_H \Psi_0(P)$. If $e \in \mathcal{E}_D(M)$, we take a copy of the degenerate model, $\Psi(P_e) = \Psi(P_{a(e)})$ and again glue $\partial_H \Psi(P_e)$ to the corresponding component of $\partial_H \Psi_0(P)$. This case involves making a choice of multicurve, $\alpha \subseteq \Sigma(e)$, to construct $\Psi(P_{a(e)})$. In principle we could take any multicurve, but to avoid some technical complications, we could take it so that no component of α is homotopic in $\Psi_0(P)$ into the vertical boundary, $\partial_V \Psi_0(P)$ (i.e. so that no curve of α ends up being an accidental parabolic). In this way, we have constructed a topological copy of, $\Psi(P)$, of $\Psi(M)$. We have already some riemannian metric on $\partial_H \Psi_0(P)$. The model ends were such that the boundary curves of $\partial_H \Psi_0(P)$ all have unit length. Each component of $\partial_V \Psi_0(P)$ is either an annulus bounded by two such curves,

which we can take to be isometric to $S(1) \times [0, 1]$; or else a torus, which we can take to be a unit square euclidean torus, $S(1) \times S(1)$ (with any marking). This gives a riemannian structure to $\partial\Psi_0(P)$, which we extend to a riemannian metric on $\Psi_0(P)$. We can choose the metric in a neighbourhood of the boundary curves of $\partial_H\Psi_0(P)$ so that the boundary, $\partial\Psi(P)$ is smoothly embedded in $\Psi(P)$.

Finally, to construct P , we note that each component of $\partial\Psi_0(P)$ is either a square torus, in which case, we glue in a standard $\mathbf{Z} \oplus \mathbf{Z}$ -cusp, or else a bi-infinite cylinder isometric to $S(1) \times \mathbf{R}$ (made up from an annular component of $\partial_V\Psi_0(P)$ together with the vertical boundary components of two model ends), in which case we glue in a standard \mathbf{Z} -cusp.

This gives us our model space, P . We now define a map $f : P \rightarrow M$, in a series of steps as follows.

First, for each $e \in \mathcal{E}_G(M)$, Lemma 12.2 gives us a universally sesquilipschitz map $f : \Psi(P_e) \rightarrow \Psi(E) = e$.

Now suppose that $e \in \mathcal{E}_D(M)$. We want to construct a map $f : \Psi(P_e) \rightarrow \Psi(M)$. This is best done by passing to the cover, $\hat{\Psi}(M)$ of $\Psi(M)$ corresponding to the end, e . Note that $\hat{\Psi}(M) \subseteq \Psi(\hat{M})$, where \hat{M} is the cover of M corresponding to e . (These need not be equal, since a cusp of M may open out in \hat{M} .) Now \hat{M} is a product manifold with base surface $\Sigma(e)$, so that $\Psi(\hat{M})$ is homomorphic to $\Sigma \times \mathbf{R}$, possibly with accidental parabolic cusps removed. In any case, the a-priori bounds theorem (Theorem 11.4) applies in this case. This means that if $\Omega \in \mathcal{W}$ then $l_M(\Omega)$ is bounded above in terms of $\kappa(\Sigma)$, $\max\{l_M(\delta) \mid \delta \in X(\alpha)\}$, and the length bound in the definition of a simply degenerate end (Proposition 11.1). Here l_M denotes the length of the homotopic closed geodesic in M , or equivalently, in \hat{M} . If this happens to be parabolic, we set it equal to 0.

We are now in a position to apply the construction of Sections 6 and 7. This gives us a partition of \mathcal{W}_I as $\mathcal{W}_0 \sqcup \mathcal{W}_1$, and a map $f : \Psi(P_e) \rightarrow \hat{M}$ which is lipschitz on the “thick part”, $\Theta(P_e) = \Psi(P_e) \setminus \bigcup_{\Omega \in \mathcal{W}_0} \text{int } T(\Omega)$ but only, for the moment, defined topologically on the thin part — the union of the Margulis tubes $T(\Omega)$ for $\Omega \in \mathcal{W}_0$.

There are a couple of minor complications in this procedure, which are most simply resolved by observing that we only need to have f defined geometrically on a neighbourhood of the end of $\Psi(P_e)$ — any lipschitz extension to the remainder of $\Psi(P_e)$ will do. The first complication is that some of the annuli in \mathcal{W}_I may correspond to accidental parabolics. The construction will still work in this case, but in any case, there are only finitely many such $\Omega \in \mathcal{W}_I$. Secondly we note that, by construction, f maps each component of $\partial_V\Psi(P_a)$ to the corresponding component of $\Psi(\hat{M})$, but it is still conceivable that $f(\Psi(P_a))$ might enter other components of $\hat{M} \setminus \hat{\Psi}(M)$. As before, f is proper, and sends P_a out an end of $\hat{\Psi}(M)$, and this end cannot contain any such regions. This problem can therefore only arise in a compact subset of $\Psi(P_a)$ and so can be fixed by the earlier observation. We now end up with a map to $\hat{\Psi}(M)$, which descends to a map $f : \Psi(P_a) \rightarrow \Psi(M)$.

Since $f : \Psi(P_a) \rightarrow \Psi(M)$ is proper, it must send $\Psi(P_a)$ out some end e' of $\Psi(M)$. If $e \neq e'$, then the corresponding base surfaces must be homotopic in $\Psi(M)$. Thus, applying Waldhausen’s cobordism theorem, we see that, in fact, $\Psi(M)$, is just a product $\Sigma \times \mathbf{R}$, and so we are in the doubly degenerate situation dealt with in Section 11. We saw there that $e = e'$ giving a contradiction. In other words, we have shown $\Psi(P_a)$ must get sent out the corresponding end of $\Psi(M)$.

We now have f defined on each of the ends of $\Psi(P)$ and hence on all of $\partial_H \Psi_0(P)$. We now extend to any lipschitz map of $\Psi_0(P)$ into $\Psi(M)$, in the right homotopy class, such that each component of $\partial_V \Psi_0(P)$ gets sent to the corresponding component of $\Psi(M)$.

This gives us a proper, end-respecting, homotopy equivalence $f : \Psi(P) \longrightarrow \Psi(M)$, for the moment only defined topologically on the margulis tubes of the degenerate ends. In particular, f is surjective.

Suppose $e \in \mathcal{E}_D(M)$. Since f is proper, we can find a neighbourhood $\Psi(M_e) \cong \Sigma(e) \times [0, \infty)$ of this end in $\Psi(M)$ such that $f^{-1}\Psi(M_e) \subseteq \Psi(P_e)$. We can also find a neighbourhood, $\Psi_1(P_e) \subseteq \Psi(P_e)$ so that $f(\Psi_1(P_e)) \subseteq \Psi(M_e)$. We can also assume that all the properties (W1)–(W9) of Section 4 hold in $\Psi_1(P_e)$. To understand this end, we are thus effectively reduced to considering the map $f|_{\Psi_1(P_e)}$ into $\Psi(M_e)$. Since we only need to control the geometry of the map in some neighbourhood of the end, we can deem any finite set of Margulis tubes in $\Psi(M_e)$, and their preimages in $\Psi_1(P_e)$, to lie in the the respective “thick parts”. In particular, we can assume that $\partial_H \Psi(M_e)$ and $\partial_H \Psi_1(P_e)$ lie in the thick part, and that $f(\partial_H \Psi_1(P_e))$ is homotopic in $\Psi(M_e)$ to $\partial_H \Psi(M_e)$. Now all the arguments of Section 9 and 10 go through as before. For the pushing argument of Section 10, we need to assume that our points lie sufficiently far out the end, in order to push our path into a band, but we only need to verify the sesquilipschitz property on some neighbourhood of the end.

We can thus extend f to a uniformly lipschitz map on each of the Margulis tubes in $\Psi_1(P_e)$, and we deduce that $f|_{\Psi_1(P_e)}$ is universally sesquilipschitz to $\Psi(M_e)$. We can take f to be any lipschitz map in the right homotopy class in the remaining Margulis tubes in $\Psi(P_e)$.

We are now ready to show:

Lemma 12.4 : *The map $f : \Psi(P) \longrightarrow \Psi(M)$ is universally sesquilipschitz.*

Proof : By construction, f is a proper lipschitz map. For each end $e \in \mathcal{E}(M)$ we can choose any product neighbourhood, $\Psi_1(M_e)$, so any two distinct $\Psi_1(M_e)$ are distance $\eta > 0$ apart for some constant $\eta > 0$. If $e \in \mathcal{E}_D(M)$ we can also take $\Psi_1(M_e) \subseteq \Psi(M_e)$ as defined above. Let $\Psi_1(P) \subseteq \Psi(P)$ be a core containing the preimage of an η -neighbourhood of $\Psi(M) \setminus \bigcup_{e \in \mathcal{E}(M)} \Psi_1(M_e)$. Since $\Psi_1(M)$ is compact, the map $f|_{\Psi_1(P)}$ is sesquilipschitz onto its range.

We want to show that the lift of f to the universal covers of $\Psi(P)$ and $\Psi(M)$ is a quasi-isometry. It is sufficient to bound the distance between two points in the domain that get sent to points at most η apart in the range. But this is now easy given that the are such bounds on each component of the lifts of $\Psi_1(P)$ and each $\Psi_1(P_e)$. \diamond

We finally need to define $f : P \longrightarrow M$. In other words, we need to extend f over each cusp R of P . Let R' be the corresponding cusp in M .

If R is a $\mathbf{Z} \oplus \mathbf{Z}$ -cusp, we simply extend the bilipschitz homeomorphism, $f|_{\partial R} : \partial R \longrightarrow \partial R'$ to a bilipschitz homeomorphism $f : R \longrightarrow R'$ by sending each geodesic ray to a geodesic ray.

Suppose R is a \mathbf{Z} -cusp. Thus ∂R is a bi-infinite cylinder, and each of its ends is a vertical boundary components of model ends. For a geometrically finite model end, such

a boundary component will be geodesically embedded. To deal with the general situation, we need to pass to the covers of $\Psi(P)$ and $\Psi(M)$ corresponding to ∂R . In a degenerate end, the same argument as Lemma 4.8 shows that its intersection with the lift of ∂R is quasi-isometrically embed in the lift of the end. Since the two ends lift to disjoint sets, it now follows that ∂R is quasi-isometrically embedded in the cover of $\Psi(P)$. It now follows that the map $f|_{\partial R} : \partial R \rightarrow \partial R'$ is a quasi-isometry with respect to the induced euclidean path metrics. We can thus extend $f|_{\partial R}$ to a universally sesquilipschitz map $f : R \rightarrow R'$ exactly as in Section 11.

We have now defined $f : P \rightarrow M$.

Proof of Theorem 12.1 : We know that $f : P \rightarrow M$ is a proper lipschitz homotopy equivalence, that it respects the decompositions of P and M into non-cuspidal parts and cusps, and that it is universally sesquilipschitz between the non-cuspidal parts and between corresponding cusps. It now follows easily that f is itself universally sesquilipschitz. \diamond

As a consequence, we immediately get:

Proposition 12.5 : *Suppose that M, M' are complete indecomposable hyperbolic 3-manifolds and that there is a homeomorphism from M to M' that sends cusps of M into cusps of M' and conversely. Suppose that the induced map between the non-cuspidal parts sends each geometrically finite end to a geometrically finite end and each degenerate end to a degenerate end. Suppose that (under the induced homeomorphisms of base surfaces) the end invariants of corresponding pairs of degenerate ends are equal. Then there is a an equivariant quasi-isometry between the universal covers of M and M' .*

Proof : We can use the same model manifold P for both M and M' . Theorem 12.1 tells us that there are universally sesquilipschitz maps $P \rightarrow M$ and $P \rightarrow M'$. The lifts then give us equivariant quasi-isometries between the universal covers. \diamond

Note that if there are no geometrically finite ends, then Sullivan's theorem tells us that M and M' are isometric, exactly as in Theorem 11.12.

In general we need to take account of the geometrically finite end invariants:

Theorem 12.6 : *Let M, M' be as in Proposition 12.5, and assume, in addition that the corresponding geometrically finite end invariants are also equal. Then the homeomorphism between M and M' is homotopic to an isometry.*

One way to prove Theorem 12.6 would be construct a model space using geometrically finite model ends that take account of the end invariants as in [Mi4]. In this case, one would show that the quasiconformal extension of the quasi-isometry given by Proposition 12.5 would be conformal. This is the approach taken in [BrocCM1].

Given Proposition 12.5 as stated, one can also proceed as follows.

Write $M = \mathbf{H}^3/\Gamma$ and $M' = \mathbf{H}^3/\Gamma'$, where $\Gamma \cong \pi_1(M) \cong \pi_1(M') \cong \Gamma'$, and write $D(\Gamma)$ and $D(\Gamma')$ for the discontinuity domains. By Ahlfors's finiteness theorem, $D(\Gamma)/\Gamma$ and $D(\Gamma')/\Gamma'$ are (possibly disconnected) Riemann surfaces of finite type. Our indecom-

possibility assumption tells us that each component of either discontinuity domain is a disc.

Now Proposition 12.5 gives us a quasi-isometry from \mathbf{H}^3 to itself which is equivariant with respect to these actions. This extends to an equivariant quasiconformal map, $f : \partial\mathbf{H}^3 \rightarrow \partial\mathbf{H}^3$. This maps $D(\Gamma)$ to $D(\Gamma')$ and descends to a quasiconformal map $\bar{f} : D(\Gamma)/\Gamma \rightarrow D(\Gamma')/\Gamma'$. Since the geometrically finite end invariants are equal, \bar{f} is homotopic to a conformal map $\bar{g} : D(\Gamma)/\Gamma \rightarrow D(\Gamma')/\Gamma'$. We can now lift \bar{g} to an equivariant conformal map $g : D(\Gamma) \rightarrow D(\Gamma')$. We set g to be equal to f on the limit sets. We thus get an equivariant bijection $g : \partial\mathbf{H}^3 \rightarrow \partial\mathbf{H}^3$, which is conformal on the discontinuity domains and a homeomorphism of limit sets.

If we knew that g were quasiconformal, we would see that it was conformal, and hence be finished. However, it is not immediately clear even that g is continuous. We are saved by the following:

Lemma 12.7 : *Suppose that $U \subseteq \mathbf{C}$ is a proper simply connected domain. Suppose that $f : U \rightarrow U$ moves each point a bounded distance with respect to the Poincaré metric. Then the extension of f by the identity of $\mathbf{C} \setminus U$ is continuous. Moreover, if $f|_U$ is quasiconformal, then its extension is uniformly quasiconformal.*

Here of course, the ‘‘Poincaré metric’’ refers to the unique complete curvature -1 metric in the conformal class.

Proof : Write d_e for the euclidean metric. Suppose that $z \in U$ and $d(z, \partial U) = r$. Using the Koebe quarter theorem to compare with the Poincaré metric on the disc of radius r centred at z , we get the well-known estimate $|ds| \geq \frac{2}{r}|dz|$, where $|ds|$ is infinitesimal Poincaré metric on U . Integrating, we deduce that if $z, w \in U$ are distance at most k apart in the Poincaré metric, then $d_e(z, w) \leq (e^{2k} - 1) \max\{d_e(z, \partial U), d_e(w, \partial U)\}$. Continuity of g at ∂U now follows easily. The fact that it is quasiconformal follows, for example, using the above estimate to control the metric quasiconformal distortion of g on ∂U . \diamond

Proof of Theorem 12.6 : Let $h = g^{-1} \circ f : \partial\mathbf{H}^3 \rightarrow \partial\mathbf{H}^3$. By construction, h is Γ -equivariant, and quasiconformal on $D(\Gamma)$ and the identity on the limit set. Let U be a component of $D(\Gamma)$, conformally a disc. Now it is well-known that a quasiconformal map of the disc is a quasi-isometry of the Poincaré metric. (This is based on the fact that the modulus of any embedded annulus that separates two points from infinity is bounded above in terms of the hyperbolic distance between them.) It thus extends to a homeomorphism of the ideal boundary. Since the map h is equivariant with respect a finite co-area action, it follows that it must be the identity on the ideal boundary, and hence moves every point a bounded distance in the Poincaré metric. Thus by Lemma 12.7 it extends to a continuous map that is the identity on the boundary ∂U of U in $\partial\mathbf{H}^3$. Now since this holds for every such component, Lemma 12.7 tells us that h is continuous and quasiconformal on $\partial\mathbf{H}^3$. Since f is quasiconformal, it follows that g is quasiconformal. But it is conformal on $D(\Gamma)$ and hence, applying Sullivan’s result [Su] it is conformal everywhere. Thus there is a hyperbolic isometry conjugating the Γ action to the Γ' action. \diamond

We can make a few remarks. A consequence of tameness is the Ahlfors conjecture which says that either the limit set is all of $\partial\mathbf{H}^3$ or else has measure 0 (see [Can]). In the latter case (namely where there is at least one geometrically finite end) Sullivan's theorem is redundant. We are only using the fact that a quasiconformal map that is conformal almost everywhere is conformal everywhere. Also, it is conjectured, but not proven in general, that the limit set is locally connected. In this case, ∂U , is naturally the continuous image of its ideal boundary, which simplifies the argument somewhat.

References.

- [Ag] I.Agol, *Tameness and hyperbolic 3-manifolds* : preprint, Chicago (2004).
- [Ah] L.V.Ahlfors, *Finitely generated Kleinian groups* : Amer. J. Math. **86** (1964) 413–429, and **87** (1965) 759.
- [BeK] L.Bers, I.Kra (eds.), *A crash course on Kleinian groups* : Lecture notes in Mathematics No. 400, Springer Verlag (1974).
- [Bon] F.Bonahon, *Bouts des variétés hyperboliques de dimension 3* : Ann. of Math. **124** (1986) 71–158.
- [Bow1] B.H.Bowditch, *Hyperbolic 3-manifolds and the geometry of the curve complex* : in “European Congress of Mathematics, Stockholm, June 27 – July 2, 2004” (ed. A.Laptev) European Mathematical Society Publishing House (2005) 103–115.
- [Bow2] B.H.Bowditch, *Intersection numbers and the hyperbolicity of the curve complex* : to appear in J. reine angew. Math.
- [Bow3] B.H.Bowditch, *Systems of bands in hyperbolic 3-manifolds* : preprint, Southampton (2003).
- [Bow4] B.H.Bowditch, *Length bounds on curves arising from tight geodesics* : preprint, Southampton (2003).
- [Bow5] B.H.Bowditch, *End invariants of hyperbolic 3-manifolds* : preprint, Southampton (2005).
- [Bow6] B.H.Bowditch, *Notes on tameness* : preprint, Southampton (2005).
- [Brocb] J.F.Brock, K.Bromberg, *Geometric inflexibility and 3-manifolds that fiber over the circle* : in preparation.
- [BrocbES] J.F.Brock, K.Bromberg, R.Evans, J.Souto, *Grafting, drilling and ending laminations* : in preparation.
- [BrocbCM1] J.F.Brock, R.D.Canary, Y.N.Minsky, *Classification of Kleinian surface groups II: The ending lamination conjecture* : preprint (2004).
- [BrocbCM2] J.F.Brock, R.D.Canary, Y.N.Minsky, *The classification of finitely generated kleinian groups* : in preparation.

- [Brow] E.M.Brown, *Unknotting in $M^2 \times I$* : Trans. Amer. Math. Soc. **123** (1966) 480–505.
- [CalG] D.Calegari, D.Gabai, *Shrinkwrapping and the taming of hyperbolic 3-manifolds* : J. Amer. Math. Soc. **19** (2006) 385–446.
- [Can] R.D.Canary, *Ends of hyperbolic 3-manifolds* : J. Amer. Math. Soc. **6** (1993) 1–35.
- [FHS] M.Freedman, J.Hass, P.Scott, *Least area incompressible surfaces in 3-manifolds* : Invent. Math. **71** (1983) 609–642.
- [Gr1] 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.
- [Gr2] M.Gromov (edited by J.LaFontaine and P.Pansu), *Metric structures for Riemannian and non-Riemannian spaces* : Progress in Mathematics No. 152, Birkhäuser (1998).
- [Ham1] U.Hamenstädt, *Train tracks and the Gromov boundary of the complex of curves* : in “Spaces of Kleinian groups” (eds. Y.Minsky, M.Sakuma, C.Series) London Math. Soc. Lecture Notes No. 329 (2005) 187–207.
- [Ham2] U.Hamenstädt, *Geometry of the complex of curves and of Teichmüller space* : Preprint, Bonn (2005).
- [Har] 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.
- [He1] J.Hempel, *3-manifolds* : Ann. of Math. Studies, No. 86, Princeton University Press (1976).
- [He2] J.Hempel, *Residual finiteness for 3-manifolds* : in “Combinatorial group theory and topology (Alta, Utah, 1984)” (ed. S.M.Gersten, J.R.Stallings) Ann. of Math. Studies No. 111, Princeton Univ. Press (1987) 379–396.
- [Ka] M.Kapovich, *Hyperbolic manifolds and discrete groups* : Progress in Mathematics No. 183, Birkhäuser (2001).
- [Kl] E.Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space* : preprint, Michigan (1999).
- [Mar] A.Marden, *The geometry of finitely generated kleinian groups* : Ann. of Math. **99** (1974) 383–462.
- [MasM1] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103–149.
- [MasM2] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.
- [Mc] D.McCullough, *Compact submanifolds of 3-manifolds with boundary* : Quart. J. Math. **37** (1986) 299–307.
- [Mi1] Y.N.Minsky, *Teichmüller geodesics and ends of hyperbolic 3-manifolds* : Topology **32** (1993) 625–647.

- [Mi2] Y.N.Minsky, *On rigidity, limit sets, and ends of hyperbolic 3-manifolds* : J. Amer. Math. Soc. **7** (1994) 539–588.
- [Mi3] Y.N.Minsky, *The classification of punctured torus groups* : Ann. of Math. **149** (1999) 559–626.
- [Mi4] Y.N.Minsky, *The classification of Kleinian surface groups I : Models and bounds* : preprint, Stony Brook (2002).
- [O1] J.-P.Otal, *Thurston’s hyperbolization of Haken manifolds* : Surveys in differential geometry, Vol. III, 77–194, International Press (1998).
- [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.
- [Sc1] P.Scott, *Finitely generated 3-manifold groups are finitely presented* : J. London Math. Soc. **6** (1973) 437–440.
- [Sc2] P.Scott, *Compact submanifolds of 3-manifolds* : J. London Math. Soc. **7** (1973) 246–250.
- [Sc3] P.Scott, *Subgroups of surface groups are almost geometric* : J. London Math. Soc. **17** (1978) 555–565.
- [Sh] K.J.Shackleton, *Tightness and computing distances in the curve complex* : preprint, Southampton (2004).
- [So] T.Soma, *Existence of polygonal wrapping in hyperbolic 3-manifolds* : preprint, Tokyo Denki (2005).
- [Su] D.Sullivan, *Quasiconformal homeomorphisms and dynamics II. Structural stability implies hyperbolicity for kleinian groups* : Acta Math. **155** (1985) 243–260.
- [T] W.P.Thurston, *The geometry and topology of 3-manifolds* : notes, Princeton (1979).
- [Wal] F.Waldhausen, *On irreducible 3-manifolds that are sufficiently large* : Ann. of Math. **87** (1968) 56–88.
- [Wan] S.Wang, *Non-zero degree maps between 3-manifolds* : in “Proceedings of the International Congress of Mathematicians, Vol. II (Beijing, 2002)” (ed. T.Li) Higher Ed. Press (2002), 457–468.