Coarse hyperbolic models for 3-manifolds

Brian H. Bowditch
Mathematics Institute, University of Warwick,
Coventry, CV4 7AL, Great Britain.
[September 2010]

0. Introduction.

The tameness and ending lamination conjectures together tell us that a hyperbolic 3-manifold with finitely generated fundamental group is determined by its topology and a finite number of "end invariants". In this paper we describe how some of this theory generalises to a much broader class of metrics. To simplify the discussion, we focus on the particular case where the 3-manifold is homotopy equivalent to a compact surface. We will state the main result (Theorem 0.7) in the "doubly degenerate" case. This case illustrates the main features of the argument, though further generalisations are possible, as we will briefly discuss. The ending lamination conjecture is closely related to the large scale geometry of Teichmüller space, and one of the main motivations for this study is its potential applications in that direction, for example to the Weil-Petersson metric.

To be more precise, let Σ be a compact orientable surface. Let M be an orientable complete riemannian 3-manifold, with a preferred homotopy equivalence $M \longrightarrow \Sigma$. We give some hypotheses under which such a manifold will serve as a "model" of a (constant curvature) hyperbolic 3-manifold. The main requirements can be paraphrased by saying that M has locally bounded geometry and a thick-thin decomposition with "standard" thin part, and that the universal cover is Gromov hyperbolic. The last assumption turns out to be equivalent to asserting that the thick part is hyperbolic relative to the thin part. There are many variations on the hypotheses that would work as well, as we will elaborate in Section 2. We begin with a more precise formulation.

Given $x \in M$, we define the *essential systole* of M at x, denoted $\mathrm{sys}(M,x)$, to be the length of the shortest homotopically non-trivial loop in M passing through x. A free homotopy class of closed curves in M is *parabolic* if it has arbitrarily short representatives in M. It *peripheral* if its image in Σ can be homotoped into $\partial\Sigma$. We denote by D^n the unit euclidean n-ball.

We make the following assumptions on M:

(M1) There is some $c_0 > 0$ such that all sectional curvatures of M lie between $-c_0$ and c_0.

(M2) There is some $\eta_0 > 0$ such that if $x \in M$ with $\mathrm{sys}(M,x) \le \eta_0$ then M has all sectional curvatures at x equal to -1.

(M3) $\tilde M$ is h_0-hyperbolic for some $h_0 \ge 0$.

(M4)(a) $(\forall\xi)(\exists\xi')$ such that any ξ-lipschitz map $f : \partial D^2 \longrightarrow \tilde M$ extends to a ξ'-lipschitz map $f : D^2 \longrightarrow \tilde M$.

(b) $(\forall\xi)(\exists\xi')$ such that any ξ-lipschitz map $f : \partial D^2 \longrightarrow \tilde M$ extends to a ξ'-lipschitz map $f : D^2 \longrightarrow \tilde M$.

(M5) A closed curve in M is parabolic in M if and only if it is peripheral in Σ.

We refer to a 3-manifold satisfying (M1)–(M5) as a *coarse hyperbolic 3-manifold*. By its *parameters*, we mean the various constants (or functions) featuring in the hypotheses (including the topological type of Σ). We elaborate on these hypotheses in Section 1.

Note that (M1)–(M4) are all automatic for a genuine (constant curvature) hyperbolic 3-manifold. In this case, the only parameter is the topological type of Σ.

Of these hypotheses, clearly (M2) is the most artificial. One justification is that this condition is an immediate consequence of the way in which one might expect to construct model manifolds in practice. We shall elaborate on this later. We first describe how it gives rise to a thick-thin decomposition of M.

Given $\eta > 0$, we write:

$$\Theta(M, \eta) = \{x \in M \mid \mathrm{sys}(M, x) \geq \eta\}$$

$$\Upsilon(M, \eta) = \{x \in M \mid \mathrm{sys}(M, x) \leq \eta\}.$$

These are closed subsets, which we refer to as the "thick" and "thin" parts respectively.

Proposition 0.1 : *There is some $\eta_1 \leq \eta_0$ depending on the parameters of M, such that each component of $\Upsilon(M, \eta_1)$ is a standard Margulis region. Its fundamental group injects into $\pi_1(M)$. Moreover, no essential curve can be freely homotoped into distinct components of $\Upsilon(M, \eta_1)$.* $\diamondsuit$

By a *standard Margulis region* we mean a subset isometric to a Margulis region in a genuine (constant curvature) hyperbolic manifold. This is either a Margulis tube or Margulis cusp. In this case, since there is no $\mathbf{Z} \oplus \mathbf{Z}$ subgroup, all cusps will be $\mathbf{Z}$-cusps.

A consequence of Proposition 0.1 and hypothesis (M3) is that there is exactly one Margulis cusp for each boundary curve of Σ.

Given $\eta \leq \eta_1$, let $\mathcal{T} = \mathcal{T}(M, \eta)$ be the set of Margulis tubes, and let $\mathcal{P} = \mathcal{P}(M, \eta)$ be the set of Margulis tubes, Let $\Psi(M, \eta) = M \setminus \mathrm{int} \bigcup \mathcal{P}(\eta)$ be the *non-cuspidal part* of M. Note that each component of $\partial\Psi(M, \eta)$ is homotopic to a component of $\partial\Sigma$. In fact, there is a relative homotopy equivalence $(\Psi(M, \eta), \partial\Psi(M, \eta)) \longrightarrow (\Sigma, \partial\Sigma)$.

In what follows, we fix some positive $\eta \leq \eta_1$, and abbreviate $\Psi = \Psi(M, \eta)$. This depends only on the parameters of M.

Since we are assuming that Σ and M are orientable, it follows that Ψ has two ends, arbitrarily designated "positive" and "negative" and denoted e_+ and e_-.

Definition : The end e of Ψ is *simply degenerate* if there is some $l \geq 0$ and a sequence of essential closed curves, $(\delta_i)_i$, realised in Ψ, with $\delta_i \to e$, such that for each i, $\mathrm{length}(\delta_i) \leq l$, and δ_i is simple.

Here $\delta_i \to e$ means the δ_i eventually leaves every compact set some fixed closed neighbourhood of e. By "simple" we mean homotopic to a simple closed curve in Σ. A number of equivalent formulations of simple degeneracy will be discussed in Section 1 and proven in Section 4.

Definition : We say that e is *topologically finite* if it has a neighbourhood homeomorphic to $\Sigma \times [0, \infty)$.

(If one wants to circumvent the Poincaré conjecture in the following discussion one can simply add the hypothesis that M be irreducible.)

We show:

Theorem 0.2 : *If e is simply degenerate, then it is topologically finite.*

To proceed, we need to recall that notion of the "curve graph", that is, the 1-skeleton of the curve complex defined in [H].

Let $X(\Sigma)$ be the set of homotopy classes of essential non-peripheral simple closed curves in Σ. The curve graph $\mathcal{G} = \mathcal{G}(\Sigma)$ has vertex set $V(\mathcal{G}) = X(\Sigma)$, with two curves in $X(\Sigma)$ adjacent in $\mathcal{G}(\Sigma)$ if they have minimal possible geometric intersection number (that is, 2 for a four-holed sphere, 1 for a one-holed torus, and 0 for all more complex surfaces). It turns out that $\mathcal{G}$ is hyperbolic in the sense of Gromov [MaM1]. Its boundary, $\partial \mathcal{G}$ can be identified with the set of arational laminations on Σ [K], though we won't formally need to know that here.

Given $\gamma \in X(\Sigma)$, let $l_M(\gamma) \in [0, \infty)$ be the infimum of lengths of realisations of γ in M. Given $\gamma \in X(\Sigma)$, write $X(M, l) = \{\gamma \in X(\Sigma) \mid l_M(\gamma) \leq l\}$.

Theorem 0.3 : *There is some $l_0 \geq 0$, depending only on the parameters, such that $X(M, l_0)$ is non-empty and uniformly quasiconvex in $\mathcal{G}(\Sigma)$. Moreover, for all $l \geq 0$, there is some $t \geq 0$ depending only on l and the parameters such that $X(M, l) \subseteq N(X(M, l_0), r)$.*

Here "uniformly" means depending only on the parameters of M. The notation $N(., t)$ denotes metric t-neighbourhood (here in $\mathcal{G}(\Sigma)$).

Theorem 0.4 : *Suppose that e is a simply degenerate end of $\Psi(M)$. Then there is some $a \in \partial \mathcal{G}$ with the following property. Suppose $(\delta_i)_i$ is a sequence of curves realised in $\Psi(M)$ with $\delta_i \to e$ and with* $\mathrm{length}(\delta_i) \leq l$ *for some $l \geq 0$. Suppose that each δ_i is simple in Σ. Then the corresponding elements of $X(\Sigma)$ tend to a in $\mathcal{G} \cup \partial \mathcal{G}$.*

We can make a number of remarks on this theorem.

First, the existence of such a sequence, $(\delta_i)_i$, is precisely the definition of a simply degenerate end, and so the element $a \in \partial \mathcal{G}$ is unique. We denote it $a(e)$, and refer to it as the *end invariant* of e.

In fact, we will see that any such sequence $(\delta_i)_i$ remains a bounded distance from a geodesic ray in $\mathcal{G}$. Also, we can always find such a sequence so that $\mathrm{length}(\delta_i) \leq l_0$ for all i, where l_0 depends only on the parameters of M (cf. Theorem 0.3).

To simplify our discussion, we will focus on the *doubly degenerate* case. This will avoid various qualifications in the statement and proofs, though the arguments are applicable more generally (see the discussion in Sections 3 and 4).

Definition : We say that M is *doubly degenerate* if both e_- and e_+ are simply degenerate.

By Theorem 0.2, it follows that Ψ is topologically finite, hence (by Waldhausen's h-cobordism theorem) homeomorphic to $\Sigma \times \mathbf{R}$.

Theorem 0.5 : Suppose that M is doubly degenerate. Then:
(1) $a(e_-) \neq a(e_+)$
(2) $X(M, l_0)$ is a uniformly bounded distance in $\mathcal{G}(\Sigma)$ from some (hence any) bi-infinite geodesic from $a(e_-)$ to $a(e_+)$.

(It is known from [MaM2] that such a bi-infinite geodesic always exists.)

Recall that $\mathcal{T}$ is the set of η_1-Margulis tubes in Ψ. The thin part of Ψ consists of the disjoint union of elements of $\mathcal{T}$, no two of which are homotopic in Ψ (see Proposition 0.1).

The following result generalises that of Otal for hyperbolic 3-manifolds [Ot].

Theorem 0.6 : *If M is doubly degenerate, then $\mathcal{T}$ is unlinked in Ψ.*

One way to formulate this is to say that there is an indexing set, $I \subseteq \mathbf{Z}$, a collection, $(A_i)_{i \in I}$, of essential annuli in Σ, a collection, $(J_i)_{i \in I}$, of disjoint closed intervals in $\mathbf{R}$, and a homeomorphism from $\Sigma \times \mathbf{R}$ to M sending $\bigcup_{i \in I}(A_i \times J_i)$ onto $\bigcup \mathcal{T}$. For elaboration, see [Bow5]. Theorem 0.6 supposes that η has been chosen sufficiently small in relation to the parameters.

In particular, each tube $T \in \mathcal{T}$ is homotopic to a unique $\delta \in X(\Sigma)$.

We can finally state the main result of this paper — a version of the ending lamination conjecture for coarse hyperbolic 3-manifolds:

Theorem 0.7 : *Suppose that M and M' are doubly degenerate coarse hyperbolic 3-manifold with the same base surface, Σ, and with ends $e_\pm$ and $e'_\pm$ respectively. If $a(e_-) = a(e_+)$, then there is a uniform equivariant quasi-isometry from $\tilde{M}$ to $\tilde{M}'$.*

Implicit in this statement are the preferred homotopy equivalences from M and M' to Σ. The quasi-isometry is equivariant with respect to the action of $\pi_1(\Sigma)$ by covering translations. The term "uniform" means depending only on the parameters of M and M'.

In fact, the quasi-isometry can be assumed to respect the universal covers of the non-cuspidal parts, $\Psi(M)$ and $\Psi(M')$. It also respects the positive and negative ends of the non-cuspidal parts, in a sense that will be made precise later (see Section 6).

It seems likely that one can, in fact, construct a uniform bilipschitz map from M to M' also sending $\Psi(M)$ to $\Psi(M')$, (cf. [BrocCM]), though shall not explore that here.

In the case where both M and M' are hyperbolic 3-manifolds, this gives the ending lamination conjecture for type preserving orientable surface groups — the existence of an equivariant quasi-isometry in this case gives rise to an isometry from M to M'.

One can also view Theorem 0.7 as telling us that hypotheses (M1)–(M5) defining a coarse hyperbolic 3-manifold are sufficient for it to serve as a "model" for a genuine hyperbolic 3-manifold, as used in the proof of the ending lamination conjecture [Mi,BrocCM] (see also [Bow5],[S]).

In fact, the proof of Theorem 0.7 follows the argument of the above. Given distinct $a_-, a_+ \in \partial\mathcal{G}$, one constructs a combinatorial model space, $P = P(a_-, a_+)$, depending only on a_- and a_+. If M is a doubly degenerate coarse hyperbolic 3-manifold with $a(e_\pm) = a_\pm$, then we construct a uniformly lipschitz map $P \longrightarrow M$ such that the lift to the universal covers, $\tilde{P} \longrightarrow \tilde{M}$, is a uniform quasi-isometry. Thus, if M' is another such, we also get an equivariant quasi-isometry $\tilde{M} \longrightarrow \tilde{M}'$.

The construction of P is very similar to that in [Mi], though we shall follow the account given later in [Bow5]. We show how the arguments of [Bow5] go through with the weaker hypotheses. The key to this is a version of the a-priori bounds theorem (cf. [Mi]) applicable in this generality, see [Bow7].

We note that in the case where M is closed and the systole of M is bounded below (so that we can take $\Theta(M) = \Psi(M) = M$), Theorem 0.7 is effectively already known from [Mo] and independently from [Bow4]. The formulations are a little different, and there are some technical points to be addressed to make the correspondence precise. Essentially, a bi-infinite path, π, in the thick part of Teichmüller space gives rise to a riemannan metric on $\Sigma \times \mathbf{R}$, well defined up to bilipschitz homeomorphism. It is shown that the universal cover of this 3-manifold is Gromov hyperbolic if and only if π remains a bounded distance from a Teichmüller geodesic. From this one can deduce that if we have two such manifolds with the same end invariants, then their covers are equivariantly quasi-isometric. It would require a bit of extra work to verify that any coarse hyperbolic 3-manifold with empty thin part, as we have defined it, arises in this way (that is, up to equivariant quasi-isometry of covers), and we will not address that issue here.

We also note that a class of coarse hyperbolic 3-manifolds has also been explored in [Ba]. There it is assumed that the 3-manifold covers a compact 3-manifold with hyperbolic fundamental group, though it seems likely that much of this can be applied coarse hyperbolic 3-manifolds, as we have defined them, under the assumption of a lower bound on systole.

1. Comments on the hypotheses.

We shall discuss the significance of the various assumptions in turn, and describe possible variations thereof.

1.1. Bounded geometry.

The only significance of hypothesis (M1) is that it gives us locally bounded geometry. The conclusion we want to draw from it can be phrased as follows. We say that a subset $Q \subseteq \tilde{M}$ is *r-separated* if $d(x, y) \geq r$ for all distinct $x, y \in Q$.

Lemma 1.1.1 : $(\forall r, s)(\exists n)$ *such that if* $Q \subseteq \tilde{M}$ *is* r*-separated and has diameter at most* s*, then* $|Q| \leq n$*.*

This is a standard packing argument. In fact, only a bound on Ricci curvature is required in (M1).

1.2. The thick-thin decomposition.

Hypothesis (M2) gives rise to a thick-thin decomposition with standard thin parts. Much of this applies to any group action on a Gromov k-hyperbolic space, H (see [Grom,GhH]).

We begin with a general observation about quasiconvex sets in H. Recall that if Y is any geodesic metric space and $Q \subseteq Y$ is a closed subset, then Q is *r-quasiconvex* in Y if every geodesic in Y with endpoints in Q lies entirely in $N(Q,r)$. It is *convex* if it is 0-quasiconvex.

Lemma 1.2.1 : *Given k there is some h' with the following property. Suppose that H is k-hyperbolic, and $Q \subseteq H$ is closed and connected. Suppose that for some $r \geq 0$, Q is r-quasiconvex in $N(Q, r+h')$ with respect to the induced path metric on $N(Q, r+h')$. Then Q is r-quasiconvex in H.*

Proof : Suppose $x \in Q$. Let $P \subseteq Q$ be the set of $y \in Q$ such that any geodesic in H from x to y lies in $N(Q,r)$. We claim that $P = Q$. Certainly $x \in P$, so $P \neq \emptyset$. Suppose $y, z \in Q$ with $d_H(y,z) \leq h$, say. We can choose $h' \geq h$ so that any pair of H-geodesics, $[x,y]$ and $[x,z]$, from x to y and z respectively are less than Hausdorff distance h' apart. If $y \in P$, then $[x,y] \subseteq N(Q,r)$ and so $[x,z] \subseteq N(Q, r+h')$. Now $[x,z]$ is intrinsically geodesic also in $N(Q, r+h')$ and so in fact, $[x,z] \subseteq N(Q,r)$ (this neighbourhood being the same whether measured in $N(Q, r+h')$ or in H itself). Thus, $z \in P$. By connectedness of Q, it now follows that $P = Q$. Since $x \in Q$ was arbitrary, it follows that Q is r-quasiconvex in H. $\diamondsuit$

In particular, if Q is convex in $N(Q,h')$ then it is convex in H. The main application will be when $N(Q,h')$ is isometric to a convex subset of hyperbolic space $\mathbf{H}^3$.

Let η_3 be the usual Margulis constant of $\mathbf{H}^3$. (Thus, for example, if M were any hyperbolic 3-manifold, then the thin part $\Upsilon(M,\eta_3)$ is a disjoint union of standard Margulis tubes and cusps.)

Suppose that $M = H/\Gamma$ is a riemannian manifold satisfying (M2), with constant η_0. Set $\eta_4 = \min(\eta_0/3, \eta_3)$, and suppose that $\eta \leq \eta_4$.

Lemma 1.2.2 : *Each component of $\Upsilon(M,\eta)$ is intrinsically isometric to a standard Margulis region.*

Proof : Let T be a component of $\Upsilon(M,\eta)$. Since $\eta < \eta_0/2$, any shortest loop in M through x lies in an open subset of M of constant curvature -1. From this, it's easily seen that ∂T is piecewise smooth. In particular, T is locally simply connected, and so we can construct its universal cover $\tilde{T}$. Thus $T = \tilde{T}/G$ where $G = \pi_1(T)$. Note that the inclusion of T into M induces a homomorphism of G into Γ. We also have a local isometry $q : \tilde{T} \longrightarrow \tilde{M}$, which is a covering map to its range — a component of the preimage of T in $\tilde{M}$.

Now $\tilde{T}$ is locally hyperbolic, so we get a developing map $p : \tilde{T} \longrightarrow \mathbf{H}^3$. This gives rise to a holonomy action of G on $\mathbf{H}^3$. We claim that this is elementary, i.e. either parabolic (fixing a point of $\partial \mathbf{H}^3$, or loxodromic (fixing setwise a bi-infinite geodesic in $\mathbf{H}^3$).

To see this, choose any $x \in \tilde{T}$, and let $G_x = \{g \in G \mid d_{\tilde{T}}(x, gx) \le \eta\}$. We first note that G_x is non-trivial, for if γ is a shortest essential curve in M passing through x, then $\gamma \subseteq \Upsilon(M, \eta)$. Thus, $\gamma \subseteq T$ and it has length at most η and so it determines a non-trivial element of G_x. Also, each element of G_x displaces the point $p(x) \in \mathbf{H}^3$ a distance at most $\eta \le \eta_3$. Thus, the Margulis Lemma tells us that $\langle G_x \rangle \le G$ is elementary. Let α_x be the fixed point or axis of $\langle G_x \rangle$. Since the above holds for all $x \in \tilde{T}$, we see by continuity, that $\alpha_x = \alpha$ must, in fact, be independent of x. Since it is determined canonically, it must be preserved by all of G. Thus, G is elementary, as claimed.

Let $\beta = p^{-1}(\alpha) \subseteq \tilde{T}$. Thus β is a (possibly empty) union of geodesics. There is a vector field on $\mathbf{H}^3$ pointing towards α, and singular at α. This pulls back to a vector field on $\tilde{T}$ singular at β. This vector field has an orthogonal foliation by constant distance surfaces or horospheres. Note that the sets G_x are locally constant as we move x in a leaf, and non-decreasing as we flow x in the direction of the field. We see that the flow lines in $\tilde{T}$ either continue forever or else eventually run into β. In the latter case, it then follows that β is in fact, a single bi-infinite geodesic. We also see that, in either case, the orthogonal leaves to the foliation are complete and foliate $\tilde{T}$ (outside β). It is now easy to see that p is in fact injective, and so the action of G on $\mathbf{H}^3$ is properly discontinuous. We can thus identify $\tilde{T}$ with $p(\tilde{T}) \subseteq \mathbf{H}^3$ and T with $p(\tilde{T})/G$. The latter is a standard Margulis region as claimed. ◊

Let T be a component of $\Upsilon(M, \eta)$ and let $G = \pi_1(T)$ and $\phi : G \longrightarrow \Gamma$ be as in the proof of Lemma 1.2.2. In general, ϕ need not be injective. However, it will be if we add a few more assumptions on M. Suppose, for example, that Γ is torsion-free. Since $\phi(G) \subseteq \Gamma$ is non-trivial, ϕ will be injective if $\phi(G) \cong \mathbf{Z}$. The only case it remains to worry about is if T is a $\mathbf{Z} \oplus \mathbf{Z}$-cusp, and $\phi(G) \cong \mathbf{Z}$. Now T retracts onto the torus ∂T. Thus, $M \setminus \operatorname{int} T$ has toroidal boundary which is not π_1-injective. It follows that if $\pi_2(M)$ is trivial, then $M \setminus \operatorname{int} T$ is a solid torus and so $\Gamma \cong \mathbf{Z}$. We conclude:

Lemma 1.2.3 : *If Γ is torsion free and not cyclic, and $\pi_2(M)$ is trivial, then each component of $\Upsilon(M, \eta)$ is π_1-injective.* ◊

We can thus identify $\tilde{T}$ with a component of the preimage of T in $\tilde{M}$.

Now suppose that $H = \tilde{M}$ is h-hyperbolic, and let h' be the constant of Lemma 1.2.1. Using Lemma 1.2.2, we see that there is some $\eta_5 \le \eta_4$ such that $N(\Upsilon(M, \eta_5), h') \subseteq \Upsilon(M, \eta_4)$. We get:

Lemma 1.2.4 : *If $\eta \le \eta_5$ and T is a component of $\Upsilon(M, \eta)$, then $\tilde{T}$ is a convex subset of $\tilde{M}$.*

Proof : By Lemmas 1.2.2 and 1.2.3, $N(\tilde{T}, h')$ is isometric to a convex subset of $\mathbf{H}^3$, and $\tilde{T} \subseteq N(\tilde{T}, h')$ is convex. Thus, by Lemma 1.2.1, $\tilde{T}$ is convex in H. ◊

Now any infinite cyclic subgroup of Γ is either loxodromic (moves a bi-infinite geodesic in H a uniformly bounded distance) or parabolic (fixes a point in ∂H). From the structure of $\tilde{T}$, we see that if T is a Margulis tube, then G is loxodromic, and if T is a Margulis cusp, then G is parabolic.

Suppose that $G \leq G' \leq \Gamma$, with $G' \cong \mathbf{Z}$. From the general structure of loxodromic and parabolic groups, we see that if $g \in G'$, then $\tilde{T}$ and $g\tilde{T}$ are some bounded Hausdorff distance, h'' say, apart. We can now find some $\eta_1 \leq \eta_5$ such that $N(\Upsilon(M, \eta_1), h'') \subseteq \Upsilon(M, \eta_5)$. Thus, if $\eta \leq \eta_1$, it follows that we must, in fact, have $\tilde{T} = g\tilde{T}$. In other words, $g \in G$, and it follows that $G' = G$. Thus:

Lemma 1.2.5 : *If $\eta \leq \eta_1$, and T is a component of $\Upsilon(M, \eta)$, then $\pi_1(T)$ is a maximal infinite cyclic subgroup of Γ.* $\diamondsuit$

We can also assume that h'' is such that any two infinite bi-infinite geodesics in H are Hausdorff distance at most h'' apart. From this it follows that:

Lemma 1.2.6 : *If $\eta \leq \eta_1$, then any two homotopic components of $\Upsilon(M, \eta)$ are equal.* $\diamondsuit$

In the case where Γ is a surface group, we can summarise what we have shown as follows:

Proposition 1.2.7 : *Suppose M satisfies (M2) and (M3) and is homotopy equivalent to a compact surface, Σ. Then there is some η_1, depending only on the parameters, such that if $\eta \leq \eta_1$, then any component of $\Upsilon(M, \eta)$ isometric to a standard Margulis region, and is homotopic to a primitive curve in Σ. There is at most one region corresponding to any given curve. Each component of the lift to $\tilde{M}$ is convex in $\tilde{M}$, and its stabiliser is loxodromic or parabolic in $\tilde{M}$, depending on whether it is a tube or $\mathbf{Z}$-cusp.* $\diamondsuit$

Note that we are not claiming, at this point, that the curve is simple in Σ.

1.3. Hyperbolicity.

We have already brought the hyperbolicity of $\tilde{M}$ into play in Section 1.2. In Section 2, we will discuss how this can be equivalently formulated in terms of the relative hyperbolicity of the thick part.

We shall note here how hyperbolicity gives rise to a formulation of the thick-thin decomposition that is applicable without reference to riemannian geometry. This will allow us to apply the general results of [Bow7].

Given $x \in \tilde{M}$ and $r \geq 0$, let $\Gamma_r(x) = \{g \in \Gamma \mid d(x, gx) \leq r\}$.

Lemma 1.3.1 : *Suppose that M satisfies (M1), (M2) and (M3). Then for all $r \geq 0$, there is some $\nu \in \mathbf{N}$ such that if $x \in \tilde{M}$, then either $|\Gamma_r(x)| \leq \nu$ or else $\langle \Gamma_r(x) \rangle$ is infinite cyclic. Here ν depends only on the parameters and r.*

Proof : Let η_1 be the constant introduced in Section 1.2. Choose $\eta_6 \leq \eta_1$ such that $N(\Upsilon(M, \eta_6)) \subseteq \Upsilon(M, \eta_1)$. Using (M1), we see that if ν is sufficiently large in relation to r, then if $|\Gamma_r(x)| \geq \nu$ then there must be some non-trivial $g \in \Gamma$ with $d(x, gx) \leq \eta_6$. Thus, $x \in \Upsilon(M, \eta_6)$, and so $N(x, r) \subseteq \Upsilon(M, \eta_1)$. It follows that $\langle \Gamma_r(x) \rangle$ is non-trivial and stabilises a component of $\Upsilon(M, \eta_1)$. It now follows from Section 1.2, that this is infinite cyclic. $\diamondsuit$

The conclusion of Lemma 1.3.1 can be viewed as a coarse version of the Margulis lemma. It is a hypothesis of the "a-priori bounds" theorem in [Bow7] which we will be applying later.

1.4. Isoperimetric inequalities.

A number of remarks can be made regarding the "isoperimetric inequalities" given by (M4).

Although they are global assumptions, given hyperbolicity, they can be reduced to bounded ones. More specifically, there is some $\xi_0 = \xi_0(k)$, depending only on the hyperbolicity constant k, such that if (M4) holds for all $\xi \leq \xi_0$, then it holds for all $\xi \geq 0$. (This is a fairly straightforward excercise, for example, subdividing the curve or 2-sphere into small ones by spheres towards some fixed basepoint.) Moreover, if we assume (M1), then the statement is also automatic on a small scale. Thus, given (M1) and (M3) together, we can replace (M4) by:

(M4$'$) There is some ξ_0, sufficiently large in relation to the hyperbolicity constant of (M3), and some $\xi_0' \geq \xi_0$, such that:

(a) any ξ_0-lipschitz map $f : \partial D^2 \longrightarrow \tilde{M}$ extends to a ξ_0'-lipschitz map $f : D^2 \longrightarrow \tilde{M}$.

(b) any ξ_0-lipschitz map $f : \partial D^2 \longrightarrow \tilde{M}$ extends to a ξ_0'-lipschitz map $f : D^2 \longrightarrow \tilde{M}$.

We also note that if we also assume (M2), then it is possible to phrase the isoperimetric inequalities in a manner intrinsic to the thick part, Θ, of M. For this we need to assume that Θ is constructed so that all Margulis tubes are sufficiently deep in relation to the hyperbolicity constant. In this case, it is not hard to see that discs and balls can be pushed off Margulis tubes while only increasing the lipschitz constant by to controlled amount. To see that inqualities in Θ give inequalities in M, there is a complication in part (a) in that a curve homotopically trivial in M need not in general be homotopically trivial in Θ. However, if the curve is short in relation to the depth of the tubes, then this can in fact be shown to be the case. This will be discussed further in Section 2 (see Lemma 2.2).

One final observation is that one can replace isoperimetric inequalities with isodiametric inequalities. In other words, curves or spheres of bounded diameter bound respectively discs and balls of bounded diameter. A similar discussion with regard to hyperbolicity and restricting to the thick part applies in this case. The main results as stated in Section 0 still hold. The only difference is that the map from the combinatorial model P to M need only be continuous. The lift $\tilde{P} \longrightarrow \tilde{M}$ is still a uniform quasi-isometry. While this is in some ways more natural, we shall stick with the lipschitz version here since it makes the exposition simpler, and ties in with the account in [Bow5].

1.5. Type preserving maps.

The "strictly type preserving" assumption (M5) allows us to elaborate on Proposition 1.2.7. Suppose $\eta \leq \eta_1$. Let $\mathcal{T} = \mathcal{T}(M, \eta)$ and $\mathcal{P} = \mathcal{P}(M, \eta)$ be the sets of Margulis tubes and cusps respectively. From Proposition 1.2.7, we see that the homotopy class of a Margulis region has abribrarily short representatives in M if and only if it is a cusp. Thus (M5) tells us that the elements of $\mathcal{P}$ are in bijective correspondence with peripheral curves of Σ.

Let $\Psi = \Psi(M, \eta) = M \setminus \operatorname{int} \bigcup \mathcal{P}$. Thus M deformation retracts onto Ψ, and each component of $\partial \Psi$ is a euclidean cylinder. It follows that there is a homotopy equivalence of Ψ to Σ sending $\partial \Psi$ to $\partial \Sigma$. Now any self homotopy equivalence of Σ sending $\partial \Sigma$ to itself must be a relative homotopy equivalence of $(\Sigma, \partial \Sigma)$. It follows that there must in fact be a relative homotopy equivalence of $(\Psi, \partial \Psi)$ to $(\Sigma, \partial \Sigma)$.

Now let $\Psi_0 \subseteq \Psi$ be a relative Scott core of $(\Psi, \partial \Psi)$ (see [Mc]). This is homotopy equivalent to $(\Sigma, \partial \Sigma)$ and hence homeomorphic to $(\Sigma \times [0,1], \partial \Sigma \times [0,1])$. It now follows that Ψ has exactly two ends, each of which retracts onto a relative boundary component of Ψ_0 in Ψ, and respecting the manifold boundaries.

We shall designate the ends arbitrarily as e_- and e_+. Each has a neighbourhood $E_\pm$ with relative boundary $\partial_\Psi E_\pm$ homeomorphic to Σ. The pair $(E_\pm, E_\pm \cap \partial \Psi)$ deformation retracts onto $(\partial_\Psi E_\pm, \partial_\Psi E_\pm \cap \partial \Psi) \cong (\Sigma, \partial \Sigma)$.

1.6. Degenerate ends.

Suppose we have fixed a preferred homotopy equivalence from M to Σ. Given a free homotopy class, γ, of non-trivial non-peripheral closed curves in Σ, we write γ^* for a realisation of γ in M as a curve of minimal length. Such a realisation is well defined up to uniformly bounded distance in M. Indeed, its lift to $\tilde{M}$ is well defined up to uniformly bounded hausdorff distance. In the case where γ happens to be homotopy equivalent to a Margulis tube, then γ^* has to be the core of this tube. In all other cases, the above is a general observation using the fact that $\tilde{M}$ is hyperbolic, to show that the lift of γ^* is a uniform quasigeodesic.

We defined an end e of Ψ to be "simply degenerate" if there is a sequence of simple closed curves of bounded length in M going out e. In fact, we shall we see that we can always assume such curves to have minimal length. In other words, that there is a sequence, $(\gamma_i)_i$, in $X(\Sigma)$ such that $\gamma_i^* \to e$, and $\operatorname{length}(\gamma_i^*)$ is bounded. Indeed, we can also take this bound to depend only on the parameters of M.

In fact, we can drop the length requirement altogether in the above:

Proposition 1.6.1 : *Let e be an end of Ψ. Suppose there exist a sequence $(\gamma_i)_i$ in $X(\Sigma)$ and points $x_i \in \gamma_i^*$ such that $x_i \to e$. Then e is simply degenerate.*

We shall prove the above statements in Section 3.

In contrast, we could make the following definition:

Definition : We say that the end, e, is *geometrically finite* if it has a neighbourhood U with $U \cap \gamma^* = \emptyset$ for all closed curves in Σ.

This is equivalent to the usual definition in constant or pinched negative curvature. If we negate this, we arrive at the formulation of a degenerate end given by Proposition 1.6.1, except that the curves are no longer simple. It is natural ask:

Question : Is each end of M either geometrically finite or simply degenerate?

Note that negation of the above definition is equivalent to the hypothesis of Proposition 1.6.1 where the requirement that the the curves γ_i are simple is dropped. The issue is therefore replacing the set of curves with simple ones, as acheived in [Bon] in constant curvature -1. As observed in [C], this generalises to pinched negative curvature. Also, recently Barnard [Ba] has shown that this holds for certain Gromov hyperbolic spaces in the case where there is a lower bound on injectivity radius.

Also, assuming (M1)–(M5) it seems reasonable to ask:

Question : Is a geometrically finite end topologically finite?

We will show in Section 3 that a simply degenerate end is topologically finite.

2. Relative hyperbolicity.

The aim of this section is to explain how the main geometric hypotheses can be separated into two largely independent parts. The first asserts that the thin part is "standard", and the second that the thick part has bounded geometry, is relatively hyperbolic and satisfies certain isoperimetric bounds. In practice, models are constructed by finding first some thick part, and then gluing in model Margulis regions to complete the picture. If this is done sensibly, then the key point is that the thick part is relatively hyperbolic.

The notion of relative hyperbolicity was defined by Gromov [Grom]. A number of accounts have been given since, in different contexts. See for example [Fa,Bow3] and the references therein.

In the following discussion, we will use the fact that hyperbolicity is essentially a local property.

Proposition 2.1 :
(1) $(\forall k)(\exists k')(\forall r)$ *any metric r-ball in a k-hyperbolic spaces is intrinsically k'-hyperbolic.*
(2) $(\forall k)(\exists r, k')$ *such that if H is simply connected space and each metric r-ball is intrinsically k-hyperbolic, then H is k'-hyperbolic.*

This general result was known to Gromov [Grom]. Part (1) is elementary. Part (2), in this form, is a consequence of the statement in [Bow1], for example. In fact, r and k' can be taken to be fixed multiples of k.

Since we will be using a variation later, we recall the idea of the proof.

We use a version of the linear isoperimetric inequality. Suppose γ is a closed curve in H. We choose some set, A, of n equally spaced points around γ, where $n = [\frac{1}{k} \operatorname{length} \gamma] + 1$. A *spanning disc*, (Δ, f), consists of a cellulation, Δ, of the disc, D, and a map $f : \Delta^1 \longrightarrow G$, where Δ^1 is the 1-skeleton of Δ, with $f|\partial D = \gamma$, and $f(V_\partial(\Delta)) = A$, where $V_\partial(\Delta)$ is the set of 0-cells lying in ∂D. Its *mesh* is the maximal length of $f(\partial c)$ as c varies over the set of 2-cells of Δ. We fix some $\mu \geq 0$, sufficiently large in relation to the hyperbolicity constant, and choose f to be "minimal" among all spanning discs of mesh at most μ, that is so that $f(\Delta^1)$ had minimal length. (Such must exist by simple connectivity.) If μ is chosen large enough, one verifies from the local hyperbolicity assumption that any 2-cell, c, with at most 13 edges satisfies $\operatorname{length} f(\partial c) < \mu/2$. ("Interior" means means that $c \cap \partial D = \emptyset$.) One deduces that any two interior 2-cells with at most 13 edges must be disjoint. A simple combinatorial argument now shows that the total number of 2-cells is linearly bounded in terms of n, and hence $\operatorname{length}(\gamma)$. It follows that H is globally k'-hyperbolic, where k' depends only on k. There are some technical qualifications to the above construction in order for it to work smoothly, though these play no essential role.

These statements give rise to a coarse version of hyperbolic Dehn surgery as we now discuss. (For related results in the context of hyperbolic groups, see [Os,GrovM].)

We recall the following construction of Gromov. Suppose that P is a any riemannian manifold, and write $g(r)dr$ for the infinitesimal distance. We define a riemannian metric on $F \times [0, \infty)$, given infinisesimally as ds, by $ds^2 = g(r)^2 dr^2 + e^{-2t} dt^2$, where t is arc length in $[0, \infty)$. We denote the resulting manifold by $B(P)$. This is complete and uniformly hyperbolic. We write $B_t(P) = P \times [t, \infty) \subseteq B(P)$. This is convex on $B(P)$. If P is euclidean, then $B(P)$ is (constant curvature) hyperbolic.

If R is a riemannian manifold with boundary ∂R. We write $C(R)$ for the manifolds obtained by gluing a cusp, $B(P)$, to each component, P, of ∂R. (We can smooth out in a small neighbourhood of P.)

Definition : We say that R is *relatively k-hyperbolic* if $C(R)$ is k-hyperbolic.

Note that, by Lemma 1.2.1, there is some $t > 0$, depending only on k such that $B_t(P)$ is convex in R for each component P of ∂R. Also $B(P)$ is itself uniformly quasiconvex.

A related construction, when each P is a euclidean torus, is that of hyperbolic Dehn filling.

Let Δ be a euclidean torus. Given a preferred class of meridean, m, on Δ, there is a unique standard Margulis tube, $T(\Delta) = T(\Delta, m)$ with $\partial T(\Delta) = \Delta$, such that m is homotopically trivial. If δ is the core geodesic we write $s = s(\Delta) = d_T(\partial T, \delta)$ for the *depth* of T. Thus, the length of the euclidean realisation of the meridian in ∂T is exactly $2\pi \sinh s$ (that is, the length of a hyperbolic circle of radius s). If we change m, so that s increases, T looks more and more like $B(\Delta)$, in the sense of geometric convergence. More precisely, given any $t \geq 0$ and $\mu > 1$, there is some $s \geq 0$ such that if $s(\Delta) \geq s$, then there is a μ-bilipshitz homeomorphism from the t-neighbourhood of ∂T in $T(\Delta)$ to that in $B(\Delta)$. Note that $B(\Delta)$ is a Margulis $\mathbf{Z} \oplus \mathbf{Z}$-cusp.

Suppose that R is a complete riemannian manifold such that each component of

∂R is a euclidean torus equipped with a preferred meridean. We construct a manifold $F(R)$ by gluing in a standard Margulis tube to each boundary component (smoothing in a neighbourhood of the boundary, if we want). Let $s(R)$ be the infimum of the depths of these tubes. We write $\tilde{C}(R)$ and $\tilde{F}(R)$ for the universal covers of $C(R)$ and $F(R)$ respectively.

Lemma 2.2 : $(\forall k, l)(\exists s)$ *with the following property. Suppose that R is a riemannian manifold such that each component of ∂R is a euclidean torus with a preferred meridian, and suppose that $s(R) \geq s$. Suppose that $\tilde{C}(R)$ is k-hyperbolic and that γ is a curve in R of length at most l that is homotopically trivial in $F(R)$. Then γ is homotoplically trivial in R.*

Proof : We use a variation on the argument in [Bow1] outlined above. We omit some of the technical details from our account here. One can fill these in, referring back to the original.

Let $\hat{R} \subseteq \tilde{F}(R)$ be the preimage of R under the covering map $\tilde{F}(R) \longrightarrow R$. We thus have regular coverings $\tilde{R} \longrightarrow \hat{R} \longrightarrow R$. For the moment, we consider any curve, γ, in R that is homotopically trivial in $F(R)$. This lifts to a closed curve, $\hat{\gamma}$, in $\hat{R}$. There is a "spanning disc", (Δ, f) for $\hat{\gamma}$, in $\hat{R}$. That is, a map $f : \Delta^1 \longrightarrow \hat{R}$ from the 1-skeleton of a cellulation, Δ, of the disc, D, with $f|\partial D = \hat{\gamma}$. We insist that $f(V_\partial(\Delta))$ is a fixed set of l equally spaced points around γ, where $n = [\frac{1}{k} \operatorname{length}(\gamma)] + 1$. Moreover, if c is a 2-cell of Δ then either length $f(\partial c) \leq \mu$ and ∂c is homotopically trivial in $\hat{R}$, or else $f(\partial c) \subseteq \partial\hat{R}$. Here, $\mu \geq 0$ is a constant fixed sufficiently large in relation to k. We refer to a 2-cell of the second type as *meridional*. The existence of such a spanning disc is a simple consequence of the fact that $\hat{\gamma}$ is homotopically trivial in $\tilde{f}(R)$.

Among all such spanning discs we choose one, (Δ, f), that is "minimal" in the following sense. First, we minimise the number of meridional 2-cells. Among these, we minimise the total length of $f(\Delta^1)$. Finally, among these, we minimise the total number of 2-cells. (The last is really just a technical requirement, as used in [Bow1].) Provided $s(R)$ is large enough in relation to k, such a minimal spanning disc satisfies the same combinatorial condition referred to above, namely any two interior 2-cells, each with at most 13 edges must be disjoint.

To see this, first note that if c is meridional 2-cell, then $f(\partial c)$ is homotopically non-trivial in the component, P, of $\partial\hat{R}$ containing $f(\partial c)$. (Otherwise, we can push it off P, and then subdivide into 2-cells of length at most μ, thereby reducing the number of meridional discs.) Recall that P is a euclidean cylinder whose width can be made arbitrarily large by choosing $s(R)$ large enough. In particular, we can assume that any essential curve in P has length at least $p\mu$, for some fixed $p \geq 14$. All the adjacent 2-cells are non-meridional and so have length at most μ. Thus c has at least $p \geq 14$ edges.

One now shows that any non-meridional 2-cell with at most 13 edges has length less than $\mu/2$ (where we have chosen μ large enough in relation to k). The geometric argument is the same as that in [Bow1], except that instead of knowing that $\hat{R}$ is locally hyperbolic, we use the fact that ∂c is homotopically trivial in $\hat{R}$, and so lifts to $\tilde{R} \subseteq \tilde{C}(R)$ and then apply hyperbolicity of $\tilde{C}(R)$.

It now follows that any two interior 2-cells of Δ with at most 13 edges are disjoint. (Otherwise we could reduce the length of $f(\Delta^1)$ by eliminating a common edge.) We conclude (by Lemma 8.4.1 of [Bow1]) that the number of 2-cells of Δ is linearly bounded above in terms of n (hence length(γ)).

In particular, if length(γ) $\leq l$, this places a bound, q, on the number of 2-cells depending only on r and k. But we are free in the above argument to choose s sufficiently large so that any meridional 2-cell has at least p edges, where we can take $p \geq q$. It will then follow that there are, in fact, no meridional 2-cells at all. Thus $\hat{\gamma}$ is homotopically trivial in R. (We need some additional information, namely that we can assume that the 1-skeleton of a minimal spanning disc is 3-edge connected, as discussed in [Bow1].) ◇

Proposition 2.3 : *Suppose that R is a riemannian manifold such that each component of ∂R is a euclidean torus with a preferred meridian.*

(1) $(\forall k)(\exists s, k')$ such that if $\tilde{C}(R)$ is k-hyperbolic and $s(R) \geq s$, then $\tilde{F}(R)$ is k'-hyperbolic.

(2) $(\forall k, l)(\exists s, k')$ such that if $\tilde{F}(R)$ is k-hyperbolic, $s(R) \geq s$, and each curve in R of length at most l that is homotopically trivial in $F(R)$ is also homotopically trivial in R, then $\tilde{C}(R)$ is k'-hyperbolic.

Proof : (Sketch)

(1) One can elaborate on the argument of Lemma 2.2 to obtain a linear isoperimetric inequality directly. Alternatively, we can apply Propostion 2.1 as follows.

Let $\hat{R}$ be the preimage of R in $\tilde{F}(R)$. We write $\tilde{R}_t$ and $\hat{R}_t$ respectively for the t-neighbourhoods of $\hat{R}$ in $\tilde{C}(R)$ and $\hat{R}$ in $\tilde{F}(R)$. If $s(R) > t$, there is a natural covering map $\tilde{R}_t \longrightarrow \hat{R}_t$ extending $\tilde{R} \longrightarrow \hat{R}$, defined by sending geodesics in $\tilde{R}_t$ perpendicular to $\tilde{R}$ to geodesics in $\hat{R}_t$ perpendicular to $\hat{R}$. Moreover, by taking $s(R)$ sufficiently large in relation to t, this can be assumed arbitrarily close to an isometry (that is λ-bilipschitz for λ arbitrarily close to 1).

By Proposition 2.1(1), there is some k'', depending only k, so that for all $r \geq 0$, every metric r-ball in $\tilde{C}(R)$ is k''-hyperbolic. We fix $r \geq 0$ sufficiently large in relation to k, as determined below. Let N is a metric r-ball in $\tilde{F}(R)$. If $N \cap \hat{R} = \emptyset$, then N is isometric to an r-ball in $\mathbf{H}^3$, hence uniformly hyperbolic. Suppose that $N \cap \hat{R} \neq \emptyset$, so that $N \subseteq \hat{R}_{2r}$. We can suppose that $\tilde{R}_{2r} \longrightarrow \hat{R}_{2r}$ is a covering map. Now applying Lemma 2.2, we can take $s(R)$ large enough so that the mimimal displacements of the covering translations of $\tilde{R} \longrightarrow \hat{R}$ are as large as we want in relation to r. In particular, we can assume that the minimal displacement of the covering translations of $\tilde{R}_{2r} \longrightarrow \hat{R}_{2r}$ is at least $3r$, say. Let N' be a component of the preimage of N in $\tilde{R}_{2r}$. This is close to an r-ball in $\hat{R}_{2r}$, and also the map $N' \longrightarrow N$ is close to an isometry. We can thus assume that N is k'''-hyperbolic, where k''' depends only on k'', and hence ultimately only on k. In other words every r-ball in $\tilde{F}(R)$ is k'''-hyperbolic. Thus, by Proposition 2.1(2), if we chose r sufficiently large in relation to k, it follows that $\tilde{F}(R)$ is k'-hyperbolic, where k' depends only on k.

(2) As above, except that we no longer have the equivalent of Lemma 2.2, so we have taken this as a hypothesis. ◇

As noted earlier, the main point of this section is to note that we can effectively separate out the assumptions on M to hypotheses on the thick and thin parts. We assume that $\Theta(M)$ has bounded geometry and systole bounded below. We assume that its universal cover, $\tilde{\Theta}(M)$ is intrinsically relatively hyperbolic, and satisfies isoperimetric inequalities of type (M4) or (M4′). We assume that each component of the thin part, $\Upsilon(M)$, is a standard Margulis region and that each Margulis tube is sufficiently deep in relation to the Margulis constant (so that we can apply Propostion 2.3).

3. Tameness.

In this section, we elaborate on the notion of a simply degenerate end, and show that such an end is topologically finite.

Let M be a coarse hyperbolic manifold. Recall that $\Psi(M,\eta)$ and $\Theta(M,\eta)$ are the η-non-cuspidal and η-thick parts of M, respectively. We will eventually settle on the parameters of M, though intermediate arguments may take us through a sequence of thick and non-cuspidal parts, getting larger and larger. In fact, it will be convenient to allow for different constants to define different Margulis regions, provided these are kept within certain bounds. More precisely, we will have $\Theta(M) \subseteq \Psi(M) \subseteq M$, with $\Theta(M,\eta') \subseteq \Theta(M) \subseteq \Theta(M,\eta)$ and $\Psi(M,\eta') \subseteq \Psi(M) \subseteq \Psi(M,\eta)$, where $\eta' \leq \eta$ depend only on the parameters of M. Certainly we will want that all Margulis regions are standard, and that all their lifts to $\tilde{M}$ are convex. This can be arranged by the results of Section 1. More conditions will be imposed later. Note that the gaps between different Margulis regions with different constants are very simple geometrically, and so it is easy to control the various constructions during these adjustments.

We recall the notion of a reduced metric as used in [Bow5]. (This is implicit in the work of [Bon], though not formulated in this way.) This is a riemannian pseudometric, equal to d on $\Theta(M)$ and 0 on each Margulis tube. In other words, given any path π in $\Psi(M)$, write $\rho - \mathrm{length}(\pi) = \mathrm{length}(\pi \cap \Theta(M))$. Given $x, y \in \Psi(M)$, set $\rho(x,y)$ to be the minimal ρ-length of any path connecting x to y in $\Psi(M)$. (This is attained.) Since Margulis tubes are compact, ρ is proper. That is, any closed ρ-bounded subset of $\Psi(M)$ is compact. In particular, if $(x_i)_i$ is any sequence going out an end of $\Psi(M)$, then $\rho(x_0, x_i) \to \infty$.

There is a preferred relative homotopy class maps $(\Sigma, \partial\Sigma) \longrightarrow (\Psi(M), \partial\Psi(M))$. Given a non-trivial non-peripheral closed curve in γ in Σ, we can realise it as a curve, γ^*, of minimal length in M. We observed in Section 1 that γ^* is well defined up to bounded distance in M. (Indeed so it its lift to $\tilde{M}$.) These bounds hold with respect to the original metric, d.

Lemma 3.1 : *Suppose γ_1, γ_2 are freely homotopic curves in $\Psi(M)$ each of length at most $\rho \geq 0$. Then ρ-diam$(\gamma_1 \cup \gamma_2)$ is bounded above in terms of r and the parameters of M.*

Proof : It is enough to show that ρ-diam$(\gamma_1 \cup (\gamma^* \cap \Psi(M)))$ is bounded, where γ^* a shortest curve in the homotopy class. Consider the lifts $\tilde{\gamma}_1$ and $\tilde{\gamma}^*$ to $\tilde{M}$. If the stable length is bounded below (in relation to the Margulis constant) then $\tilde{\gamma}_1$ and $\tilde{\gamma}^*$ are a bounded Hausdorff distance apart. If the stable length is small, then $\tilde{\gamma}^*$ will be the axis of a Margulis region, and $\tilde{\gamma}_1$ will be a bounded distance away from this region. Either way, ρ-diam$(\gamma_1 \cup (\gamma^* \cap \Psi(M))$ is bounded above. $\diamondsuit$

The following requires some more involved machinery, as developed in [Bow6].

Lemma 3.2 : *If $\gamma \in X(\Sigma)$, then $\gamma^* \subseteq \Psi(M)$, and ρ-diam$(\gamma^*) \leq r_0$. Moreover, there is a relative homotopy equivalence, $f : (\Sigma, \partial\Sigma) \longrightarrow (\Psi(M), \partial\Psi(M))$ with ρ-diam$(\gamma^* \cup f(\Sigma)) \leq r_0$. Here r_0 depends only on the parameters of M.*

The above involve enlarging the original non-cuspidal part, $\Psi(M)$, but only by an amount controlled in terms of the parameters.

Proof : The proof uses the result from [Bow6] that γ^* remains close to a track realised in M. Here we give a simplified description, only relating facts relevant to the proof.

We can think of a track as an embedded graph $\tau \subseteq \Sigma$ with $\partial\Sigma \subseteq \tau$, and with each component of $\Sigma \setminus \tau$ a topological disc whose closure meets at most one component of $\partial\Sigma$. The total number of edges of τ is bounded above in terms of type(Σ). A realisation of τ consists of a map $f : \tau \longrightarrow M$ which extends to a map $f : \Sigma \longrightarrow M$ in the preferred homotopy class. (Thus, $f(\partial\Sigma)$ is homotopic to $\partial\Psi(M)$, though we don't assume for the moment that $f(\Sigma) \subseteq \Psi(M)$.)

The track, τ, comes together with a (possible empty) set, $\mathcal{S}$, of disjoint circuits embedded in τ. Each component of $\partial\Sigma$ lies in Σ, and we write $\mathcal{S}_\partial \subseteq \mathcal{S}$ for these components. Each $\delta \in \mathcal{S}$ comes with a regular annular neighbourhood $A(\delta) \subseteq \Sigma$, such that $\tau \cap A(\delta)$ consists of δ together with an interval in each adjacent edge. These annuli are disjoint, and we write $A(\mathcal{S}) = \bigcup_{\delta \in \mathcal{S}} A(\delta)$.

Given $\gamma \in X(\Sigma)$, the main result of [Bow6] gives us a track $\tau \subseteq \Sigma$, a set, $\mathcal{S}$, of circuits in τ, and a system, $(A(\delta))_{\delta \in \mathcal{S}}$ of annuli, and a realisation $f : \tau \longrightarrow M$, extending to $f : \Sigma \longrightarrow M$, with the following properties.

(1) $f(A(\mathcal{S})) \cap \Theta(M) = \emptyset$. This means that for each $\delta \in \mathcal{S}$, $f(A(\delta))$ lies inside a corresponding Margulis region.

(2) The total length of $f(\delta \setminus A(\mathcal{S}))$ is bounded.

(3) The total length of $f(\partial A(\mathcal{S}))$ is bounded.

(4) There is a representative $\gamma^\wedge$ of γ in $\tau \setminus \partial\Sigma \subseteq \Sigma$ such that $f(\gamma^\wedge)$ is a bounded distance from γ^*. (Indeed their lifts to $\tilde{M}$ are Hausdorff close.) Note that, necessarily, $\gamma^n \cap A(\delta) = \emptyset$ for all $\delta \in \mathcal{S}_\partial$.

Here all the bounds depend only on the parameters of M.

These properties are stated somewhat differently in [Bow6], though they are simple consequences of the main result there.

A simple consequence of the above is that $f(\delta \setminus A(\mathcal{S}))$ cannot penetrate too deeply into any Margulis region. The same therefore goes for $f(\gamma^\wedge)$, and hence for γ^*. Thus, after shrinking the Margulis tubes a controlled amount, we can assume that $\gamma^* \subseteq \Psi(M)$.

Suppose that $\delta \in \mathcal{S}_\partial$, and let δ' be the other boundary curve of $A(\delta)$. Since $f(\delta)$ does not penetrate too far into the corresponding **Z**-cusp, C, and has bounded length, it must lie a bounded distance from a closed euclidean geodesic δ''. This may involve modifying f on adjacent edges of τ, but only a bounded amount. Moreover, any other edge of τ can only penetrate a bounded distance into C, so we can project outward, increasing its length a bounded amount. Thus, we can assume that $f(\tau) \subseteq \Psi(M)$. We can now homotope $f(\Sigma)$ so that $f(\Sigma) \subseteq \Psi(M)$, and $f : (\Sigma, \partial\Sigma) \longrightarrow (\Psi(M), \partial\Psi(M))$ is a relative homotopy equivalence in the preferred class.

We know that the total length of $f(\tau \cap \Theta(M))$ is bounded. It follows easily that ρ-diam(γ^*) is bounded.

It remains to modify f to bound ρ-diam$(f(\Sigma))$. We shall keep f fixed on $\tau \cup A(\mathcal{S})$.

Let D be a component of $\Sigma \setminus (\delta \cup A(\mathcal{S}))$. Thus, D is a disc, and $\partial D \subseteq \tau \cup \partial A(\mathcal{S})$. But the length of $f(\tau \cup \partial A(\mathcal{S}))$, and hence of $f(\partial D)$ is bounded. Using (M4) we can extend $f|\partial D$ to a map $f : D \longrightarrow M$ with diam$(f(D))$ bounded. Since it penetrates each **Z**-cusp a bounded amount, we can project it back into $\Psi(M)$ increasing the diameter at most a bounded amount. In other words, we can assume that $f(D) \subseteq \Psi(M)$.

Performing this construction for each such component, we obtain a map $f : \Sigma \longrightarrow \Psi(M)$ with $f(\partial\Sigma) \subseteq \partial\Psi(M)$, and with diam $f(R)$ bounded for each component, R, of $\Sigma \setminus A(\mathcal{S})$.

Now $f(A(\delta))$ lies in a Margulis tube for each $\delta \in \mathcal{S} \setminus \mathcal{S}_\partial$. It then follows that ρ-diam $f(\Sigma)$ is bounded. Since we have not changed $f(\tau)$ (other than the earlier modification involving **Z**-cusps) we see that ρ-diam$(\gamma^* \cup f(\Sigma))$ is bounded as required. ◊

Another consequence of this construction is as follows.

Lemma 3.3 : *There are constants, r_0, l_0, depending only on the parameters of M such that if $\gamma \in X(\Sigma)$, then there is some $\beta \in X(M, l_0)$ such that ρ-diam$(\gamma^* \cup \beta^*)$ is bounded.*

Proof : Suppose first that $\mathcal{S} = \mathcal{S}_\partial$. In this case, the d-diameter of $f(\tau)$ is bounded. Moreover, since each component of $\Sigma \setminus \tau$ is a disc meeting only one component of $\partial\Sigma$, we can find a closed path β in $\tau \setminus \partial\Sigma$, meeting any edge at most twice, and which is homotopic to an essential non-peripheral curve in Σ. In other words, $\beta \in X(\Sigma)$. Now, since length$(f(\beta))$ is bounded, it follows by Lemma 3.1, that ρ-diam$(\gamma^* \cup \beta^*)$ is bounded as required.

If $\mathcal{S} \neq \mathcal{S}_\partial$, we can take β to be any element of $\mathcal{S} \setminus \mathcal{S}_\partial$. ◊

Remark : In fact, the construction of quasiprojection described in [Bow7] gives us such a β obtained by shortcutting γ in a particular way. In this way, one can arrange, in addition, that $d_{\mathcal{G}(\Sigma)}(\beta, \gamma)$ is bounded in terms of the parameters. This allows us to strengthen certain statements made later, but is not essential to the main results stated in this paper.

Lemma 3.4 : *Suppose that e is an end of $\Psi(M)$ and there is a sequence of curves, $\gamma_i \in X(\Sigma)$ and $x_i \in \gamma_i^*$ with $x_i \to e$, there is another sequence, β_i in $X(M, l_0)$ with $\beta_i^* \to e$, where l_0 depends only on the parameters of M.* ◊

Proof : These statements follow directly from Lemmas 3.2 and 3.3 respectively, and the fact that the pseudometric ρ is proper. $\diamondsuit$

In particular, it follows that e is simply degenerate as we have defined it, justifying a statement made in Section 1.

Proof of Theorem 0.2 : Let e be a simply degenerate end of $\Psi(M)$. We want to show that e is topologically finite. By standard topology (cf. [T,Bon]), it is enough to find a sequence, f_i of relative homotopy equivalences, $f_i : (\Sigma, \partial\Sigma) \longrightarrow (\Psi(M), \partial\Psi(M))$, with $f_i(\Sigma) \to e$.

For this, we take any sequence of $\gamma_i \in X(\Sigma)$ with $\gamma_i^* \to e$ as given by the definition of "simply degenerate". Lemma 3.2 gives us a relative homotopy equivalence, $f_i : (\Sigma, \partial\Sigma) \longrightarrow (\Psi(M), \partial\Psi(M))$ with ρ-diam$(\gamma^* \cup f_i(\sigma))$ bounded. It follows that $f_i(\Sigma) \to e$ as required. $\diamondsuit$

To go further, we need another consequence of the construction in Lemma 3.2.

Lemma 3.5 : *Suppose $\beta, \gamma \in X(\Sigma)$ are adjacent in $\mathcal{G}(\Sigma)$, then ρ-*diam$(\beta^* \cup \gamma^*) \le r_0$*, where r_0 depends only on the parameters of M.*

Proof : In fact, the result of [Bow6] applies to any multicurve in Σ in particular to $\beta \cup \gamma$. Then $\beta^* \cup \gamma^*$ is close to $f(\tau)$. Thus ρ-diam$(\beta^* \cup \gamma^* \cup f(\tau))$ is bounded. $\diamondsuit$

It immediately follows that for any $\beta, \gamma \in X(\Sigma)$, ρ-diam$(\beta^* \cup \gamma^*) \le r_0 d_{\mathcal{G}(\Sigma)}(\beta, \gamma)$. This in turn implies:

Lemma 3.6 : *Suppose that e is an end of M and that $(\gamma_i)_i$ is a sequence in $X(\Sigma)$ with $\gamma_i^* \to e$, then $d_{\mathcal{G}(\Sigma)}(\gamma_0, \gamma_i) \to \infty$.* $\diamondsuit$

Recall that $\mathcal{G}(\Sigma)$ is Gromov hyperbolic with boundary, $\partial\mathcal{G}(\Sigma)$.

Given a simply degenerate end, e, of M, we write $A_l(e)$ for the set of $a \in M$ such that there is and some sequence $\gamma_i \in X(M, l)$ with $\gamma_i^* \to e$ in M and $\gamma_i \to a$ in $\mathcal{G}(\Sigma) \cup \partial\mathcal{G}(\Sigma)$.

Lemma 3.7 : $A_{l_0}(e) \neq \emptyset$.

Proof : This follows from Lemmas 3.4 and 3.6, together with the general fact that an unbounded sequence in any Gromov hyperbolic space has a subsequence that converges to a boundary point. $\diamondsuit$

We will see in Section 4, that $A_l(e)$ for all independent of l for all sufficiently large l in relation to the parameters. Indeed, one could take any sequence γ_i with $\gamma_i^* \to e$, without constraint on the length, though this requires more work.

We will also eventually see, in Section 6, that $A_l(e)$ consists of a single point, giving rise to the notion of "end invariant".

As noted in Section 1, the statement that an end, e, of $\Psi(M)$ is not geometrically finite means that there is a sequence $(\gamma_i)_i$ of non-trivial non-peripheral closed curves in

Σ, and points $x_i \in \gamma_i^*$ with $x_i \to e$. Here, the γ_i are not assumed simple, and it is no longer necessarily the case that $\gamma_i^* \subseteq \Psi(M)$, nor that the ρ-diameters of $\gamma_i^* \cap \Psi(M)$ are bounded. One can ask whether it is possible to replace the γ_i by a sequence of simple closed curves, which would then imply that e is simply degenerate. This is achieved in [Bon] in the constant curvature case, and it is noted in [C] that the argument also works for pinched negative curvature. The case where $\tilde{M}$ is only assumed hyperbolic is discussed in [Ba], though under the assumption that the systole is bounded below. In fact, the author assumes that M covers a compact 3-manifold with hyperbolic fundamental group, though it seems likely that some form of bounded geometry of M is all that is required of this last assumption.

4. End invariants.

In this section, we continue the discussion of invariants of simply degenerate ends. We need to bring the "a-priori bounds" results of [Bow7] into play. These generalise the statements of [Mi]. We have already verified that the relevant hypotheses hold. Namely, the locally bounded geometry condition used there is given by Lemma 1.1.1, and the thick-thin decomposition by Lemma 1.3.1.

One immediate consequence is the following, stated as Theorem 0.3 in the introduction.

Theorem 4.1 : *There is some $l_0 \geq 0$ such that $X(M, l_0)$ is non-empty and uniformly quasiconvex in $\mathcal{G}(\Sigma)$. Moreover, given any $l \geq 0$, there is some $r \geq 0$ depending only on l and the parameters, such that $X(M, l) \subseteq M(X(M, l_0), r)$.* ◊

Throughout this section, "uniform" means depending only on the parameters of M.

We may as well take l_0 in the above to be the same constant as that featuring in Section 3. One immediate consequence of Theorem 4.1 is that if e is a simply degenerate end of $\Psi(M)$, then $A_l(e) = A_{l_0}(e)$ for all $l \geq l_0$. We shall therefore abbreviate $A(e) = A_{l_0}(e)$.

To go further, we need the notion of a "tight geodesic" in $\mathcal{G}(\Sigma)$. This was introduced in [MaM2]. Here we refer to the slightly more general definition used in [Bow7] for the purposes of quoting results.

Theorem 4.2 : *Suppose that $\gamma_0\gamma_1 \cdots \gamma_n$ is a tight geodesic in $\mathcal{G}(\Sigma)$ with $\gamma_0, \gamma_p \in X(M, l)$. Then $\gamma_i \in X(M, l')$ for all $i \in \{0, \ldots, n\}$, where l' depends on l and the parameters of the action.* ◊

Theorem 4.3 : *There is some r_0 depending only or the parameters of the action such that $\gamma_0, \ldots, \gamma_p$ is a tight geodesic in $\mathcal{G}(\Sigma)$ and $r + r_0 \leq i \leq p - r - r_0$, where $r = \max\{d(\gamma_0, X(M, l)), d(\gamma_p, X(M, l))\}$, then $\gamma_i \in X(M, l')$, where l' depends only on l and the parameters of M.* ◊

Let e be a simply degenerate end of $\Psi(M)$.

Lemma 4.4 : *Suppose $a \in A(e)$ and $l \geq 0$. Suppose that $(\gamma_i)_i$ is a sequence of curves in $X(M, l)$ with $\gamma_i \to a$. Then $\gamma_i^* \to e$.*

Proof : Let e' be the other end of $\Psi(M)$ (which we need not assume to be simpy degenerate). Let E, E' be neighbourhoods of e, e' respectively with $\rho(E, E') \geq r_0$, where r_0 is the constant of Lemma 3.3. Let K be the closure of $\Psi(M) \setminus (E \cup E')$. Thus, K is compact. Note that if $\alpha, \beta \in X(\Sigma)$, with $\alpha^* \subseteq E'$ and $\beta^* \subseteq E$, and $(\epsilon_i)_i$ is any path in $\mathcal{G}(\Sigma)$ from α to β, then $\epsilon_i^* \cap K \neq \emptyset$ for some i.

Suppose, for contradiction that $\gamma_i^* \not\to e$. Now γ_i is unbounded in $\mathcal{G}(\Sigma)$. Since the γ_i^* have bounded length, they must escape any compact subset of M. Passing to a subseqence, we can therefore assume that $\gamma_i^* \to e'$.

By assumption, there is a sequence $(\delta_j)_j$ in $X(M, l_0)$ with $\delta_j \to a$ in $\mathcal{G}(\Sigma)$ and with $\delta_j^* \to e$ in $\Psi(M)$. For any given j, let $(\beta_{ij})_i$ be a tight geodesic from γ_j to δ_j. By Theorem 4.2, there is some $l \geq 0$ such that $\beta_{ij} \in X(M, l)$ for all i, j. Moreover, by the earlier observation, for any j, there is some i_j such that $\alpha_j \cap K \neq \emptyset$, where $\alpha_j = \beta_{i_j j}$. Now the α_j^* have bounded length, and remain in a compact region of M, and so lie in finitely many homotopy classes. In other words, $\{\alpha_j | j \in \mathbf{N}\}$ is finite.

In summary, we have $\gamma_j \to a$, $\delta_j \to a$, and for each j, we have a geodesic from γ_j to δ_j which meets a fixed finite subset of $\mathcal{G}(\Sigma)$. This is easily seen to contradict the hyperbolicity of $\mathcal{G}(\Sigma)$. $\diamondsuit$

Lemma 4.5 : *For any $a \in A(e)$ there is a tight geodesic ray, $(\gamma_i)_{i=1}^{\infty}$ in $\mathcal{G}(\Sigma)$ converging to a with $\gamma_i \in X(M, l_0)$ for all i, where l_0 depends only on the parameters of M.*

Proof : From the definition of $A(e)$, there is sequence $(\delta_i)_i$ in $X(M, l_0)$ with $\delta_i \to a$ and $\delta_i^* \to e$. For each i, let $(\beta_{ij})_{j=1}^{n_i}$ be a tight geodesic from δ_0 to δ_i. Thus, $n_i \to \infty$, and by Proposition 4.2, $\beta_{ij} \in X(M, l')$ for some uniform $l' \geq 0$. Given any j, we have $d_{\mathcal{G}(\Sigma)}(\delta_0, \beta_{ij}) = i$ (for all sufficiently large i so that β_{ij} is defined) and so, by Lemma 3.3, $\rho(\delta_0^*, \beta_{ij}^*)$ is bounded in terms of i. Thus, the curves β_{ij}^* all lie in a compact region of M. Since their lengths are bounded, they lie in finitely many homotopy classes. That is, $\{\beta_{ij} | j \geq 0\}$ is finite. Passing to a diagonal subsequence, we can find a sequence $(\gamma_i)_i$ so that for all i, $\beta_{ij} = \gamma_i$ for all sufficiently large j. Now $(\gamma_i)_i$ is a tight geodesic converging to a with $\gamma_i \in X(M, l')$ for all i. To simplify notation, we can reset $l_0 = l'$. $\diamondsuit$

Lemma 4.6 : *Suppose $(\gamma_i)_i$ is any tight geodesic converging on some $a \in A(e)$. Then for all sufficiently large i, we have $\gamma_i \in X(M, l_0)$ for some uniform l_0.*

Proof : Let $(\beta_i)_i$ be the geodesic given by Lemma 4.5, so that $\beta_i \in X(M, l_0)$. For for all sufficiently large i, $(\gamma_i)_i$ lies uniformly close to $(\beta_i)_i$ (after shifting indices), and so it follows using Theorem 4.3 that $\gamma_i \in X(M, l')$ for some uniform l'. Reset $l_0 = l'$. $\diamondsuit$

The following result will turn out to be vacuous once we know that $A(e)$ is a singleton. However it is needed (or more precisely the variation Lemma 4.12) to prove this fact.

Lemma 4.7 : *Suppose $a, b \in A(e)$ with $a \neq b$. Let $(\gamma_i)_i$ be any bi-infinite geodesic from a to b. Then $\gamma_i \in X(M, l_0)$ for all i, where $l_0 \geq 0$ is uniform.*

Proof : By Lemma 4.6, we have $l_M(\gamma_i)$ and $l_M(\gamma_{-i})$ uniformly bounded for all sufficiently large i. Now apply Theorem 4.2. $\diamondsuit$

We note that there is always a tight geodesic between any pair of distinct boundary points of $\mathcal{G}(\Sigma)$ (see [MaM2] or [Bow2]).

Suppose now that M is doubly degenerate with ends e_- and e_+.

Lemma 4.8 : $A(e_-) \cap A(e_+) = \emptyset$.

Proof : Suppose $a \in A(e_-) \cap A(e_+)$. Let $(\gamma_i)_i$ be any tight geodesic ray tending to a. By Lemma 4.4, we have $\gamma_i^* \to e_-$ and $\gamma_i^* \to e_+$, giving a contradiction. $\diamondsuit$

The following follows exactly as in Lemma 4.7.

Lemma 4.9 : *Suppose $a \in A(e_-)$ and $b \in A(e_+)$ and that $(\gamma_i)_i$ is a bi-infinite geodesic from a to b. Then $\gamma_i \in X(M, l_0)$ for all i.* $\diamondsuit$

For applications in Section 6, we will also need to consider "hierarchies". All we need to know for this paper is that to any $\alpha, \beta \in X(\Sigma)$ we can associate a canonical finite subset, $\mathcal{H}(\alpha, \beta) \subseteq X(\Sigma)$. Here we shall refer to the description in [Bow7] which is simplified from [MaM2]. (From the definition, it turns out that $\mathcal{H}(\alpha, \beta)$ contains every tight geodesic from α to β, and is contained in a bounded neighbourhood of such a geodesic.)

Hierarchies satisfy a number of finiteness properties. In particular, we have:

Lemma 4.10 : *Suppose α_i, β_i are sequences in $\mathcal{G}$ tending to points of $\partial\mathcal{G}(\Sigma)$, and that $B \subseteq X(\Sigma)$ is bounded. Then there is a finite subset $C \subseteq B$ such that for all sufficiently large i, we have $\mathcal{H}(\alpha_i, \beta_i) \cap B \subseteq C$.* $\diamondsuit$

Thus, if $\alpha_i \to a \in \partial\mathcal{G}(\Sigma)$ and $\beta_i \to b \in \partial\mathcal{G}(\Sigma)$, then we can pass to a diagonal subsequence so that $\mathcal{H}(\alpha_i, \beta_i)$ is eventually constant on any bounded subset of $\mathcal{G}(\Sigma)$. This gives rise to a "hierarchy" between a and b. More formally, we can define this canonically as follows. We set $\mathcal{H}(a, b)$ to be the set of $\gamma \in X(\Sigma)$ such that given and neigbourhoods U and V of a and b respectively in $\mathcal{G}(\Sigma) \cup \partial\mathcal{G}(\Sigma)$, there exist $\alpha \in U \cap X(\Sigma)$ and $\beta \in V \cap X(\Sigma)$ such that $\gamma \in \mathcal{H}(\alpha, \beta)$. Thus, $\mathcal{H}(a, b)$ is a locally finite subset of $\mathcal{G}(\Sigma)$. If $a \neq b$, this is non-empty. (It is a bounded Hausdorff distance from any bi-infinite geodesic from a to b.) In practice, we could constrain α and β to lie in a tight geodesic from a to b, which may simplify some of the arguments, but is not essential.

We quote the following from [Bow7]:

Theorem 4.11 : *Given $l \geq 0$ there is some $l' \geq 0$, and depending only on l and the parameters of M such that if $\alpha, \beta \in X(M, l)$ then $\mathcal{H}(\alpha, \beta) \subseteq X(M, l')$.*

We draw the following conclusions which follow by arguments similar to those of Lemmas 4.7 and 4.9.

Lemma 4.12 : *Suppose $a, b \in A(e)$ with $a \neq b$. Then $\mathcal{H}(a, b) \subseteq X(M, l_0)$.* ♢

Lemma 4.13 : *Suppose that M is doubly degenerate, and $a \in A(e_-)$ and $b \in A(e_+)$. Then $\mathcal{H}(a, b) \subseteq X(M, l_0)$.* ♢

We remark that we can, in fact, eliminate the length bound altogether from the definition of $A(e)$. Indeed, if $(\gamma_i)_i$ is any sequence in $X(\Sigma)$ with $\gamma_i^* \to e$ and $\gamma_i \to a \in \partial\mathcal{G}(\Sigma)$, then $a \in A(e)$. This follows from the remark after Lemma 3.3 — after moving each of the γ_i a bounded distance in $\mathcal{G}(\Sigma)$, we can assume that the $l_M(\gamma_i)$ are all (unformly) bounded. This does not change the limit point of the sequence in $\partial\mathcal{G}(\Sigma)$. We shall not be explicitly using this fact in this paper, though it ties in out notion of "end invariant" with that used in various places elsewhere.

5. Systems of quasiconvex sets.

To prove the main results of this paper, we will follow through the arguments of [Bow5]. Most of this is simple reinterpretation, as we discuss in Section 6. The only point requiring any significant modification is the construction of the lipschitz map from the combinatorial model to the coarse hyperbolic 3-manifold. In [Bow5], this entailed a discussion of systems of convex sets in $\mathbf{H}^3$, here replaced by uniform quasiconvex sets in $H = \tilde{M}$.

First we recall the notion of a truncated simplicial complex. Let Π be a 3-dimensional simplicial complex. We write Π^i for the set of i-simplices. Let $|\Pi|$ be the realisation using regular euclidean simplices all of whose side lengths are 3. The *truncated realisation*, $R(\Pi)$, is the closed subset of $|\Pi|$ obtained by removing a regular simplex of side length 1 about each vertex of each simplex and then taking the closure in $|\Pi|$. Given $x \in \Pi^0$, let $D(x)$ be the boundary of the component of $|\Pi| \setminus R(\Pi)$ containing x. Thus, $D(x)$ is a 2-dimensional simplicial complex. In applications, Π will be locally finite away from Π^0, and so $R(\Pi)$ will be proper.

Let H be a proper k-hyperbolic space. All "quasiconvex" subsets will be assumed to be quasiconvex for some fixed constant. To simplify notation, we may as well take this to be the same as the hyperbolicity constant. Given two such subets, $P, Q \subseteq H$, we set $\mathrm{par}(P, Q) = \mathrm{diam}(N(P, k') \cap N(Q, k'))$, where k' is fixed sufficiently large in relation to the hyperbolicity constant such that for all $r \geq k'$, $\mathrm{diam}(N(P, r) \cap N(Q, r))$ is bounded above in terms of r. In other words, $\mathrm{par}(P, Q)$ serves as a convenient measure of the extent to which P and Q remain "parallel" in H.

Suppose that Π is a 3-dimensional simplicial complex, and that to each $x \in \Pi^0$ we have associated a quasiconvex set $Q(x) \subseteq H$. We assume:

(A1′) If $x, y \in \Pi^0$ are distinct, then $Q(x) \cap Q(y) \subseteq \partial Q(x) \cap \partial Q(y)$ and $\mathrm{par}(Q(x), Q(y))$ is bounded above.

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

(Here (A1′) is a slight weakening of the corresponding statement in [Bow5].)

Given $xy \in \Pi^1$, let $\beta(xy)$ be a shortest geodesic for $Q(x)$ to $Q(y)$ (just a point if they intersect).

Lemma 5.1 : Assuming (A1′) and (B1), if $xyz \in \Pi^2$ then $\mathrm{diam}(\beta(xy) \cup \beta(yz) \cup \beta(zx))$ is bounded.

Proof : The argument given in [Bow5] is easily quasified. ♢

In [Bow5] various constructions that were canonical can only be assumed equivariant here. For this reason, we introduce group actions at this point.

Let Γ be a group acting simplicially on Π and by isometries on H. Given $x \in \Pi^0$, we write $\Gamma(x) = \{g \in \Gamma \mid gx = x\}$. We assume:

(D1) $(\forall x \in \Pi^0)(\forall g \in \Gamma)(Q(gx) = gQ(x))$.

(D2) The setwise stabiliser of each element of Π^1 and of Π^2 is trivial.

(Note that (D2) implies that Γ acts freely and isometrically of $R(\Pi)$.)

We will need to assume that the sets $Q(x)$ are have reasonably nice intrinsic geometry. In practice they arise from loxodomic axes or from Margulis regions. The following, somewhat artificial assumption, will serve for our purposes.

(D3) If $x \in \Pi^0$, then $Q(x)$ is either a uniform quasigeodesic in H, or else convex in H and intrinsically isometric to a convex subset of $\mathbf{H}^3$, whose boundary is the topological boundary of H.

At this point, we need to assume the lipschitz isoperimetric inequalities of (M4) as laid out in the introduction. (These are need to construct our lipschitz maps, though some form of isodiametric inequality would serve if we only wanted quasi-isometric maps.)

Lemma 5.2 : *Under the assumptions (A1′), (B1), (D1)–(D3), there is an equivariant unformly lipschitz map, $\phi : R(\Pi) \longrightarrow H$. 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.*

Proof : The proof follows as in [Bow5]. We replace the coning construction by applications of the isoperimetric inequality (M4). This is no longer canonical. However our assumptions (D1) and (D2) mean that they can be carried out equivariantly. ♢

The next task is to push $\phi(R(\Pi))$ off the interior of the sets $Q(x)$. We only need to worry about those of the second type. For this, we introduce another collection, $(A(x))_{x \in \Pi^0}$ of nonempty quasiconvex sets, with $A(gx) = gA(x)$ for all $x \in \Pi^0$ and $g \in \Gamma$. We assume that for each $x \in \Pi^0$, there is some $t(x) \geq 0$ such that $Q(x) = N(A(x), t(x))$. If $Q(x)$ is of the first type (a quasigeodesic), then $t(x) = 0$ and $Q(x) = A(x)$. Note that for $g \in \Gamma$, $t(gx) = t(x)$. We choose suitable constants $t_0 \leq t_1$ as in [Bow5] and take a Γ-invariant partition of $\Pi^0 = \Pi^0_0 \sqcup \Pi^0_1$ such that $t(x) \leq t_0$ for all $x \in \Pi^0$ and $t(x) \geq t_1$ for all $x \in \Pi^0_0$.

Lemma 5.3 : *With the above hypothesis we can find a Γ-equivariantly uniformly lipschitz map $\phi : R(\Pi) \longrightarrow H$ such that:*

(1) If $x \in \Pi^0_1$, 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 modify the map given by Lemma 5.2, as in [Bow5]. (If $Q(x)$ is of the first type, then no modification is necessary.) $\diamondsuit$

For our application we will not be given (B1) directly, so we need to consider other hypotheses that imply it. At this point we diverge somewhat from the account given in [Bow5] since we no longer have available the involutions used there.

We suppose that the boundary of each of the $Q(x)$ lie in the thick part. More precisely:

(D4) $(\forall r)(\exists \nu)(\forall x \in \Pi^0)(\forall a \in \partial Q(x))$ we have $|\{g \in \Gamma \mid d(a, ga) \leq r\}| \leq \nu$.

Lemma 5.4 : *We assume (A1′), (D1)–(D4). Let $xy \in \Pi^1$. Suppose that $g \in \Gamma(x)$ has infinite order, and set $z = gy$. Then $\mathrm{par}(\beta(xy), \beta(xz))$ is bounded above in terms of the hyperbolicity constant and the parameters of (D4).*

Less formally, this means that either $\beta(xy)$ has bounded length or $\beta(xy)$ and $\beta(xz)$ diverge after a bounded distance.

Proof : Let a be at the intersection of $Q(x)$ and $\beta(xy)$, so $a \in \partial Q(x)$. Suppose $\beta(xy)$ and $\beta(xz) = g\beta(xy)$ remain parallel over a large distance, s. If $\nu \in \mathbf{N}$ and if s is large enough in relation to ν times the hyperbolicity constant, then $d(a, g^i a)$ is bounded by some uniform constant r, for each $i \in \{0, \ldots, \nu\}$. By taking ν large enough in relation to r, we get a contradiction to property (D4). $\diamondsuit$

Next, we suppose we have a Γ-equivariant subset $\Pi^2_0 \subseteq \Pi^2$ satisfying:

(D5) If $xyz \in \Pi^2_0$, then there are infinite order elements $g(x, y, z) \in \Gamma(x)$, $g(y, z, x) \in \Gamma(y)$ and $g(z, x, y) \in \Gamma(z)$ with $g(x, y, z)g(y, z, x)g(z, x, y) = 1$ in Γ.

Lemma 5.5 : *We assume (A1′), (D1)–(D5). If $xyz \in \Pi^2_0$, then the lengths of $\beta(xy)$, $\beta(yz)$ and $\beta(zx)$ are all bounded above in terms of the parameters.*

Proof : It is enough to bound $\beta(xy)$. Let $g = g(x, y, z) \in \Gamma(x)$ and $h = g(y, z, x) = \Gamma(y)$, so that $gh = g(z, x, y)^{-1} \in \Gamma(z)$. Let $w = g^{-1}z = hz$. Since $\Gamma(x) \cap \Gamma(z)$ is trivial, $w \neq z$.

Using (D3) we can find uniformly quasigeodesic paths, $\alpha(x), \alpha(y), \alpha(z), \alpha(w)$ in H, respectively contained in $Q(x), Q(y), Q(z), Q(w)$, so that

$$\alpha(x) \cup \beta(xz) \cup \alpha(z) \cup \beta(zy) \cup \alpha(y) \cup \beta(yw) \cup \alpha(w) \cup \beta(wx)$$

forms a closed path — an "octagon". Consecutive edges of the octagon can remain parallel only over a bounded distance (in terms of the hyperbolicity constant). Moreover, by Lemma 5.4, $\beta(xz)$ and $\beta(xw)$ have the same length, and remain parallel over a bounded distance. Similarly, $\beta(yz)$ and $\beta(yw)$ have the same length, and remain parallel over a bounded distance.

It is now an exercise in hyperbolic spaces to show that if $\alpha(x)$ and $\alpha(y)$ are far apart, then $\alpha(z)$ and $\alpha(w)$ remain close over a large distance. In particular, if $\beta(xy)$ is very long, then $\mathrm{par}(Q(z), Q(w))$ is large, contradicting (A1′).

This shows that $\beta(xy)$ has bounded length, as claimed. $\diamondsuit$

Next we assume:

(D6) If $xyz \in \Pi^2$, at least two of the edges xy, yz, zx lie in simpices of Π^2_0.

(D7) Each edge of Π^1 lies in at least two simplices of Π^2.

Lemma 5.6 : *We assume (A1′), (D1)–(D7). If $xy \in \Pi^1$, then the length of $\beta(xy)$ is bounded above in terms of the parameters of the hypotheses.*

Proof : By (D7) there are distinct $z, w \in \Pi^0$ with $xyz, xyw \in \Pi^2$. Now if $\beta(xy)$ is very long, then by Lemma 5.5, xy cannot lie in any simplex of Π^2_0. Thus, by (D7), xy, yz, xw, yw must all lie in simplices in Π^2_0. By Lemma 5.5 again, the lengths of each of $\beta(xy)$, $\beta(yz)$, $\beta(xw)$ and $\beta(yw)$ are bounded.

As in the proof of Lemma 5.5, we consider the octagon

$$\alpha(x) \cup \beta(xz) \cup \alpha(z) \cup \beta(zy) \cup \alpha(y) \cup \beta(yw) \cup \alpha(w) \cup \beta(wx).$$

This time, we note that all the β-edges have bounded length, and that $\alpha(y)$ and $\alpha(z)$ are far apart. Again by a (simpler) excercise in hyperbolic spaces, $\alpha(z)$ and $\alpha(w)$ remain close over a large distance, giving a contradiction as before. $\diamondsuit$

In other words, we have shown (A1′), (D1)–(D7) imply (B1). In particular, in Lemma 5.3, we can substitute hypothesis (B1) with (D1)–(D7).

We finally come to the application of all this to coarse hyperbolic 3-manifolds. To be clear about our assumptions, we start again at the beginning.

Let M be a coarse hyperbolic 3-manifold, and let $\Gamma = \pi_1(M)$.

Suppose that Γ also acts on a 3-dimensional simplicial complex, Π, and let $R(\Pi)$ and $(D(x))_{x \in \Pi^0}$ be as constructed earlier. Given $x \in \Pi^0$, write $\Gamma(x)$ for the stabiliser of x in Γ. Let Π^2_0 be a Γ-equivariant subset of Π^2.

We suppose:

(E1) Every edge Π is contained in at least two 2-simplices of Π.

(E2) If $x \in \Pi^0$, then $\Gamma(x)$ is infinite cyclic.

(E3) If $x, y \in \Pi^0$ are distinct, then $\Gamma(x) \cap \Gamma(y)$ is trivial.

(E4) The setwise stabiliser of any element of Π^2 is trivial.

(E5) If $xyz \in \Pi^2_0$, we can choose generators, $g(x), g(y), g(z)$ respectively of $\Gamma(x), \Gamma(y), \Gamma(z)$ such that $g(x)g(y)g(z) = 1$.

(E6) If $xyz \in \Pi^2$, then at least two of its edges lie in Π^2_0.

Note that Γ acts freely on $R(\Pi)$. We write $\hat{\Pi} = \Pi/\Gamma$ and $R(\hat{\Pi}) = R(\Pi)/\Gamma$. Given any $x \in \hat{\Pi}^0$, we have a subset $\hat{D}(x) = D(x)/\Gamma(x)$ of $R(\hat{\Pi})$. Note that to each $x \in \hat{\Pi}^0$, we have a free homotopy class of closed curve, $\gamma(x)$, in M and hence also in Σ. (We do not assume, for the moment, that this is simple in Σ.)

We now make the following "a-priori bounds" assumption:

(E7) There is some $L \geq 0$ such that if $x \in \hat{\Pi}^0$, then $l_M(\gamma(x)) \leq L$.

Given any $x \in \hat{\Pi}^0$ we write $\gamma^*(x)$ for a shortest representative in M. If $\eta \leq \eta_1$, we write $T_\eta(x)$ for the η-Margulis tube about $\gamma^*(x)$. In Section 1, we saw that this is a standard tube.

Note that if $y \in \Pi^0$, there are canonical lifts $\tilde{\gamma}^*(x)$ and $\tilde{T}^*_\eta(x)$ to $\tilde{M}$.

Lemma 5.7 : *Let M be a coarse hyperbolic 3-manifold and $\Gamma = \pi_1(M)$. Suppose that Γ acts on a simplicial complex, Π, satisfying (E1)–(E7) above. Then there are positive constants ξ, η_2, η_3 with $\eta_3 \leq \eta_2$, depending only on the parameters of M and of (E7) such that the following holds. We can write Π^0 as a Γ-invariant disjoint union $\Pi^0 = \Pi^0_0 \sqcup \Pi^0_1$, and find a Γ-equivariant ξ-lipschitz map $\phi : R(\Pi) \longrightarrow \Psi(\tilde{M})$ to the lift of the η_2-non-cuspidal part of M, such that:*

(1) If $x \in \Pi^0_1$, then $l_M(\gamma(x)) \leq \eta_3$ and $\phi(D(x)) \subseteq \tilde{\gamma}^*(x)$.

(2) If $x \in \Pi^0_0$, then $l_M(\gamma(x)) \leq \eta_2$ and $\phi(D(x)) \subseteq \partial\tilde{T}_{\eta_2}(x)$.

(3) For all $x \in \Pi^0_0$, $\tilde{T}_{\eta_2}(x) \cap \phi(R(\Pi)) \subseteq \partial\tilde{T}_{\eta_2}(x)$.

Proof : We choose η_2 and η_3 and define Π^0_0 and Π^0_1 similarly as in [Bow5]. If $x \in \Pi^0$, we set $Q(x) = \tilde{\gamma}^*(x)$ if $l_M(\gamma(x)) \geq \eta_2$ and $Q(x) = \tilde{T}_{\eta_2}(x)$ if $l_M(\gamma(x)) < \eta_2$. By the results of Section 1.2, we see that $(Q(x))x \in \Pi^0$ satisfy the conditions laid out above. We now proceed as in [Bow5], applying Lemmas 5.6 and 5.3 to construct a lipschitz map, and then modify inside Margulis regions, either projecting to the boundary or to the axis, depending on the depths of the respective Margulis regions. ◊

Passing to the quotient, we get a map $f : R(\hat{\Pi}) \longrightarrow \Theta(M)$, where $\Theta(M) = \Psi(M) \setminus \bigcup_{x \in \hat{\Pi}^0} \operatorname{int} T_{\eta_2}(x)$.

As in [Bow5] we see:

Lemma 5.8 : *There is some $\eta_4 \geq 0$, depending only on the parameters of M and on the constant L of (D7) such that if $f(r(\hat{\Pi}) = \Theta$, then $\Theta(M, \eta_2) \subseteq \Theta(M) \subseteq \Theta(M, \eta_4)$.* $\diamondsuit$

In particular, it then follows that Θ is a "thick part" of M.

6. A coarse version of the ending lamination conjecture.

We now have most of the machinery in place to give a version of the ending lamination conjecture for coarse hyperbolic 3-manifolds. This will follow the arguments of [Bow5] and we shall give a brief overview of that paper, explaining how the relevant notions are reinterpreted. Here we discuss mainly the doubly degenerate surface group case, though there is scope for applying these ideas to more general manifolds.

We shall use exactly the same "combinatorial model" is in [Bow5], which is, in turn, essentially just a reformulation of that used in [Mi].

Given $a_-, a_+ \in \partial\mathcal{G}(\Sigma)$, with $a_- \neq a_+$, one constructs a "combinatorial model", $P = P(a_-, a_+)$. This is a riemannian manifold, with a preferred "non-cuspidal part", $\Psi(P)$, homeomorphic to $\Sigma \times \mathbf{R}$. Thus $\Psi(P)$ had two ends, e_-^P and e_+^P, which we designate as "negative" and "positive" respectively, and are associated to the "end invariants" a_- and a_+ respectively. Each component of $P \setminus \Psi(P)$ is a standard Margulis $\mathbf{Z}$-cusp, with fixed universal constant. Note that P need not be canonical. We only need that its construction only uses makes use of the data a_- and a_+.

Suppose that M is a coarse hyperbolic 3-manifold (i.e. satisfying (M1)–(M5) of the introduction). We write $\Psi(M)$ for its non-cuspidal part (with fixed constant depending only on the parameters). We designate the ends of M as e_-^M and e_+^M. We aim to prove:

Proposition 6.1 : *Suppose that M is doubly degenerate and $a_- \in A(e_-^M)$ and $a_+ \in A(e_+^M)$. Let $P = P(a_-, a_+)$ be a combinatorial model constructed from a_- and a_+. Then there is a proper uniformly lipschitz map, $f : P \longrightarrow M$, with $f^{-1}(\Psi(M)) = \Psi(P)$, which sends e_-^P to e_-^M and e_+^P to e_+^M, and such that the lifts $\tilde{f} : \tilde{P} \longrightarrow \tilde{M}$ and $\tilde{f}|\Psi(\tilde{P}) : \Psi(\tilde{P}) \longrightarrow \Psi(\tilde{M})$ to universal covers are both uniform quasi-isometries. Here "uniform" means depending only on the parameters of M.*

We will eventually see (Proposition 6.8) that $A(e_\pm^M) = \{a_\pm\}$, and so P depends only on the "end invariants", $a(e_\pm^M)$. However, for the moment, $a_\pm$ may be any elements of $A(e_\pm^M)$. Recall that, by Lemma 4.8, $A(e_-^M) \cap A(e_+^M) = \emptyset$, and so the construction of P makes sense.

For the proof, we will need to recall various facts about P that arise directly from its construction.

There is a closed submanifold, $\Lambda(P) \subseteq \Psi(P)$, such that each component of $\Lambda(P)$ is isometric to the interior of a standard Margulis tube. The set of such tubes is unlinked in $\Psi(P) \cong \Sigma \times \mathbf{R}$, and no two tubes are freely homotopic in $\Psi(P)$. They are therefore in bijective correspondence to a certain subset, $X(P) \subseteq X(\Sigma)$. We write $T(\gamma) \subseteq \Psi(P)$ for the closed tube corresponding to $\gamma \in X(P)$.

The submanifold, $\Lambda(P)$, has bounded geometry, with injectivity radius bounded below. (This is the bit constructed as a union of bricks.) In fact, there is a uniformly bilipschitz homeomorphism from $\Lambda(P)$ to $R(\hat{\Pi})$, where $\hat{\Pi} = \Pi/\Gamma$, and where Π is a 3-dimensional simplicial complex admitting an action of $\Gamma = \pi_1(\Sigma)$. Here $R(\hat{\Pi}) = R(\Pi)/\Gamma$ is the truncated realisation described in Section 5. In fact, Π and Γ satisfy all of the conditions (E1)–(E7) laid out there. Moreover, to each $x \in \hat{\Pi}^0$, we have associated a unique curve $\gamma(x) \in X(P)$ such that the homeomorphism sends $T(\gamma(x))$ to $\hat{D}(x)$, where $\hat{D}(x)$ is the 2-dimensional subcomplex of $R(\hat{\Pi})$ associated to x.

We also note that $X(P) \subseteq \mathcal{H}(a_-, a_+)$, where $\mathcal{H}(a_-, a_+)$ is the hierarchy associated to $a_-, a_+ \in \partial\mathcal{G}(\Sigma)$ described in Section 4. In fact, there is a bi-infinite geodesic $(\gamma_i)_i$ in $\mathcal{G}(\Sigma)$ from a_- to a_+ with $\gamma_i \in X(P)$ for all i. As $i \to \pm\infty$, $T(\gamma_i) \to e^P_\pm$.

Suppose that $Y \subseteq X(P)$ is a subset. We set $\Lambda(P, Y) = \Psi(P) \setminus \bigcup_{\gamma \in T} \operatorname{int} T(\gamma)$. (Thus, $\Lambda(P, \emptyset) = \Psi(P)$ and $\Lambda(P, X(P)) = \Lambda(P)$.) We shall assume that for each $\gamma \in X(P) \setminus Y$, $T(\gamma)$ has bounded depth. (A particular choice of Y will be given later, when the construction is applied.) This implies that the injectivity radius of $\Lambda(P, Y)$ is bounded below. We define a reduced pseudometric $\rho = \rho_P$ on $\Psi(P)$, similarly as for $\Psi(M)$, by deeming each tube $T(\gamma)$ for $\gamma \in Y$ to have ρ-diameter 0. We see that ρ is a proper pseudometric on $\Psi(P)$. It turns out that $(\Psi(P), \rho)$ is quasi-isometric to the real line. In fact, we can be more explicit.

Given $x \in \Psi(P)$, we can find an embedded surface, S, in $\Psi(P)$, such that the inclusion $(S, \partial S)$ into $(\Psi(P), \partial\Psi(P))$ is a relative homotopy equivalence, and such that the ρ-diameter of $\{x\} \cup S$ is uniformly bounded above. In fact, we can assume that $S \subseteq \Lambda(P, Y)$.

Now, since ρ is proper, we can find a proper map $\pi : \mathbf{R} \longrightarrow (\Psi(P), \rho)$ which is uniformly quasi-isometric (that is, for all $t, u \in \mathbf{R}$, $\rho(\pi(t), \pi(u))$ is bounded above and below by uniform increasing linear functions of $|t - u|$), and such that $\pi(t) \to e^P_\pm$ as $t \to \pm\infty$. (In fact, we can take π to be as close as we want to being geodesic with respect to ρ.) Now, given any $t \in \mathbf{R}$, we can find a surface, $S_t \subseteq \Lambda(P, Y) \subseteq \Psi(P)$ of the type described above such that the ρ-diameter of $\{\pi(t)\} \cup S_t$ is bounded above. From this, we can deduce that π is, in fact, a uniform quasi-isometry from $\mathbf{R}$ to $(\Psi(P), \rho)$. Recall that this applies to any subset $Y \subseteq X(P)$ so that the injectivity radius of $\Lambda(P, Y)$ is bounded below. We will settle on a particular subset later.

We will also need a fact about the non-distortion of components of $\partial\Lambda(P, Y)$ in $\Lambda(P, Y)$. This will be mentioned again when it is applied.

Now let M be a coarse hyperbolic 3-manifold and let $a_- \in A(e^M_-)$ and $a_+ \in A(e^P_+)$, and let $P = P(a_-, a_+)$ be a combinatorial model constructed from a_- and a_+.

Lemma 6.2 : *There is some uniform $L \geq 0$ such that $X(P) \subseteq X(M, L)$.*

Proof : We know from the construction that $X(P) \subseteq \mathcal{H}(a_-, a_+)$. By Lemma 4.13, we have $\mathcal{H}(a_-, a_+) \subseteq X(M, L)$. $\diamondsuit$

This is the "a-priori bounds" property (E7) listed in in Section 5. As in Section 5, we now obtain a subset $\Theta(M) \subseteq \Psi(M)$, and a lipschitz map of $\Lambda(P)$ to $\Theta(M)$. More precisely:

Lemma 6.3 : *There is a proper uniformly lipschitz map, $f : \Lambda(P) \longrightarrow \Theta(M)$ such that $f^{-1}(\partial\Psi(M)) \subseteq \partial\Psi(P)$ and $f^{-1}(\partial\Theta(M)) \subseteq \partial\Lambda(P)$, and which extends continuously to a relative homotopy equivalence $(\Psi(P), \partial\Psi(P)) \longrightarrow (\Psi(M), \partial\Psi(M))$. Moreover, for each $\gamma \in X(P)$ either $f(\partial T(\gamma))$ lies in the boundary component of $\Theta(P)$, or else $f(\partial T(\gamma)) = \gamma^*$, where γ^* is a shortest representative of γ in M.*

Here, of course, we can assume that the relative homotopy equivalence $(\Psi(P), \partial\Psi(P)) \longrightarrow$ $(\Psi(M), \partial\Psi(M))$ commutes with the natural homotopy equivalence of these pairs to $(\Sigma, \partial\Sigma)$.

Proof : By construction, $\Lambda(P)$ is uniformly bilipshitz equivalent to a truncated complex, $R(\hat{\Pi})$. We can therefore apply Lemma 5.7 to get a uniformly lipschitz map from $R(\hat{\Pi})$ to $\Psi(M)$. Precomposing with the bilipschitz map to $\Lambda(P)$, we get a lipschitz map $f : \Lambda(P) \longrightarrow \Psi(M)$. The fact that $\Lambda(P) \subseteq \Theta(M)$ and the other properties listed, follow from Lemma 5.7. $\diamondsuit$

We set $Y \subseteq X(P)$ to be the set of curves $\gamma \in X(P)$ such that $f(\partial T(\gamma))$ lies in a boundary component of $\Theta(M)$, and write $\Theta(P) = \Lambda(P, Y) \subseteq \Psi(P)$. If $\gamma \in X(P) \setminus Y$, then $f(\partial T(\gamma)) = \gamma^*$, which has length bounded below by η_3 and bounded above by L. Since $f|\partial T(\lambda)$ is lipschitz, we see, as in [Bow5] that $T(\gamma)$ has bounded depth, and that $f|\partial T(\gamma)$ extends to a uniformly lipschitz map from $T(\gamma)$ to γ^*. (This only depends on the intrinsic geometry of γ^* as a circle.) We thus obtain a uniformly lipschitz map $f : \Theta(P) \longrightarrow \Theta(M)$ with $f^{-1}(\partial\Theta(M)) \subseteq \Theta(P)$. We can extend this to a homotopy equivalence $f : (\Psi(P), \partial\Psi(P)) \longrightarrow (\Psi(M), \partial\Psi(M))$, though for the moment, it is only defined topologically in $\Psi(P) \setminus \Theta(P)$.

As in [Bow5] we see that f sends $e^P_\pm$ to $e^M_\pm$. (This stems from the fact that we have a bi-infinite geodesic $(\gamma_i)_i$ in $\mathcal{G}(\Sigma)$, with realisations as curves of bounded length in $\Psi(P)$ and $\Psi(M)$ respectively, and tending out $e^P_\pm$ and $e^M_\pm$ resepectively, as $i \to \pm\infty$.) One then deduces that f is surjective. By Lemma 5.7, we see that the injectivity radius of $\Theta(M)$ is uniformly bounded below, and so is a "thick part" of M.

We can now prove Theorem 0.6, namely that the components of $\Psi(M) \setminus \Theta(M)$ are unlinked in $\Psi(M)$. This results from the following topological statement:

Lemma 6.4 : *Suppose that $\mathcal{L}$ and $\mathcal{L}'$ are locally finite collections of disjoint essential non-peripheral closed curves in $\Sigma \times \mathbf{R}$. Suppose that $f : \Sigma \times \mathbf{R} \longrightarrow \Sigma \times \mathbf{R}$ is a surjective proper homotopy equivalence of the pair $(\Sigma \times \mathbf{R}, \partial\Sigma \times \mathbf{R})$. Suppose that $f^{-1}(\bigcup \mathcal{L}') = \bigcup \mathcal{L}$ and that $f|\bigcup \mathcal{L} : \bigcup \mathcal{L} \longrightarrow \bigcup \mathcal{L}'$ is a homeomorphism. If $\mathcal{L}$ is unlinked in $\Sigma \times \mathbf{R}$, then so is $\mathcal{L}'$.*

Proof : This argument is an application of [FrHS], or more precisely, the topological tower construction used there in the case of a homotopy equivalence of a surface to a 3-manifold. The account given in [O] can be reinterpreted in a general topological setting as outlined in [Bow5]. The idea is as follows.

Since $\mathcal{L}$ is unlinked, we can assume that each $\gamma \in \mathcal{L}$ lies in a surface $S(\gamma)$ of the form $\Sigma \times \{t(\gamma)\}$ where $t : \mathcal{L} \longrightarrow \mathbf{R}$ is injective with discrete image. For each $\gamma \in \mathcal{L}$ whe have $\gamma' \subseteq f(S(\gamma))$, where $\gamma' = f(\gamma) \in \mathcal{L}'$. We surger $f(S(\gamma))$ as in [Ot] using the methods of

[FrHS] to obtain an embedded surface $S'(\gamma) \supseteq \gamma'$, whose inclusion into $\Sigma \times \mathbf{R}$ is a relative homotopy equivalence $(S'(\gamma), \partial S'(\gamma)) \longrightarrow (\Sigma \times \mathbf{R}, \partial\Sigma \times \mathbf{R})$. Moreover, we can arrange that the surfaces $S'(\gamma)$ are all disjoint. Since $f | \bigcup \mathcal{L} : \bigcup \mathcal{L} \longrightarrow \bigcup \mathcal{L}'$ is a homeomorphism, each element of $\mathcal{L}'$ lies in such a surface, and so $\mathcal{L}'$ is unlinked. $\diamondsuit$

Proposition 6.5 : *The components of $\Psi(M) \setminus \Theta(M)$ are unlinked in $\Psi(M)$.*

Proof : We have defined $f : \Psi(P) \longrightarrow \Psi(M)$. By construction, the components of $\Psi(P) \setminus \Theta(P)$ are unlinked in $\Psi(P)$. Moreover, $f^{-1}(\Theta(M)) = \Theta(P)$, and the result follows from Lemma 6.4. $\diamondsuit$

This now proves Theorem 0.6, since we can assume that η was chosen such that $\Theta(M) \subseteq \Theta(M, \eta)$, so that each $T \in \mathcal{T}$ lies in a component of $\Psi(M) \setminus \Theta(M)$.

It now follows as in [Bow5] that $f | \Theta(P) : \Theta(P) \longrightarrow \Theta(M)$ is a homotopy equivalence, and homotopic to a homeomorphism.

The remainder of the proof of Proposition 6.1 follows as in [Bow5], where only the bounded geometry of $\Theta(M)$ was used.

Briefly, one shows that the lift, $\tilde{f} : \tilde{\Theta}(P) \longrightarrow \tilde{\Theta}(M)$ to universal covers is a uniform quasi-isometry. Moreover, we can homotope f a bounded amount near $\partial\Theta(P)$ so that it becomes uniformly bilipschitz from $\partial\Theta(P)$ to $\partial\Theta(M)$. (This uses a fact about the non-distortion of $\partial\Theta(P)$ in $\Theta(P)$, alluded to earlier.) One can now extend in bilipschitz fashion over the Margulis tubes and $\mathbf{Z}$-cusps, giving a map $f : P \longrightarrow M$ with the required properties.

This proves Proposition 6.1.

The following corollary is immediate:

Corollary 6.7 : *Let M, M' be doubly degenerate coarse hyperbolic 3-manifolds with $A(e_-^M) \cap A(e_-^{M'}) \neq \emptyset$ and $A(e_+^M) \cap A(e_+^{M'}) \neq \emptyset$. Then there is a uniform equivariant quasi-isometry from $\tilde{M}$ to $\tilde{M}'$, sending $\Psi(\tilde{M})$ to $\Psi(\tilde{M}')$.* $\diamondsuit$

The quasi-isometry respects the ends of $\Psi(M)$ and $\Psi(M')$ in the sense that we can find sequences $x_i^\pm$ in $\Psi(\tilde{M})$ and $y_i^\pm$ in $\Psi(\tilde{M}')$ related by the quasi-isometry and whose projections to $\Psi(M)$ and $\Psi(M')$ tend out $e_\pm^M$ and $e_\pm^{M'}$ respectively.

We remark that in the course of proving Proposition 6.1, it is shown that $f : (\Psi(P), \rho_P) \longrightarrow (\Psi(M), \rho_M)$ is a quasi-isometry, where ρ_P and ρ_M are the reduced metrics on $\Psi(P)$ and $\Psi(M)$ respectively. In particular, we see:

Proposition 6.8 : *$(\Psi(M), \rho_M)$ is quasi-isometric to the real line.* $\diamondsuit$

A couple of the trickier statements in the introduction remain to be proven. First, we have the uniqueness of the end invariant.

For this, let e be a simply degenerate end of a coarse hyperbolic 3-manifold, M. We want to show:

Proposition 6.9 : *$A(e) = \{a\}$ for some $a \in \partial\mathcal{G}$.*

Proof : Let e' be the other end (about which we make no assumptions). We can find a map $\pi_M : \mathbf{R} \longrightarrow \Psi(M)$ that is uniformly quasi-isometric with respect to the metric ρ_M on $\Psi(M)$. (Indeed we can take it abritrarily close to geodesic.) Moreover, $\pi_M(t) \rightarrow e$ as $t \rightarrow \infty$ and $\pi_M(t) \rightarrow e'$ as $t \rightarrow -\infty$. The existence of π_M is an easy consequence of the fact that ρ_M is proper.

Suppose, for contradiction that we have $a, b \in A(e)$ with $a \neq b$. Let $P = P(a, b)$ be a combinatorial model.

We now proceed as in the proof of Proposition 6.1. We get a lipschitz map, $f : \Lambda(P) \longrightarrow M$ with $f(\Lambda(P)) \subseteq \Theta(M)$. Setting $Y = \{\gamma \in X(P) \mid f(\partial T(\gamma)) \subseteq \partial\Theta(M)\}$ and $\Theta(P) = \Lambda(P, Y)$ we extend to a lipschitz map $f : \Theta(P) \longrightarrow \Theta(M)$, and then continuously to a map $f : \Psi(P) \longrightarrow \Psi(M)$, which is a relative homotopy equivalence of $(\Psi(P), \partial\Psi(P))$ to $(\Psi(M), \partial\Psi(M))$. This time, however, both ends of $\Psi(P)$ get sent out the same end, e, of $\Psi(M)$.

Let $\pi_P = \pi : \mathbf{R} \longrightarrow \Psi(P)$ be a quasi-isometric map to $(\Psi(P), \rho_P)$ as described earlier. Given $t \in \mathbf{R}$, there is a surface, $S_t \subseteq \Theta(P)$ such that the ρ_P-diameter of $\{\pi_P(t)\} \cup S_t$ is uniformly bounded above. Now $f(S_t)$ must separate the ends of M, and so must meet $\pi_M(\mathbf{R})$. Choose $t' \in \mathbf{R}$ so that $\pi_M(t') \in f(S_t)$ and set $\sigma(t) = t'$. Now the ρ_M-diameter of $f(S_t)$ is also bounded above, so this is well-defined up to an additive constant. Moreover, the map $\sigma : \mathbf{R} \longrightarrow \mathbf{R}$ is quasi-isometric. These statements follow exactly as in [Bow5] (where they formed an essential part of the proof that $\tilde{f} : \tilde{\Theta}(P) \longrightarrow \tilde{\Theta}(M)$ is a quasi-isometry in the doubly degenerate case).

But now, since f sends both ends of $\Psi(P)$ out e, we see that $\sigma(t) \rightarrow +\infty$ as $t \rightarrow -\infty$ and as $t \rightarrow +\infty$. This gives a contradiction, thereby proving Proposition 6.9. $\Diamond$

We can now write $A(e) = \{a(e)\}$, where $a(e) \in \partial\mathcal{G}(\Sigma)$ is the *end invariant* of e.

Note that we have now proven Theorem 0.4, and the definition of "end invariant" arising from it agrees with that defined above.

In the light of this, Theorem 0.5(1) arises from Lemma 4.8 and Theorem 0.7 arises from Corollary 6.7.

It still remains to prove Theorem 0.5(2).

We begin by remarking that we can now always reduce to the case of a (constant curvature) hyperbolic 3-manifold, since any coarse hyperbolic 3-manifold is equivalent to such, in the sense that their universal covers are equivariantly quasi-isometric. In fact, if M is doubly degenerate, there is a unique doubly degenerate hyperbolic 3-manifold, M', with the same pair of end invariants (using Thurston's double limit theorem). Theorem 0.7 then tells us that $\tilde{M}$ and $\tilde{M}'$ are uniformly equivariantly quasi-isometric. In particular, curves of bounded length in M will also be of bounded length in M'. Theorem 0.5(2) is already known for such manifolds, and the general case then follows.

One can argue more directly along the following lines. First note that $X(M, l_0)$ is unformly quasiconvex, and it follows that any bi-infinite geodesic in $\mathcal{G}(\Sigma)$ from $a(e_+)$ to $a(e_-)$ lies in a bounded neighbourhood of $X(M, l_0)$. Conversely, there is a bi-infinite geodesic $(\gamma_i)_i$ in $\mathcal{G}(\Sigma)$ with length$(\gamma_i^*) \leq l_0$ for all i and with $\gamma_i^* \rightarrow e_\pm$ as $i \rightarrow \pm\infty$. Note that $\rho_M(\gamma_i^*, \gamma_{i+1}^*)$ is bounded for all i. It now follows that any point of M is a bounded

ρ_M-distance from some γ_i^* (cf. Propostion 6.8). To complete the argument, we need the following general lemma:

Lemma 6.10 : *Given $l, h \geq 0$, there is some $r \geq 0$ depending only on l, h and the parameters of M such that if $\beta, \gamma \in X(M, l)$ with $\rho_M(\beta^*, \gamma^*) \leq h$, then $d_{\mathcal{G}(\Sigma)}(\beta, \gamma) \leq r$.*

Here M can be any coarse hyperbolic manifold. Again, in the case of a doubly degenerate manifold, it can be deduced from the corresponding statement in constant curvature.

Alternatively, suppose $\gamma \in X(M, l)$. Then by [Bow6], as in the proof of Lemma 3.2, γ is carried on a train track τ which admits a lipschitz map into M. Since γ^* has bounded length, it passes through a bounded number of branches of τ. Moreover, τ either contains a curve, δ, homotopic to a Margulis tube, or else τ has bounded diameter. In the former case, $d_{\mathcal{G}(\Sigma)}(\gamma, \delta)$ is bounded. In the latter case, we can find thick structure σ, on Σ, and a uniformly lipschitz map of (Σ, σ) into $\theta(M)$, such that γ has bounded length also in (Σ, σ). We can now transfer this picture back to the model space P. We see that γ lies in a surface in the thick part of the model that passes through only a bounded number of brick or Margulis tubes of bounded depth. It follows that γ has bounded intersection number with some curve δ arising the hierarchy used to construct the model. In either of the above cases, we see that γ is a bounded distance in $\mathcal{G}(\Sigma)$ from a curve in a bi-infinite geodesic from used in the construction of the model. The result now follows.

References.

[Ba] J.Barnard, *Geometrically infinite surfaces in 3-manifolds with hyperbolic fundamental group* : preprint, Norman (2007).

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

[Bow1] B.H.Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 64–167.

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

[Bow3] B.H.Bowditch, *Relatively hyperbolic groups* : preprint, Southampton (1999).

[Bow4] B.H.Bowditch, *Stacks of hyperbolic spaces and ends of 3-manifolds* : preprint, Southampton (2002).

[Bow5] B.H.Bowditch, *Geometric models for hyperbolic 3-manifolds* : preprint, Southampton (2005).

[Bow6] B.H.Bowditch, *Train tracks and surface groups acting on hyperbolic spaces* : preprint, Warwick 2010.

[Bow7] B.H.Bowditch, *Surface group actions and length bounds* : preprint, Warwick 2010.

[BrBES] J.F.Brock, K.Bromberg, R.Evans, J.Souto, *Grafting, drilling and ending laminations* : in preparation.

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

[C] R.D.Canary, *Ends of hyperbolic 3-manifolds* : J. Amer. Math. Soc. **6** (1993) 1–35.

[Fa] B.Farb, *Relatively hyperbolic groups* : Geom. Funct. Anal. **8** (1998) 810–840.

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

[GhH] E.Ghys, P.de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Mathematics No. 83, Birkhäuser (1990).

[Grom] M.Gromov, *Hyperbolic groups* : in "Essays in group theory", (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer (1988) 75–263.

[GrovM] D.Groves, J.Manning, *Dehn filling in relatively hyperbolic groups* : preprint (2007)

[H] 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.

[K] E.Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space* : preprint, Michigan (1999).

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

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

[Mc] D.McCullough, *Compact submanifolds of 3-manifolds with boundary* : Quart. J. Math. **37** (1986) 299-307.

[Mi] Y.N.Minsky, *The classification of Kleinian surface groups I : Models and bounds* : Ann. of Math. **171** (2010) 1–107.

[Mo] L.Mosher, *Stable Teichmüller quasigeodesics and ending laminations* : Geom. Topol. **7** (2003) 33-90.

[Os] D.V.Osin, *Peripheral fillings of relatively hyperbolic groups* : Invent. Math. **167** (2007) 295–326.

[Ot] 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.

[S] T.Soma, *A geometric approach to the ending lamination conjecture* : preprint, Tokyo Metropolitan University (2008).

[T] W.P.Thurston, *The geometry and topology of three-manifolds* : preprint, Princeton (1979).