Train tracks and surface groups acting on hyperbolic spaces

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

0. Introduction.

Let Σ be a compact surface, and let M be a path metric space with $\pi_1(M) \cong \pi_1(\Sigma)$ and with the universal cover $\tilde{M}$ Gromov hyperbolic. We show that, under certain assumptions, the shortest realisation in M of a multicurve Σ can be approximated by a train track in M. In the case where M is a hyperbolic 3-manifold, this is well known (see for example [Bon]). The main result is stated as Theorem 0.1 here. Among other things, one can deduce that, under a certain assumption on the "systole", any shortest realisation of a multicurve in the quotient has finite diameter (Corollary 0.2). Under the same assumption, we recover the result of [Ba] regarding the existence of short standard generating sets (Corollary 0.3), though this is much less direct than following the original proof. Our broader aim will be to generalise certain arguments from the proof of the Ending Lamination Conjecture to this broader setting. In this, train tracks will play the role of pleated surfaces. These issues will be discussed in [Bow1,Bow2].

Let (H, d_H) be a k-hyperbolic path-metric space, and suppose that Γ is a group acting by isometry. Since the main results are quasi-isometry invariant, we can, without any essential loss of generality, make a few simplifying assumptions for the purposes of exposition.

We will assume that H is simply connected geodesic space, and that Γ acts freely on H with discrete orbits, so that we get a quotient space $M = H/\Gamma$, with $\pi_1(M) = \Gamma$, and $H = \tilde{M}$. We also assume that, for all $g \in \Gamma$, the minimum, $|g| = \min\{d_H(y, gy) \mid y \in H\}$, is attained at some $x \in H$. Thus, we get a g-invariant *axis*, $\bigcup_{i \in \mathbf{Z}} g^i \alpha$, where α is any geodesic from x to gx in H. This projects to a closed curve in M, which has minimal length in its homotopy class. We justify the above assumptions at the end of this section.

Suppose now that $\Gamma = \pi_1(\Sigma)$ where Σ is a compact orientable surface with (possibly empty) boundary, $\partial\Sigma$. We assume that Σ is not a closed torus. We write $\tilde{\Sigma}$ for its universal cover.

If α is an essential closed curve in Σ, we write $g(\alpha) \in \Gamma$ for the corresponding element of Γ (defined up to conjugacy in general). We set $l_M(\alpha) = |g(\alpha)|$. Thus, α is represented by a shortest curve (not necessarily unique), α_M in M, of length $l_M(\alpha)$. We can construct a locally injective map $q : \alpha \longrightarrow M$ with $q(\alpha) = \alpha_M$. We write d_α for the induced path-metric on α. If α is simple, we write $\tilde{\alpha}$ for the preimage in $\tilde{\Sigma}$. Thus q lifts to a Γ-equivariant map, $\tilde{q} : \tilde{\alpha} \longrightarrow \tilde{M} = H$.

Definition : We refer to α_M as a *reasilisation* of α in M.

More generally, this applies to a disjoint union, γ, of simple closed curves. We again get $q : \gamma \longrightarrow M$ so that $\gamma_M = q(\gamma)$ is shortest in its homotopy class. We write $l_M(\gamma) = \sum_\alpha l_M(\alpha)$, where the sum is taken over the set of components, α, of γ. (It is 0 if $\gamma = \emptyset$.) We similary refer to γ_M as a *realisation* of γ. We say that γ is a *multicurve* if it is non-empty and no component is peripheral, and no two components are homotopic.

We will use the notion of a "train track", $\tau \subseteq \Sigma$, here abbreviated to "track", and associated terminology, see for example [PH]. Here, we will assume that $\partial\Sigma \subseteq \tau$. No component of $\Sigma \setminus \tau$ can be a smooth disc, monogon, digon or smooth annulus. We say that τ *fills* Σ if each component of $\Sigma \setminus \tau$ is a topological disc or peripheral annulus (so that $\Gamma = \pi_1(\Sigma)$ is carried on τ). We write $\tilde{\tau}$ for the lift τ to $\tilde{\Sigma}$. We also recall the notions of "trainpath" and "carrying map". If γ is a multicurve, and $p : \gamma \longrightarrow \tau$ is a carrying map, then there is a unique lift $\tilde{p} : \tilde{\gamma} \longrightarrow \tilde{\tau}$. We write $E(\tau)$ and $V(\tau)$ for the sets of branches and switches respectively.

By a *loop*, $\delta \subseteq \tau$ mean a circular subtrack conisisting of a single branch and a single switch of τ. It is *simple* if at the switch, there is exactly one incident half-branch on each side of δ, and arranged so that there is a local trainpath of τ meeting δ precisely at the switch. If δ happens to be a boundary component, then there is just one branch emerging from the switch, and we drop the final condition. Note that distinct simple loops are necessarily disjoint. An annular neighbourhood, A, of δ is *simple* if $A \cap \tau$ consists of either just δ, or else γ together with an interval in each of the adjacent edges. By a *special set* of simple loops, we mean a preferred (possibly empty) set, $\mathcal{S}$, of loops, together with a set, $A(\mathcal{S}) = \{A(\delta) \mid \delta \in \mathcal{S}\}$, of disjoint simple annular neighbourhoods. We write $\tilde{\mathcal{S}}$ for the set of lifts of elements of $\mathcal{S}$ to $\tilde{\Sigma}$. If $\tilde{\delta} \in \mathcal{T}$, we write $A(\tilde{\delta})$ for the corresponding lift of $A(\delta)$. We have an obvious interpretation of $A(\tilde{\mathcal{S}})$, etc.

The following is the main result:

Theorem 0.1 : *Suppose that $\gamma \subseteq \Sigma$ is a multicurve, and $q : \gamma \longrightarrow M$ is a realisation of γ (i.e. so that $q(\gamma)$ is minimal length in its homotopy class). Suppose $\nu_0 \in \mathbf{N}$. Then there is a track, $\tau \subseteq \Sigma$, filling Σ, with a special set, $\mathcal{S}$, of loops, a path-metric, d_τ, on τ, a carrying map $p : \gamma \longrightarrow \tau$ and a map $f : \tau \longrightarrow M$ such that:*

(T1) $p : (\gamma, d_\gamma) \longrightarrow (\tau, d_\tau)$ is ξ_0-lipschitz.

(T2) For all $x \in \tilde{\gamma}$, $d_{\tilde{M}}(\tilde{q}(x), \tilde{f} \circ \tilde{p}(x)) \le h_0$.

(T3) The d_τ-length of $\tau \setminus \bigcup A(\mathcal{S})$ is at most $h_1(\nu_0)$.

(T4) $f : (\tau \setminus \bigcup \mathcal{S}, d_\tau) \longrightarrow (M, d_M)$ is $\xi_1(\nu_0)$-lipschitz.

(T5) If $\tilde{\delta} \in \tilde{\mathcal{S}}$ and $x \in A(\tilde{\delta}) \cap \tilde{\tau}$, then for all $i \in \{1, \ldots, \nu_0\}$, $d_{\tilde{M}}(\tilde{f}(x), g^i \tilde{f}(x)) \le h_2$, where $g = g(\tilde{\delta}) \in \Gamma$ corresponds to δ.

Moreover, if $\nu \ge \nu_0$, we can find some subset $\mathcal{S}_\nu \subseteq \mathcal{S}$, and a simple annulus system $A_\nu(\delta) \subseteq A(\delta)$ for $\delta \in \mathcal{S}_\nu$ such that:

(T6) The d_τ-length of $\tau \setminus \bigcup A_\nu(\mathcal{S}_\nu)$ is at most $h_1(\nu)$.

(T7) f is $\xi_1(\nu)$-lipschitz on each $\delta \in \mathcal{S} \setminus \mathcal{S}_\nu$.

(T8) If $\tilde{\delta} \in \tilde{\mathcal{S}}_\nu$ and $x \in A_\nu(\tilde{\delta}) \cap \tilde{\tau}$, then for all $i \in \{1, \ldots, \nu\}$, $d_{\tilde{M}}(\tilde{f}(x), g^i \tilde{f}(x)) \le h_2$, where $g = g(\tilde{\delta}) \in \Gamma$ corresponds to δ.

Here, the constants ξ_0, h_0, h_2, and the function ξ_1 and h_1 depend only on $l_M(\partial\Sigma)$, the hyperbolicity constant of $\tilde{M}$, and the topological type of Σ.

Less formally, this says that the realisation of a multicurve in M must factor through a train track, $f(\tau)$, in M, up to bounded distance. Moreover, the part of the train track lying outside the "thin part" of M has bounded length. The "thin part" of M can be defined by the criterion of (T5), and depends on the parameter, ν_0, chosen at the outset. This is the content of (T1)–(T5). If we later change our minds, and decide we want an even thinner thin part, we can choose a larger $\nu \geq \nu_0$. We need then only modify the constants outputted. We do not need to change the track, $f(\tau)$, in M. This is the content of (T6)–(T8).

One context in which the above can be made more explicit, is that of a hyperbolic manifold, M. In this case, we can extend $q : \gamma \longrightarrow M$ to a 1-lipschitz map, $q : \Sigma \longrightarrow M$, where Σ is equipped with some hyperbolic structure. This necessarily maps the thin part of Σ to the thin part of M. Now, γ is realised as a geodesic multicurve in Σ which can be approximated by a train track, τ, in Σ, whose length outside the thin part of Σ is bounded. Moreover, we can assume the track to be simple inside the thin part (which is a disjoint union of annuli). We then take $f = q|\tau$. For further discussion of these constructions (in dimension 3), see for example [Bon].

Note that if $\mathcal{S} = \emptyset$, then (τ, d_τ) has bounded diameter, and it follows that $\gamma_M = q(\gamma)$ has bounded diameter in (M, d_M). Since τ carries $\Gamma = \pi_1(\Sigma)$, it also follows that there is a standard generating set of Γ, represented by paths of bounded length in δ based at some point $a \in \tau$. It follows that $f(a)$ is displaced a bounded distance by each of these generators in $\tilde{M}$.

To relate this to results appearing elsewehere, we make some further definitions.

We say that a subgroup, $G \leq \Gamma$ is *elliptic* if it has a bounded orbit in H (and hence in fact a uniformly bounded orbit in H). If G has a fixed point in ∂H but is not elliptic then the fixed point set in ∂H is either a singleton and G is *parabolic* or a pair and G is *loxodromic*. We say that G is *elementary* if it is elliptic, parabolic or loxodromic.

Given $g \in \Gamma$, we write $||g|| = \lim_{n\to\infty} \frac{1}{n} d_H(x, g^n x)$, where x is any point of H, for the *stable length* of g. Clearly $||g|| \leq |g|$ in fact $|g| - ||g||$ is bounded above by some fixed multiple of the hyperbolicity constant. Moreover, $\langle g \rangle$ is necessarily elementary. If $||g|| = 0$ it is elliptic or parabolic, and if $||g|| > 0$, it is loxodromic.

We define $\mathrm{sys}_0^S(M) = \inf_\alpha \{l_M^S(\alpha)\} \geq 0$, as α varies over all essential simple closed curves in Σ. Note that if $g = g(\tilde{\delta})$ is the element of Γ featuring in (T8) of the theorem, then $||g|| \leq h_2/\nu$, and so $\nu\, \mathrm{sys}_0^S(M) \leq h_2$. If $\mathrm{sys}_0^S(M) = \eta > 0$ and $\nu > h_2/\eta$ then if follows that $\mathcal{S} = \emptyset$.

We deduce the following corollaries:

Corollary 0.2 : *Suppose that $\mathrm{sys}_0^S(M) > 0$ and $\gamma \subseteq \Sigma$ is a multicurve. Then the diameter of any realisation of γ in M is bounded above in terms of $\mathrm{sys}_0^S(M)$, $l_M(\partial\Sigma)$, the hyperbolicity constant of $\tilde{M}$ and the topological type of Σ.*

Corollary 0.3 : *Suppose that* $\mathrm{sys}_0^S(M) > 0$. *Then there is a standard generating set,* $g_1, \dots, g_m$, *for* Γ, *and some* $x \in \tilde{M}$, *such that for each* $i \in \{1, \dots, m\}$, $d_{\tilde{M}}(x, g_i x)$ *is bounded above in terms of* $\mathrm{sys}_0^S(M)$, $l_M(\partial\Sigma)$, *the hyperbolicity constant of* $\tilde{M}$ *and the topological type of* Σ.

Here "standard" can be interpreted to mean arising from a specific combinatorial representation of the surface in terms of a polygon with edge identifications. Thus there is only one standard generating set modulo the mapping class group of Σ.

In the case where $\partial\Sigma = \emptyset$, Corollary 0.2 is proven in [Ba]. Indeed the argument given there also effectively proves Corollary 0.3 in this situation. Such arguments can also be generalised, without too much difficulty to the case where $\partial \neq \emptyset$ so as to give the above statement. This would be much more direct than following the proof in the present paper, though we shall not pursue that matter here.

Finally, we justify the assumptions made earlier on Γ and H, starting from an arbitrary isometric action of a group, Γ on a Gromov hyperbolic space, H. (Here Γ could be any group, though our only interest is in the case where Γ is torsion-free.)

Fix some suitable constant c (a fixed multiple of the hyperbolicity constant) and put the product metric on $\Gamma \times H$, where any two distinct points of Γ are deemed to be distance c apart. Let H' be the 2-dimensional simplicial complex with vertex set $\Gamma \times H$, and where a pair or triple in $\Gamma \times H$ spans a simplex if its diameter is at most $2c$. The diagonal action on $\Gamma \times H$ extends to a simplicial action on H', and H and H' are equivariantly quasi-isometric. The action will be free if Γ is torsion-free. (In general, one may need modify the construction to deal with 2 and 3 torsion.) The results given above when applied to H' can readily be translated back into results about H. Indeed, the arguments can be reformulated directly in terms of the original H, though this becomes more clumsy. We therefore assume we have replaced H by H' and that the conditions of the previous paragraph hold.

1. Outline of proof.

The proof separates into two parts, which use different arguments.

Let $\gamma \subseteq \Sigma$ be a multicurve with realisation, $\gamma_M \subseteq M$. The main step is to construct a "binding" of γ in Σ. This consists of a foliation of a closed subset, $\Pi \subseteq \Sigma$, whose leaves are compact, and transverse to γ (except that we allow a closed leaf in Π to be a component of γ). We construct a map of Π into M which extends the realisation of γ.

The leaves of Π when mapped over to M have bounded diameter. More precisely, the image of any interval leaf lifted to $\tilde{M}$ has bounded diameter in $\tilde{M}$, and the image of any circular leaf lifted to $\tilde{M}$ has a fundamental domain of bounded diameter. These bounds depend only on the hyperbolicity constant of $\tilde{M}$, the topological type of Σ, and the length of the realisation of $\partial\Sigma$ in M.

The binding is constructed in Section 5. The key ingredient is the existence of a "transverse graph", Υ, as given by Proposition 2.2. The discussion of this is postponed until later.

After constructing the binding, we can construct a train track as follows. We can assume that each component of the union of the circular leaves is a closed essential annulus in Σ. This corresponds to an infinite cyclic subgroup of $\pi_1(\Sigma) \equiv \pi_1(M)$ acting on $\tilde{M}$. Such a subgroup has relatively simple geometry, and so we can explicitly construct a suitable track in each such annulus (see Lemmas 3.1 and 3.3). We can now collapse down each of the interval leaves to give us a track extending the tracks in the annuli. By construction, γ will be supported on this track. We may need to extend this track to larger track, $\tau \subseteq \Sigma$, so that it fills all of Σ (i.e. each component of $\Sigma \setminus \tau$ is simply connected). Also, one can construct a map $f : \tau \longrightarrow \tilde{M}$ so that the realisation of γ agrees, up to bounded distance, with the composition of the carrying map $\gamma \longrightarrow \tau$ and the map f. One verifies properties (T1)–(T8), thereby proving Theorem 0.

To obtain the transverse graph Υ and hence Π, we need a different construction. To give the idea, we recall a related construction where M is a hyperbolic 3-manifold (cf. [Bon]).

First we extend γ to a triangulation, χ, of Σ, where the vertex set, $V \subseteq \gamma \cup \Sigma$. One then constructs a 1-lipschitz map, $\phi : \Sigma \longrightarrow M$, extending the realisation $\gamma \longrightarrow \gamma_M$, sending each edge of χ to a geodesic segment, and each triangle to a totally geodesic simplex in M. The pullback metric on Σ is hyperbolic, possibly with cone singularites of angle greater than 2π at the vertices. One can then use this to construct a train track in Σ carrying γ. (For example, collapse down a small metric neighbourhood of γ in Σ in the induced metric.)

In our situation, where $\tilde{M}$ is Gromov hyperbolic, we still have a map of the 1-skeleton of χ into M, though we can no longer extend over 2-simplices. Instead we use the induced metric on the 1-skeleton to put a suitable metric on $\Sigma \setminus V$ (see Section 6). We can then use the thin-triangle property of $\tilde{M}$ to relate this to the geometry of M, and construct the transverse graph, Υ, giving rise to the binding Π and track τ. The construction of Υ will be discussed in Section 8.

2. Triangulations and transverse graphs.

Let Σ be a compact surface, and let χ be a triangulation of Σ. We do not require simplices to be embedded on their boundaries, so for example, we allow an edge to be incident on the same vertex at both ends. Such a triangulation lifts to a genuine triangulation, $\tilde{\chi}$, of the universal cover, $\tilde{\Sigma}$, of Σ. We write $V = V(\chi)$ and $E = E(\chi)$ respectively for the sets of vertices and edges of Σ.

From χ we can construct a pair of transerse singular smooth folliations on Σ, termed *vertical* and *horizontal*. The horizontal foliation has degree-1 singularites at each point of V, and tripod singularties at the centre of each triangle. Each edge of χ is contained in a horizontal leaf. Similarly, the vertical foliation is transverse to the interior of each edge of χ and has a tripod singularity at the centre of each triangle. We can arrange that each vertical leaf is compact (but that is not essential to the discussion, for the moment).

Definition : By a *transverse arc* in Σ we mean a smooth arc transverse to the horizontal foliation, disjoint from V, and meeting the singular points at most at its endpoints.

Definition : A *transverse graph*, Υ, is a graph embedded in Σ such that each edge is a transverse arc, and each vertex is either at a (tripod) singular point or lies in $\partial\Sigma$. Each edge is assumed to meet at least one singular point. Moreover, we assume that at each singular point there is at least one incident edge of Υ emerging between any two of the branches of the local horizontal singular leaf through this point.

Thus each vertex of $\Upsilon \setminus \partial\Sigma$ has degree at least 3, and each point of $\Upsilon \cap \partial\Sigma$ is a degree-1 vertex of Υ.

Definition : We say that a triangulation, χ, of Σ is *compatible* with a multicurve, $\gamma \subseteq \Sigma$, if $V \subseteq \Sigma$, and if each component of $\gamma \cup \partial\Sigma$ consists of a single vertex and edge of χ.

It is easily seen that any multicurve admits a compatible triangulation.

We claim:

Lemma 2.1 : *Suppose that $\gamma \subseteq \Sigma$ is a multicurve, and that χ is a compatible triangulation. Then Σ admits a riemannian metric, with each edge of $\tilde{\chi}$ in $\tilde{\Sigma}$ geodesic. Moreover, there is an equivariant map, $\tilde{\phi} : \tilde{\Sigma} \longrightarrow \tilde{M}$, sending each edge of $\tilde{\chi}$ to a geodesic in $\tilde{M}$ and each vertical leaf in each triangle of $\tilde{\chi}$ to a geodesic segment of bounded length. The quotient map $\phi : \Sigma \longrightarrow M$ sends each component of $\gamma \cup \partial\Sigma$ to a realisation in M. Moreover, we can assume this to be any prescribed realisation in M. The length bound depends only on the hyperbolicity constant, k, and the topological type of Σ.* $\diamondsuit$

In the above, we do not assume ϕ to be continuous. It will however be continuous restricted to χ and to each vertical leaf.

The proof of Lemma 2.1, is a simple exercise using the thin triangle property of $\tilde{M}$. We will give a more elaborate treatment in Section 6.

With more work, we can show:

Proposition 2.2 : *We can choose the map $\tilde{\phi} : \tilde{\Sigma} \longrightarrow \tilde{M}$, such that χ admits a transverse graph, Υ with the following properties. The graph $\Upsilon \subseteq \Sigma$ has at most n_0 edges. For each edge, ϵ, of $\tilde{\Upsilon}$, $\operatorname{diam} \phi(\epsilon) \leq l_0$. Moreover, there is some $\eta_0 > 0$, such that if $a \in E(\chi)$ and $\operatorname{length}(a \setminus D) > 2\eta_0$, then each component of $a \setminus (\Upsilon \cup D)$ has length at least η_0. Here l_0, η_0 and n_0 depend only on the topological type of Σ.*

The proof of this involves somewhat different methods, and will be posptoned until Sections 6 to 8. In the meantime, we explain how this gives rise to a track in Σ, and thereby prove Theorem 0.1. This will be achieved in Sections 3 to 5.

3. Tracks in annuli.

We will need a technical lemma regarding tracks in annuli. Suppose that Ω is a topological annulus, with boundary $\partial\Omega = \epsilon \sqcup \epsilon'$. By a *simple track*, $\omega \subseteq \Omega$, we mean a track meeting $\partial\Omega$ transversely such that Ω is a simple annulus with respect to ω, as defined in Section 0. More precisely, ω consists of a core curve, $\delta = \delta(\omega) \subseteq \operatorname{int}\Omega$ of Ω, together with two branches, ζ and ζ' connecting δ to ϵ and ϵ' respectively, and such that $\zeta \cup \zeta'$ is a trainpath. We assume that ω meets $\partial\Omega$ transversely. An annulus, $A \subseteq \Omega$ is *simple* if is simple in the sense of Section 0. Thus, $A \cap \omega$ is a simple track in A.

Suppose that $I \subseteq \epsilon$ and $I' \subseteq \epsilon'$ are intervals, each containing one of the two terminal points of ω. By *crossing arc* we mean an arc, Ω, meeting ϵ and ϵ' in points of I and I' respectively. Let β be a disjoint union of crossing arcs. A map, $p : \beta \longrightarrow \omega$ is a *carrying map* if it is smooth at the switch, and is p is homotopic to inclusion relative to $I \cup I'$ (that is sliding the endpoint in these intervals). In other words, if α is any component of β, then $p(\alpha)$ runs along ζ, then some number (possibly 0) times around δ, and then out along ζ'.

By a realisation of Ω, β, we mean a homotopically non-trivial map, $q : \Omega \longrightarrow M$, such that, for each component, α, of β, the image $q(\beta)$ is of minimal length in its homotopy class relative to its endpoints. This lifts to a map $\tilde{q} : \tilde{\Omega} \longrightarrow \tilde{M} = H$, which is equivariant with respect to the group of covering transformations of $\tilde{\Omega}$, an infinite cyclic subgroup, $G = \langle g \rangle$ of Γ. If $\tilde{\alpha}$ is a lift of a component of β, then $\tilde{q}(\tilde{\alpha})$ is a geodesic in H from $\tilde{q}(\tilde{z})$ to $\tilde{q}(\tilde{z}')$, where $\tilde{z}$ and $\tilde{z}'$ are the endpoints of $\tilde{\alpha}$. We write d_β for the induced path metric on β. We write $\tilde{\omega}$, $\tilde{\delta}$, etc. for the lifts of ω, δ etc.

Lemma 3.1 : *Suppose that $l \geq 0$. Let Ω be an annulus, let I, I' be intervals in each of the two boundary components, and let β be a disjoint union of crossing arcs. Suppose that $q : \Omega \longrightarrow M$ is a realisation of Ω, β, with* $\operatorname{length}(q(\partial\Omega)) \leq l$. *Then there is a simple track, $\omega \subseteq \Omega$, with terminal points in $I \cup I'$, with an intrinsic path metric, d_ω, a carrying map $p : \beta \longrightarrow \omega$, and a map $f : \omega \longrightarrow M$ with $f(\delta(\omega))$ freely homotopic in M to $q(\delta(\omega))$.*

(1) $p : (\beta, d_\beta) \longrightarrow (\omega, d_\omega)$ is ξ_0-lipschitz,

(2) $(\forall x \in \tilde{\beta}) \quad d_{\tilde{M}}(\tilde{q}(x), \tilde{f}(\tilde{p}(x))) \leq h_0$,

Also, given $\nu \in \mathbf{N}$ there is a subset, $A \subseteq \Omega$, which is either empty or a simple annulus, satisfying:

(3) The d_ω-length of $\omega \setminus A$ is at most $h_1(\nu)$,

(4) $f : (\tau \setminus (\delta(\omega) \cap A), d_\omega) \longrightarrow (M, d)$ is $\xi_1(\nu)$-lipschitz.

(5) $(\forall x \in \tilde{A} \cap \tilde{\omega})(\forall i \in \{1, \ldots, \nu\}) \quad d_{\tilde{M}}(\tilde{f}(x), g^i \tilde{f}(x)) \leq h_2$.

Here h_0, h_2 and the functions ξ_0, ξ_1, h_1 depend only on l and the hyperbolicity constant of $\tilde{M}$.

In fact, all we really require of the statement that $\operatorname{length}(q(\Omega)) \leq l$ is that there should be a fundamental domain for each of the two boundary components of $\partial\tilde{\Omega}$ in $\tilde{\Sigma}$ whose $\tilde{q}$-image is of bounded diameter.

The proof is based on some general observations on the geometry of an infinite cyclic subgroup, $G = \langle g \rangle$, acting on a k-hyperbolic space $H = \tilde{M}$. We note that G is parabolic, loxodromic or elliptic. If it is parabolic, it fixes a point of ∂H. If it is loxodromic, there is a bi-infinite geodesic (or path geodesic up to a small additive constant), θ, such that for all i, $g^i \theta$ is a uniformly bounded Hausdorff distance from θ (depending only on k). Such a θ is well defined up to uniformly bounded Hausdorff distance. We refer to it as an *axis* of G. If G is elliptic, there is a G-invariant uniformly quasiconvex set, $Q \subseteq G$, such that for all $x \in Q$ and $g \in G$, $d(x, gx)$ is uniformly bounded. We refer to such a set, Q, as an *almost fixed* set. We can take the above bounds to be a suitable fixed multiple of k (e.g. $100k$).

Lemma 3.2 : *Given $k \geq 0$, there exist $h_3 = h_3(k)$, and given also $\nu \geq 0$, there is some $L = L(l, \nu, k)$ with the following properties. Suppose that $G = \langle g \rangle$ acts on a h-hyperbolic space H. Suppose $y \in H$, with $d(y, gy) \leq l$.*

(1) Suppose G is parabolic with fixed point, $w \in \partial H$, and that $x \in H$ lies in a geodesic ray from y to w with $d(x, y) \geq l$. Then $d(x, g^i(x)) \leq h_3$ for all $i \in \{1, \ldots, \nu\}$.

(2) Suppose that G is loxodromic and that θ is an axis of G.

(2a) If $d(y, \theta) \geq L$, and $x \in \theta$, then $d(x, g^i(x)) \leq h_3$ for all $i \in \{1, \ldots, \nu\}$.

(2b) If x lies in a shortest geodesic from x to θ and $d(x, y) \geq L$, then $d(x, g^i(x)) \leq h_3$ for all $i \in \{1, \ldots, \nu\}$.

(3) Suppose that G is elliptic, and that Q is an almost fixed set. If x lies in a shortest geodesic from x to Q and $d(x, y) \geq L$, then $d(x, g^i(x)) \leq h_3$ for all $i \in \{1, \ldots, \nu\}$.

Proof : The proof is an elementary exercise in hyperbolic spaces, noting that for all $i \in \{1, \ldots, \nu\}$, we have $d(y, g^i y) \leq \nu l$. $\diamondsuit$

To apply this to Lemma 3.1, let $\tilde{q} : \tilde{\Omega} \longrightarrow H = \tilde{M}$ be a lift of Ω, equivariant with respect to $G = \langle g \rangle$. We first consider the case where β consists of just one arc, say from $z \in \epsilon$ to $z' \in \epsilon'$. We can take $I = \{z\}$ and $I = \{z'\}$, so the terminal points of ω will also be at z and z'. Thus, $\tilde{q}(\tilde{\beta})$ is a geodesic from $y = \tilde{q}(\tilde{z})$ to $y' = \tilde{q}(\tilde{z}')$, which it will be convenient to denote by $[y, y']$. We consider a number of cases separately.

Suppose first that G is parabolic with fixed point, $w \in \partial H$. Let $u \in [y, y']$ be a "nearest point" to w. More precisely, using the thin triangles property, we choose some u a bounded distance, depending only on k, from geodesics rays, $[y, w]$ and $[y', w]$ from y and y' respectively, to w. Thus, $u = \tilde{q}(\tilde{v})$ where $\tilde{v}$ is a point of $\tilde{\beta}$, mapping to a point, $v \in \beta$. We let δ be a core curve of Ω such that $\beta \cap \delta = \{v\}$, and set $\omega = \beta \cup \delta$, and orient the switch arbitrarily. (Here, δ will play no role other than to fulfil the requirement that ω be simple track.) Take d_ω to be the pull-back path metric on ω. Below, "length" in ω refers to the metic d_ω.

We define $p : \beta \longrightarrow \omega$ be be inclusion. We define $f : \omega \longrightarrow M$ to be equal to q on β, and to send δ to a shortest loop in the homotopy class relative to $q(v)$. It remains to describe A. To this end we distinguish three cases. If length(ζ) and length(ζ') are both at most L, then we set $A = \emptyset$. If they are both greater than L, we take $A \supseteq \delta$ so that length($\zeta \setminus A$) and length($\zeta' \setminus A$) are both equal to L. If length(ζ) $> L$ and length(ζ') $\leq L$, then we take $A \supseteq \delta$ so that length($\zeta \setminus A$) $= L$ and length($\zeta' \cap A$) is (arbitrarily) small. Now each point of $\tilde{q}(\tilde{A}\cap\omega)$ lies a bounded distance from a point x as featuring in Lemma 3.2(1), and is thus displaced at most a bounded distance, say h_2, by g^i for all $i \in \{1,\dots,\nu\}$. Here h_2 depends only on h_3 and k, and hence ultimately only on k, as required.

Next suppose that that G is loxodromic with axis θ. We in turn split this into two cases. Suppose first that $d(\theta,[y,y'])$ is large in relation to h. We set $u \in [y,y']$ to be a nearest point to θ. We now proceed in a similar manner as above, this time using Lemma 3.2(2b) in place of Lemma 3.2(1).

Secondly we suppose that the distance between $[y,y']$ and θ is a bounded multiple of k. Let $w,w' \in \theta$ be nearest points to y,y' respectively. Then $[y,y']$ is a bounded distance, depending on k, from the path $[y,w]\cup[w,w']\cup[w',y']$, where $[w,w']\subseteq\theta$ and $[y',w']$.

Now the bi-infinite path, $\bigcup_{i\in\mathbf{Z}}[g^jw,g^{j+1}w]$ is a bounded Hausdorff distance from θ. (It is quasigeodesic, but need not be uniformly so.) Moreover, $d(w,gw)$ is bounded above in terms of l and h. Thus, there is some $j\in\mathbf{Z}$ with $d(w',g^jw)$ bounded. Without loss of generality we can assume $j\geq 0$. Let ϕ be the path $[w,gw]\cup\cdots\cup[g^{j-1}w,g^jw]$. We see that $[y,y']$ remains a bounded distance from $[y,w]\cup\phi\cup[g^jw,y']$ (this time, depending on l as well as on h). We refer to the direction from w to g^jw along $[w,gw]$ as the "forward" direction.

Now let ω be a simple track in Ω. We can construct a map $f:\omega\longrightarrow M$, so that its lift, $\tilde{f}:\tilde{\omega}\longrightarrow H$, sends a lift of the path δ to $[w,gw]$ and lifts of the branches, $\tilde{\zeta}$ and $\tilde{\zeta}'$ to $[y,w]$ and $[g^{-j}y,w]$ respectively. We orient the cusp between ζ and δ so that it points in the forward direction. We define the metric d_ω by taking the induced path-metric on $\alpha\cup\zeta'$, and taking the induced path metric on δ scaled so that the length of δ is equal to its stable length in M. We define a carrying map, $p:\beta\longrightarrow\omega$ such that $\tilde{f}\circ\tilde{p}:\tilde{\beta}\longrightarrow H$ sends $\tilde{\beta}$ to $[y,w]\cup\phi\cup[g^jw,y]$. We take f to be linear with respect to the path metric d_ω, and the metric on M. This means that it will be uniformly bilipschitz on $\zeta\cup\zeta'$. If we can bound j (in terms of r and l), it will also be lipschitz on δ (with constant depending on ν). Since $\tilde{q}(\tilde{\beta})=[y,y']$ remains a bounded distance from this path, we can assume it to move each point a bounded distance.

To construct A we split into three cases, similarly as before, depending on whether length(ζ) and length(ζ') are greater than or less than L. This time apply Lemma 3.2(2a). In the first case, where length(ζ) and length(ζ') are both at most L, needs slight modification. If the stable length of is greater than k/ν, say, then we set $A=\emptyset$. The bound on stable length places an upper bound on j in the above, so β is lipschitz on δ (with constant depending on ν). If it is less than k/ν, then any point near the axis gets translated a bounded distance by g^i for all $i\in\{1,\dots,\nu\}$. In this case, we choose $A\supseteq\delta$ so that length$(\zeta\cap A)$ and length$(\zeta'\cap A)$ are both small.

Finally suppose that G is elliptic. Let $Q\subseteq H$ be an almost fixed set. Set $u\in[y,y']$ to be a nearest point to Q. Define $\omega,\beta,\zeta,\zeta'$ and $p:\beta\longrightarrow\omega$ as in the parabolic case. Given ν and hence L we also define A as in the parabolic case. The fact that $d(x,g^ix)$ is bounded for all $x\in\omega\cap A$ and all $i\in\{1,\dots,\nu\}$ now follows from Lemma 5.2(3) (if x is far from Q) or else from the definition of Q, and the fact that it is quasiconvex, (if x is near Q).

The above argument case also deal with the case where β consist of exactly two curves, β_1 and β_2, each connecting z to z', but differing homotopically by a Dehn twist. The $\tilde{q}$-images of $\tilde{\beta}_1$ and $\tilde{\beta}_2$ connect $\tilde{z}$ to $\tilde{z}'$ and $\tilde{z}$ to $g\tilde{z}'$ respectively. Note that these remain a bounded distance apart.

The general case can be reduced to this, since we can decompose β as $\beta_1 \cup \beta_2$, so that they lift to $\tilde{\beta}_1$, a set of intervals connecting $\tilde{I}$ to $\tilde{I}'$, and $\tilde{\beta}_2$, a set of intervals connecting $\tilde{I}$ to $g\tilde{I}'$.

There is a variation on Lemma 3.1, where, instead of a set of crossing arcs realised in M, we have a core curve realised in M.

Lemma 3.3 : *Suppose that $l \geq 0$. Let Ω be an annulus, let z, z' be points in each of the boundary components. Suppose that β is some core curve of Ω. Suppose that $q : \Omega \longrightarrow M$ is a realisation with* $\mathrm{length}(q(\partial\Omega)) \leq l$*, and with $q(\beta)$ of minimal length in its homotopy class. Then there is a simple track, $\omega \subseteq \Omega$, with terminal points at $z \cup z'$, with an intrinsic path metric, d_ω, a carrying map $p : \beta \longrightarrow \omega$ with $p(\beta) = \delta(\omega)$, and a map $f : \omega \longrightarrow M$ with $f(\delta(\omega))$ homotopic to $q(\delta(\omega))$. Also, given $\nu \in \mathbf{N}$ there is a subset, $A \subseteq \Omega$, which is either empty or a simple annulus, and satisfying properties (1)–(5) of Lemma 3.1.*

(Here we can take d_ω to be equal to d_β on $\beta = \delta(\omega)$.) The proof is similar.

We can also allow for an annulus where β is a boundary component of Σ. In this case there will only be one branch, ζ, emerging.

4. Bindings.

To get from transverse graphs to tracks we pass via another notion which we will term a "binding" of a multicurve, γ. Informally, this consists of a subset of our surface, Σ, with a vertical foliation transverse to the multicurve and with all leaves compact. (We also allow for a component of γ be be a closed vertical leaf.) The eventual aim will be to construct a track carrying γ by collapsing those leaves that are intervals. The part consisting of circular leaves will be dealt with separately, using the constructions of Section 3. We may also need to "augment" the binding so that it fills the whole of Σ. Here are some more formal definitions.

Definition : A *binding* of γ consists of a subsurface $\Pi \subseteq \Sigma$ whose boundary, $\partial\Pi$, is a union of *horizonal* and *vertical* parts, $\partial\Pi = \partial_H\Pi \cup \partial_V\Pi$, together with two transverse *horizontal* and *vertical* foliations on Π satisfying the following:

(B1) $\Pi \cap \partial\Sigma$ is a union of vertical leaves.

(B2) $\partial_H\Pi \cap \partial_V\Pi$ is finite, and each component of $\partial_H\Pi$ lies in a horizontal leaf, and each component of $\partial_V\Pi$ lies in a horizontal leaf.

(B3) $\partial_H\Pi \subseteq \gamma$.

(B4) Each component of $\gamma \cap \Pi$ either lies in a horizontal leaf, or is a vertical leaf.

(B5) Each vertical leaf is compact, and each circular vertical leaf is homopically essential in Σ.

(B6) No annular component of $\Sigma \setminus \Pi$ is bounded by two vertical leaves.

Definition : By a *subbinding* of Π, we mean another binding Π' of γ, with $\Pi' \subseteq \Pi$ and where the vertical and horizontal foliations on Π restrict to Π'.

Definition : A *rectangle* is a binding, R, that is topologically a disc and where each of $\partial_H R$ and $\partial_V R$ are a pair of intervals.
A rectangle is *empty* if $R \cap \gamma \subseteq \partial_V R$.

Definition : A *vertical annulus* is a binding, Ω, of γ with $\partial_H \Omega = \emptyset$.
A vertical annulus is *isolated* if $\gamma \cap \Omega$ is either empty or a component of γ.

We talk about a *rectangle* or *vertical annulus* in a binding Π to mean subbindings that are rectangles or vertical annuli respectively.

We make the following observations. Any vertical annulus is a topological annulus. Any two distinct vertical annuli are disjoint. Any isolated annulus is a connected component of $\gamma \cup \Pi$. The binding Π is a union of maximal empty rectangles and closed vertical leaves. Any component of Π that does not meet γ is an isolated maximal annulus.

Definition : The *circular part*, $B(\Pi)$, of Π is the union of all circular vertical leaves.

We may as well assume that there are no isolated vertical leaves (since we can always thicken such a leaf to give a vertical annulus). In this case, $B(\Pi)$ is a subbinding of Π, and consists of a disjoint union of all the maximal vertical annuli. We can also assume that each component of γ that is a vertical leaf lies in the interior of $B(\Pi)$.

Definition : We say that a binding, Π, is *circle-free* if every vertical leaf is an interval.

Definition : By the *circle-free part* $S(\Pi)$ of a binding, Π, we mean the closure of $\Pi \setminus B(\Pi)$.

We can check that $S(\Pi)$ is a circle-free subbinding of Π.

Let $\Phi = \Phi(\Pi)$ be the metric completion of $\Sigma \setminus (\gamma \cup \Pi)$. This is a surface with boundary $\partial\Phi$. There is a natural map $\iota : \Phi \longrightarrow \Sigma$ which is injective on the interior, but may be injective or two-to-one on $\partial\Phi$. We have $\iota(\partial\Phi) \subseteq \gamma \cup \Pi$, and we write $\partial_H \Phi$ and $\partial_V \Phi$ for the primages in Φ of γ and $\partial_H \Pi$ respectively. Note that $\partial_H \Phi \cap \partial_V \Phi$ is finite. We refer to the components of $\partial_H \Phi$ and of $\partial_V \Phi$ as the horizontal and vertical *sides* of Φ.

Definition : A *cusp* of Φ is a vertical side of Φ that is an interval.

Definition : A component, F, of Φ is a *monogon*, respectively *digon*, if F is a topological disc, and has exactly one, respectively two, cusps.

In other words, $\partial_H F$ and $\partial_V F$ each consist respectively of an interval, or of a pair of intervals.

Definition : The *complexity*, $c(\Pi)$, of Π is defined as $c(\Pi) = \chi(\Phi) - \chi(\Sigma)$, where χ denotes Euler characteristic.

We see that the total number of cusps of Φ is exactly $2c(\Pi)$. In particular, $c(\Pi) \geq 0$.

Note that the total number of vertical and horizontal edges of Φ is bounded in terms of $c(\Pi)$ and the topological type of Σ. (In fact, it bounds the combinatorics of Π, in the sense of limiting the number of maximal annuli and rectangles in Π, and hence the manner in which they can intersect.)

Definition : We say that Π is *efficient* if there are no complementary monogons or digons.

If Π is efficient (in particular, if there are no monogons) then any crossing of γ with a closed vertical leaf is essential (so that the total number of crossings is minimal in the free homotopy class). It also follows that no two distinct maximal vertical annuli can be homotopic (otherwise there would be a digon between them).

In fact, if Π is efficient, then $c(\Pi)$ is bounded in terms of type(Σ). This is not hard to see, though we will not formally be using this fact here.

We describe a few moves on an arbitrary binding with the eventual aim of eliminating monogons and digons, so that it will become efficient. These will be performed under additional geometric constraints in Section 5.

(M1) Splitting.

Suppose that $R \subseteq \Pi$ is an empty rectangle and let Π' be the closure of $\Pi \setminus R$. Then Π' is a subbinding. Now $\partial_V R$ meets $\partial_V \Pi$ in 0, 1 or 2 intervals (cusps of Φ). If there are 0, then we would introduce a digon, and $c(\Pi') = c(\Pi) + 1$, so we will not use this operation. If it is 1, then $c(\Pi') = c(\Pi)$, and it is 2, then $c(\Pi') = c(\Pi) - 1$. We refer to either of the last two cases as a "splitting" of Π. The last case is termed a "complete splitting".

(M2) Filling digons.

Suppose that D is a digon component of $\Phi(\Pi)$. We can construct a rectangle, R, with $\partial_V R = \partial_V D$ and $\partial_H R = \partial_H D$. In other words, we equip R with transverse horizontal and vertical foliations. We write $\Pi' = \Pi \cup R$ and connect up the foliations to give horizontal and vertical foliations on Π'. This will be a binding provided that each vertical leaf of Π' is compact. Note that $c(\Pi') = c(\Pi) - 1$. We refer to this process as "filling" the digon D.

A sufficient condition for the compactness of vertical leaves is that we can push one of the horizontal sides of D, along the vertical folation, into a horizonal side of a different component of Φ. In this case, we can construct the vertical foliation on D however we want. This process of pushing will be elaborated on later.

(M3) Modifying an annulus.

Suppose that $\Omega \subseteq \Sigma$ is an essential annulus such that each component of $\gamma \cap \Omega$ is a crossing arc (i.e. connects the two boundary components of Ω), and such that each component of $\Pi \cap \partial\Omega$ is a vertical leaf of Π. We note that the closure, Π_0, of $\Pi \setminus \Omega$ is a subbinding of Π. On Π_0 we take the original foliations. On Ω, we take two new foliations with the vertical foliation consisting of circles, and the horizontal foliation of crossing arcs, including the components of $\gamma \cap \Omega$. We note that $c(\Pi') \leq c(\Pi)$.

The above assumes that Ω is embedded in Σ, though we can modify the construction to deal with an immersed annulus, $\iota : \Omega \longrightarrow \Sigma$. We assume that there is some embedded core curve, $\alpha \subseteq \Omega$ such that $\iota^{-1}\iota(\alpha) = \alpha$. We also assume that $\iota(\partial\Omega)$ meets Π only in vertical leaves. We can now apply the above construction to the component of $\Sigma \setminus \iota(\partial\Omega)$ containing $\iota(\alpha)$, perhaps adjusted slightly so that its closure Ω' is embedded. We can then set $\Pi' = \Pi \cup \Omega'$ similarly as above. We describe this more carefully when we come to apply it in the next section.

All the above operations can only reduce the complexity of the binding.

Next, we consider how bindings give rise to tracks in Σ carrying γ.

Suppose, first, that Π is a circle-free binding. Define an equivalence relation $\sim$ on γ by writing $x \sim y$ if either $x, y \in \gamma \setminus \Pi$ or if $x, y \in \gamma \cap \Pi$ and x, y lie in the same vertical leaf of Π. Set $\rho = \Pi/\sim$. This is the same as taking $\gamma \cup \Pi$ and collapsing each vertical leaf to a point. Now ρ is a graph which admits an embedding Σ that is well defined up to isotopy. In fact, we can extend the quotient map, $\gamma \cup \Pi \longrightarrow \rho$, to a map $\pi : \Sigma \longrightarrow \Sigma$ that is a homotopy equivalence, and so that π is injective on $\Sigma \setminus (\gamma \cup \Pi)$. Note that $\pi(\gamma) = \rho$, and that $\pi|\gamma$ is a local homeomorphism. There is a natural map from Φ to Σ also denoted by π, which agrees with the above on its interior, and sends $\pi(\partial\Phi) \subseteq \rho$.

Recall that a track in Σ consists of an embedded graph where all vertices have degree at least 3, and where each vertex has the structure of a swich (that is, the incident half-edges are partitioned into two non-empty subsets). We also allow for closed circle components. We see that ρ has a cononical switch structure at each vertex such that $\pi : \gamma \longrightarrow \rho$ is a carrying map. Cusps of Φ get mapped to cusps of $\Sigma \setminus \rho$ in the usual sense. In particular, monogons and bigon components of Φ correspond to monogon and bigon components of $\Sigma \setminus \rho$. Note that there can be no smooth disc or annular components of ρ. Thus, if Π is efficient, we see that ρ is a track, in the sense we have defined it. (That is a train track in the traditional sense, except that it might not be connected.)

Suppose now, that Π is any binding with $\partial\Sigma \subseteq \Pi$. Let $B = B(\Pi)$ be the circular part of Π, which we can suppose is a disjoint union of maximal vertical annuli. Let $S = S(\Pi)$ be the circle-free part. Thus $S \cap B = \partial_V S \cap \partial_V B$ consists of finite set of vertical intervals. By modifying Π, slightly, splitting along empty rectangles, we can assume that S meets each component of $\partial_V B$, if at all, in a single vertical interval. Thus, each component of $\partial B \setminus \partial\Sigma$ either is vertical circle or meets Φ in a single cusp. This will simplify a little the following discussion.

Let ρ' be constructed as above from Π' by collapsing each vertical leaf of Π' to a point. This time ρ' is a graph naturally embedded in the closure of $\Sigma \setminus B$. This may have a number of terminal vertices in ∂B. All other vertices of ρ' have the structure of a switch as before.

To complete ρ' to a track, we need some additional data. Suppose that in each non-isolated maximal annular component, Ω, of we have a simple track, ω, in Ω, as defined in Section 3. Thus, $\omega \cap \Omega$ consists of a single point in each component of $\partial\Omega$. We shall assume these terminal points to lie also in S. Thus, after sliding these points in $\Omega \cap S$, we can connect them up to terminal points of ρ', to give a smooth edge crossing $\partial\Omega$. Performing this construction for each non-isolated component of B gives us a track ρ which carries those components of γ that are not vertical leaves in isolated annuli. There is some ambiguity about how the carrying map behaves on B, which will be discussed when we apply the construction. As before, we can extend the carrying map to a map $\pi : \Sigma \longrightarrow \Sigma$, homotopic to the identity and injective on $\Sigma \setminus \Pi$.

To finish, we need to extend ρ to a track that fills Σ. For this we need some more data. Let us suppose that we also have a simple track in each islated component of B. In addition we have a graph in λ embedded in Φ such that for each component F of Φ, $F \setminus \lambda$ is a disc (cf, the notion of carrying graph [Ba]). We assume that π meets $\partial\Phi$ precisely in its terminal vertices. In fact, we can assume that λ meets each component of $\partial\Phi$ in a single point.

Now λ gives us a graph, $\pi(\lambda)$ in Σ, with its terminal vertices in ρ or in boundaries of isolated annuli. We connect these up with the simple tracks in the isolated annuli, and to gives us a graph $\tau \supseteq \rho$ in Σ, with $\partial\Sigma \subseteq \tau$. Each component of Σ is a topological disc. We can now abitrarily assign the vertices of τ that were not already in ρ the structure of switches. This cannot introduce any monogons or digons. Thus, τ has the structure of a track carrying γ and filling Σ.

Definition : By an *augmented binding* we mean binding Π together with with tracks ω in $B(\Pi)$ and λ in $\Phi(\Pi)$ of the type described above.

In other words if we have an augmented binding that is efficient (no monogons or digons in $\Phi(\Pi)$) then we can construct a track τ filling Σ as above.

We next introduce the notion of "pushing" intervals in γ along the vertical foliations. This is best described in the the cover $\tilde{\Pi}$. Given intervals I, we say that *I pushes to J* if there is an rectangle, $\tilde{R}$, in $\tilde{\Pi}$ with $\partial_H \tilde{R} = \tilde{I} \sqcup \tilde{J}$, where $\tilde{I}$ and $\tilde{J}$ are lifts of I and J respectively. (We allow for the possibility that $\tilde{R}$ is degenerate, i.e. $\tilde{R} = \tilde{I} = \tilde{J}$.) This determines a canonical homeomorphism, $\theta : I \longrightarrow J$, obtained by flowing along the leaves. Note that $\tilde{R}$ determines an immersed rectangle, R, in Π. We will abuse notation slightly by writing $\partial_H R = I \cup J$.

We now move on to consider some metric conditions on Π that will be used in Section 5.

Suppose that d_γ is a path-metric on γ.

Definition : Given $\mu > 0$, we say that d_γ is μ-pushing invariant if whenever $I, J \subseteq \gamma$ are intervals and I pushes to J, then the canonical homeomorphism from I to J is μ-bilipschitz.

Note that if Π is a circle-free binding, then we can put a path-metric, d_ρ, on the track ρ, so that the carrying map, $(\gamma, d_\gamma) \longrightarrow (\rho, d_\rho)$ is locally μ-bilipschitz.

Definition : The *horizontal size* of Π is length$(\partial_H \Phi(\Pi))$.

We now bring the space, M, into play. Let M be as in the introduction, with $\pi_1(M) = \Gamma$, and $H = \tilde{M}$ hyperbolic.

Definition : Let Π be a binding of γ. By a *realisation* of Π, we mean a map $\phi : \gamma \cup \Pi \longrightarrow M$, such that $\phi|\gamma$ is a realisation of γ (i.e. $\phi(\gamma)$ has minimal length in its homotopy class) and such that q is continuous on each vertical leaf of Π. Moreover, q has a lift $\tilde{q} : \tilde{\gamma} \cup \tilde{\Pi} \longrightarrow \tilde{M}$, projecting to ϕ. (We regard the choice of lift as part of the structure of q.)

Let d_γ be the pull back metric to γ. We assume that d_γ is μ-pushing invariant for some $\mu \geq 0$.

Suppose that α is a vertical leaf and let λ be a component of $\tilde{\alpha}$. If α is an interval, we define the *size* of α to be $\operatorname{diam}(\tilde{\phi}(\lambda))$. If α is an circle, we define the size to be the the diameter of the $\tilde{\phi}$-image of a fundamental domain for λ. This is well defined up to a factor of 2.

Definition : The *vertical size* of Π is the maximal size of a vertical leaf of Π.

We talk of the size of Π as being "bounded" if there are bounds on its complexity, vertical size and horizontal size.

5. From transverse graphs to tracks.

Let $\gamma \subseteq \Sigma$ be a multicurve, and γ_M a realisation of minimal length in M. Let χ be a compatible triangulation, as in Section 2. We get a map $\phi : \Sigma \longrightarrow M$, lifting to $\tilde{\phi} : \tilde{\Sigma} \longrightarrow \tilde{M}$ (Lemma 2.1). Let Υ be the transverse graph given by Proposition 2.2.

Starting with such a map ϕ we construct our track in three steps.

(1) Use ϕ to construct a binding, Π, of γ of bounded complexity, and a map, $\psi : \Pi \longrightarrow M$ with $\psi(\gamma) = \gamma_M$, such that the pull-back path-metric, d_γ, on Π is μ-pushing invariant for some μ, and such that Π has bounded horizontal size. We also have a carrying graph $\lambda \subseteq \Phi$ and an extension of ψ to λ such that $\psi(\lambda)$ has bounded length. Note that ψ is in the right homotopy class. All the bounds and constant μ depend only on $\operatorname{type}(\Sigma)$ and the hyperbolicity constant.

(2) We modify Π to eliminate monogons and digons, while maintaining geometric control. Thus Π becomes efficient. We also modify $\lambda \subseteq \Phi(\Pi)$ and $\psi|\lambda$ so that it has the same properties as when we started out.

(3) We use Π to construct a track, τ, filling Σ and a carrying map, $p : \gamma \longrightarrow \tau$, in the same manner as that described above. Some of the components of $B(\Pi)$ give rise to the special annuli, $A(\mathcal{S})$. We use ψ to give us a map $f : \tau \longrightarrow M$ with the properties claimed.

Step 1 :

We start with $\Sigma, \chi, \phi, \Upsilon$ as above, with $\phi(\gamma) = \gamma_M$. We can lift $\phi : \Sigma \longrightarrow M$ to $\tilde{\psi} : \tilde{\Sigma} \longrightarrow \tilde{M}$. Suppose that I, J are disjoint intervals in $\tilde{\Sigma}$ transverse to $\tilde{\gamma}$, meeting each component at most once. (In practice, these will be edges of $\tilde{\Upsilon}$.) Let $G(I, J)$ be the set of closures of components of $\tilde{\gamma} \setminus (I \cup J)$ that meet both I and J. It has a natural linear order. Let R be the union of $I \cup J \cup \bigcup G(I,J)$ together with all the bounded components of its complement in Σ. If $G(I,J) = \emptyset$, then $R = \emptyset$. If $G(I,J)$ consists of a single arc, β, then $R = \beta$. Otherwise, R is a topological disc. We can write $\partial \tilde{R} = \partial_H \tilde{R} \cup \partial_V \tilde{R}$, where $\partial_H \tilde{R}$ consists of the two extreme segments of $G(I,J)$, and where $\partial_V \tilde{R} \subseteq I \cup J$. We give $\tilde{R}$ the structure of a rectangle, with the elements of $G(I,J)$ horizontal, and so that $\partial_H \tilde{R}$ and $\partial_V \tilde{R}$ are, respectively, the horizontal and vertical boundary components. We can chose the vertical foliation so that the pushing map is linear on each component of $G(I,J)$.

Recall that the metric d_γ on γ is the pull back under the map $\phi : \gamma \longrightarrow M$. Lifting to Σ, we see that the difference in lengths between any two elements of $G(I,J)$ is bounded above by $\mathrm{diam}(\tilde{\phi}(I)) + \mathrm{diam}(\tilde{\phi}(J))$. Suppose that there is a positive lower bound on the lengths of the elements of $G(I,J)$. Then R is μ pushing invariant, where μ depends only on these bounds and the hyperbolicity constant. Moreover, we can find a realisation $\tilde{\psi} : \tilde{R} \longrightarrow \tilde{M}$, with $\tilde{\psi}$ equal to $\tilde{\phi}$ on $\tilde{\gamma} \cap \tilde{R}$ and with each vertical leaf sent to a piecewise geodesic path with breakpoints in $\tilde{\gamma}$. Using hyperbolicity, we see that such a path must have bounded length. This descends to a rectangle, R, in Σ and a map $\psi : R \longrightarrow M$, with $\psi|(R \cap \gamma) = \phi|(R \cap \gamma)$. The vertical size is bounded above in terms of the hyperbolicty constant and $\mathrm{diam}(\tilde{\phi}(I)) + \mathrm{diam}(\tilde{\phi}(J))$.

Now let C be a component of $\Sigma \setminus \Upsilon$. The horizontal foliation in C can have singularties only in V or on ∂C, and such singularties have positive degree. This means that C must be a disc containing at most one point of V or an annulus with $V \cap C = \emptyset$. For each such component we aim to construct a binding, $\Pi_C \subseteq C$. We eventually piece these together to give us our binding Π.

Suppose first that C is a disc. Let $\tilde{C}$ be a lift of C to Σ. Then $\partial \tilde{C}$ is a closed path consisting of a bounded number of edges of $\tilde{\Upsilon}$, or segment of $\partial\Sigma$. Thus $\mathrm{diam}(\tilde{\phi}(\tilde{C}))$ is bounded.

Suppose that $C \cap V = \emptyset$. Then $\tilde{C} \cap \tilde{\gamma}$ consists of a set of arcs, each of which has length bounded above (by $\mathrm{diam}(\tilde{\phi}(\tilde{C}))$) and below (by the constant, η_0, of Propostion 2.2). For each pair of distinct edges, I, J of $\tilde{\Upsilon}$ in $\partial \tilde{C}$ (with $G(I,J) \neq \emptyset$), we construct a rectangle, $\tilde{R}$, as above, and let $\tilde{\Pi}_C$ be the union of all such rectangles. We construct $\tilde{\psi} : \tilde{\Pi}_C \longrightarrow \tilde{M}$ as above. This descends to Π_C and $\psi : \Pi_C \longrightarrow M$, with bounded vertical size. Since there are boundedly many rectangles, we see that Π_C has bounded horizontal size.

If $C \cap V = \{v\}$ with $v \in V$, then all edges of χ must radiate from v, and so $C \cap \gamma$ consists of a single arc. In this case, we can just set $\Pi_C = \emptyset$.

Suppose now that C is an annulus, with boundary components ϵ_1 and ϵ_2, say. Each ϵ_i is either a boundary component of Σ, or a closed path in Υ with a bounded number of edges. It lifts to a strip, $\tilde{C}$, in $\tilde{\Sigma}$ bounded by $\tilde{\epsilon}_1 \cup \tilde{\epsilon}_2$. For each pair of edge I, J of $\tilde{\Upsilon}$ lying in $\tilde{\epsilon}_1$ and $\tilde{\epsilon}_2$ respectively, we construct a rectangle from $G(I,J)$ as before. We can do this equivariantly with respect to the covering translations of $\tilde{C}$. Their union descends to a binding, Π_0, in C. We have a map, $\psi : \Pi_0 \longrightarrow M$ of bounded vertical size. In the case where $C \cap \gamma \neq \emptyset$, the complement $C \cap \Pi_0$ consists of a bounded number of discs. Each such disc, D, is bounded by two components, say β_1, β_2 of $\gamma \cap C$. We can insert a rectangle in D, by pushing the components of $D \cap \epsilon_1$ and ϵ_2 inward a bit to give horizontal segments, an so that the vertical segments are most of β_1 and β_2. We now foliate this similarly as before, so that the pushing maps are linear, and such that the vertical leaves all close up.

This will give us our binding C. If $C \cap \gamma = \emptyset$, then we perform a similar construction, but starting with a component of $\chi \cap C$ instead of a components of $C \cap \gamma$. This will give us an isolated annulus Π_C and a map $\psi : \Pi_C \longrightarrow M$ of bounded vertical size.

We now piece together all the bindings Π_C, as C ranges over components of $\Sigma \setminus \Upsilon$, to give us our binding Π.

To construct $\lambda \subseteq \Phi$, we begin by taking $\Upsilon \cap \Phi$. This cuts Φ into discs and annuli. Each annulus arises from an annular component of $\Sigma \cap \Upsilon$ and is bounded on one side by a vertical leaf of Π and on the other by a closed path in Υ. It is crossed by some short arc in an edge of χ. By adding these edges, we can arrange the the complement of our graph is a union of disc. We can then delete enough edges so that the complement in each component of Φ consists of just a single disc. This gives us a our graph $\lambda \subseteq \Phi$.

Lifing to Σ, we have an equivariant map $\tilde{\lambda} \longrightarrow M$ sending each edge to a set of bounded diameter. We now modify this, fixing the vertex set, and sending each edge to a geodesics segment. This descends to a map $\psi : \lambda \longrightarrow M$ such that each edge has bounded length.

Step 2 :

We start with a binding, Π, for γ, that is uniformly pushing invariant, and has bounded complexity and horizontal size. We have a map $\psi : \Pi \longrightarrow M$ of bounded vertical size. We also have a carring graph $\lambda \subseteq \Phi(T)$ and an extension $\psi : \lambda \longrightarrow M$ in the right homotopy class. We first need to eliminate monogons and digons from $\Phi(\Pi)$. At the end, we will describe how to modify λ.

(a) Eliminating monogons.

Let D be a monogon component of $\Phi(\Pi)$. Thus, $\partial_H D$ and $\partial_V D$ are both intervals. Let $R \subseteq \Pi$ be the maximal empty rectangle on the other side of $\partial_V D$, in other words so that $\partial_V D \subseteq \partial_V R$.

We construct Π' as the full splitting of Π along R. Thus, $c(\Pi') \leq c(\Pi)$. Since $\Pi' \subseteq \Pi$, Π' is uniformly pushing invariant, and restricting ψ to Π' it has bounded vertical height. Note that Π' remains uniformly pushing invariant, and the vertical size cannot increase. The horizontal size has increased by $\mathrm{length}(\partial_H R)$ so we need to check that this is bounded.

Let I, J be the components of $\partial_H R$, so that $\alpha = I \cup \partial_H D \cup J$ is an arc in γ. Let $\theta : I \longrightarrow J$ be the pushing map, which we are know to be uniformly bilipschitz. Since $\psi(\Pi)$ has bounded vertical length, lifting to $\tilde{\Sigma}$, we see that $\tilde{\psi}|I$ and $\tilde{\psi} \circ \theta|I$ are a bounded distance apart. But $\tilde{\psi}|\tilde{\alpha}$ is a uniform quasigeodesic, and θ reverses orientation along α. It follows that I and J have bounded length, as required.

After performing this construction a bounded number of times (depending on the original complexity of Π) eliminate all monogons while maintaining a bound on the horizontal length of Π.

(b) Eliminating digons.

Let D be a digon component of Φ. We label the components if $\partial_H D$ arbitrarily as $\partial_H^+ D$ and $\partial_H^- D$ (the "top" and "bottom" sides) and of $\partial_V D$ as $\partial_V^+ D$ and $\partial_V^- D$ (the "left" and "right" sides).

We begin by pushing $\partial_H^+ D$ upwards as far as we can. If we runt into a component of $\partial\Sigma$, we can immediately apply move (M2) to eliminate D, so we suppose that this does not happen. Thus, we arrive an interval, I, meeting $\partial_H \Pi$ non-trivially. There is a rectangle in Π with horizontal sides $\partial_H^+ D$ and I. (One can see that this rectangle must be embedded in Π, though an immersed recangle would be sufficient for subsequent arguement.) Since $\partial_H^+ D \subseteq \partial\Phi$, the length of $\partial_H^+ D$ is bounded, and since Π is uniformly pushing invariant, so also is the length of I.

Let $K \subseteq I$ be the set of points with meet in Φ on the top side of I. We can assume that K is connected. Otherwise, there would be a emply rectangle in Π with one of its horizontal sides lying in I. Since it has bounded length, we could remove it, thereby reducing $c(\Pi)$ while maintaining a bound on the horizontal length. In now follows, in particular, that K, lies in a single component, F, of Φ.

Suppose that $F \neq D$. Then we can assume that $I \subseteq \partial_H F$ — again by splitting along a empty rectangle on the top of I. (This will not be a full splitting though). We can then fill the digon D by inserting an empty rectangle, R_0, say, to give us a binding $\Pi' \supseteq \Pi$. This reduces complexity and horizontal length. We can construct the vertical foliation so that the pushing map between the horizontal components is linear, and hence unformly bilipshitz. Lifting to $\tilde{\Pi}$, we see that that $\mathrm{diam}(\tilde{\phi}(\partial\tilde{D}))$ is bounded. We can thus extend this to a map $\tilde{R}_0$ so that the $\tilde{\psi}$-image of each vertical leaf has bounded diameter. This descends to a map $\psi : R_0 \longrightarrow M$. We see that $\psi(\Pi')$ has bounded vertical size.

We can therefore assume that I meets D non-trivially in a single interval. This must lie in $\partial_H^- D$.

Similarly, we push $\partial_H^- D$ down as far as we can. By similar reasoning we can assume that we arrive an interval I' which meets $\partial_H^+ D$ is a single non-trivial interval.

Combining these, we obtain a rectangle, R, in Π with $\partial_H^+ D \subseteq \partial_H^- R$ and with $\partial_H^- D \subseteq \partial_H^+ R$, where $\partial_H^+ R$ and $\partial_H^- R$ are the "top" and "bottom" sides of R. Indeed, without loss of generality, we can assume we are in one of the following two cases:

(i) $\partial_H^- D = \partial_H^+ R$ and $\partial_H^+ D$ is a subinterval of $\partial_H^- R$,

(ii) $\partial_H^- D$ is a subinterval of $\partial_H^+ R$ containing its right endpoint and $\partial_H^+ D$ is a subinterval of $\partial_H^- R$ meeting containing its left endpoint.

Note that we can find a simple closed curve, β, in Σ consisting of a vertical leaf of R together with an arc in D. This determines an element, $g \in \Gamma$, defined up to conjugacy. (We orient so that we run downwards along the vertical leaf.)

We aim to eliminate D by applying the move (M3) described in Section 4. That is, we want to construct an annulus, Ω, in Σ, containing D and meeting Π only in vertical leaves. This is best described in terms of the cover, $\tilde{\Sigma}$, of Σ.

We lift D to a digon, $\tilde{D}$ in $\tilde{\Phi}$, with $\partial_H^+ \tilde{D} \subseteq \partial_H^- \tilde{R}$, where $\tilde{R}$ is a lift of R. Thus, $\partial_H^- \tilde{D} \subseteq g\partial_H^+ \tilde{R}$. We aim to construct a $\langle g\rangle$-invariant strip, $\tilde{\Omega} \subseteq \tilde{\Sigma}$, which will descend to the required annulus, $\Omega \subseteq \Sigma$. We will write β^-, β^+ for its boundary components, and $\tilde{\beta}^-$ and $\tilde{\beta}^+$ for their lifts to $\tilde{\Sigma}$.

In the above, we can have either:

(i) $\partial_H^+ \tilde{R} = g^{-1}\partial_H^- \tilde{D}$ and $\partial_H^- \tilde{R} = I^- \cup \partial_H^+ \tilde{D} \cup I^+$, or

(ii) $\partial_H^+ \tilde{R} = I^- \cup g^{-1}\partial_H^- \tilde{D}$ and $\partial_H^- \tilde{R} = \partial_H^+ \tilde{D} \cup I^+$, where I^- and I^+ are intervals. Either way, I^+ is an interval in $\partial_H^- \tilde{R}$ with left endpoint in $\partial_H^+ \tilde{D}$ and containing the right endpoint of $\partial_H^+ \tilde{R}$.

If I^+ consists of a single point, when we can just set $\tilde{\beta}^+ = \partial_V^+ D \cup \partial_V^+ R$, so we can suppose that I^+ is a non-trivial interval.

We now push I^+ down as far as it will go, to an interval, J, meeting $\partial_H \tilde{\Pi}$ non-trivially. By a similar argument as earlier, we can assume that a $J \cap \partial_H \tilde{\Phi}$ consists of a single interval. (Otherwise we can split Π, decreasing $c(\Pi)$.) Thus $J \cap \partial_H^+ \tilde{\Phi} \subseteq \partial_H^+ \tilde{F}$. Thus, $\partial_H^+ \tilde{\Phi} \subseteq \partial_H^+ \tilde{F}$, where $\tilde{F}$ is a component of $\tilde{\Phi}$. Let $F \subseteq \Sigma$ be the projection. We distinguish two cases:

(a) $F \neq D$.

In this case, splitting Π if necessary, we can suppose that $J \subseteq \partial_H^+ \tilde{F}$. Now the left endpoint of I^+ is also the right endpoint of $\partial_H^+ \tilde{D}$, and the right endpoint of $\partial_H^- \tilde{D}$ is also the right endpoint of $g\partial_H^+ R$. This lies in the same vertical leaf of $\tilde{\Pi}$ as the left endpoint of J and the right endpoint of gI^+. Now gI^+ pushes down to $gJ \subseteq g\tilde{F}$. This means that there is a vertical interval, δ_1, from the left endpoint of J so some point, $x \in g\tilde{F}$. We connect x to the left endpooint of gJ by a path δ_1, lying in $g\tilde{F}$. Now $\delta_1 \cup \delta_2$ maps down to a simple closed curve, $\beta^+ \subseteq \Sigma$, meeting Π in a single leaf, and homotopic to the curve β described above.

(b) $F = D$.

We set $\tilde{F} = h\tilde{D}$, where $h \in \Gamma$. Note that $h \notin \langle g \rangle$. Let $K \subseteq J$ be the component of $J \setminus h\partial_H^+ \tilde{D}$ containing the left endpoint. Let $S \subseteq \tilde{\Pi}$ be the emptly rectangle with $\partial_H^+ S = K$. Note that $\partial_V^+ S = h\partial_V^- \tilde{D}$. Now the right endpoint of ∂I^+ pushes down into $h\partial_H^+ \tilde{R}$, so so the right endpoint of gI^+ pushes down into $gh\partial_H^+ \tilde{R}$. From this, it follows that the left endpoint of $\partial_H^- S$ lies in $gh\partial_H^+ \tilde{R}$. Thus, $\partial_H^- S$ meets both $gh\tilde{R}$ and $hg\tilde{R}$, and so $\partial_V^+ \tilde{R}$ lies in both $h\tilde{D}$ and $ghg^{-1}\tilde{D}$. It follows that $ghg^{-1}\tilde{D} = h\tilde{D}$ and so $ghg^{-1} = h$. In other words, g and h commute. Since g is primative and $h \notin \langle g \rangle$, it follows that Σ is closed torus, contrary to our assumptions.

In summary, we have found a simple closed curve β^+ in Σ homotopic to β, and meeting Π in a vertical leaf. Note that $\tilde{D}$ lies to the right of $\tilde{\beta}^+$ oriented downward, in the direction of translation of g.

We now perform a similar construction on the left of $\tilde{D}$. There are two possibilities: either

(i) $I^- \subseteq \partial_H^- R$, and we push it down as far as it will go, or

(ii) $I^- \subseteq \partial_H^+ R$, and we push it up as far as it will go.

Either way, as above we get a curve $\beta^- \subseteq \Sigma$ homotopic to β, again meeting Π in a single vertical leaf. They bound an immersed annulus Ω mapping to D with degree 1. (In fact, it's not hard to see that this has to be embedded in Σ.) We can now use the construction (M3) of Section 4 to give us $\Pi' = \Pi \cup \Omega$. In the process, we have removed the digon D, and so $c(\Pi') < c(\Pi)$.

After a bounded number of such operation we will eventually arrive at a binding Π'' with no monogons or digons. In other words, Π'' is efficient.

We started out, in addition, with a carrying graph, $\lambda \subseteq \Phi(\Pi)$, and a map, $\psi : \lambda \longrightarrow M$, in the right homotopy class, with $\psi(\lambda)$ of bounded length. To complete the picture, we need to explain how to get a carrying graph, $\lambda'' \subseteq \Phi(\Pi'')$ and a map $\psi'' : \lambda'' \longrightarrow M$.

Note that in each step of the above process, $\partial\Phi(\Pi')$ is a bounded distance from $\partial\Phi(\Pi)$. More precisely, in a splitting move, we attach a rectangle to $\Phi(\Pi)$ along one or two vertical sides. There given the bound on vertical size, there is no loss in assuming that $\lambda \cap \partial\Phi \subseteq \partial_H \Pi$, and so gives us also a graph $\lambda \subseteq \Phi(\Pi)$. Now, if the rectangle is attached along just one side, or is attached along both sides but to different components of Φ, then we can simply set $\lambda' = \lambda$ and $\psi' = \psi$. If the rectangle meets the same component of $\Phi(\Pi)$ along both its vertical sides, then we set $\lambda' = \lambda \cup \alpha$, where α is any vertical leaf of the rectangle. Since the $\tilde{\psi}$-image of a lift of the boundary of rectangle has bounded diameter, we can define $\psi' : \alpha \longrightarrow M$ so that $\tilde{\psi}'(\tilde{\alpha})$ has bounded diameter. We set $\psi'|\lambda = \psi$. This defines our map ψ' in this case. In the move that deletes a digon, there is no change to λ or ψ. Finally, in the case of addition of an annulus, $\Phi(\Pi)$ changes by the removal of a digon, and the conversion of a horizontal boundary segment of bounded length to a vertical segment. Therefore no change is necessary in this case.

After all these operations we end up with our efficient binding, Π'', the supporting graph on the complement, and a map $\psi'' : \lambda'' \longrightarrow M$. To simplify notation, we rename these as Π and λ and continue to the next step.

Step 3 :

We now have an efficient binding, Π, and a carrying graph, λ, on Φ, and a map $\psi : \Pi \cup \lambda \longrightarrow M$, as above.

Let $B(\Pi)$ and $S(\Pi)$ be the circular components of Π as defined in Section 4. We can assume that $B(\Pi)$ is a disjoint union of maximal vertical annuli, and that $\Phi(P)$ meets each component of $\partial B(P)$ in a single vertical edge of Φ. (This will simplify the subsequent discussion.)

Let Ω be a component of $B(\Pi)$, and let $\partial\Omega \cup \epsilon \sqcup \epsilon'$. Suppose that Ω is not isolated. Let $I = \epsilon \cap S(\Pi)$ and let $I' = \epsilon' \cap S(\Pi)$. Thus, I, I' are intervals, $\epsilon \setminus I$ and $\epsilon' \setminus I'$ and cusps of $\Phi(\Pi)$. Now $\gamma \cap \Omega$ consists of a set of crossing arcs. The map $\psi : \Omega \longrightarrow M$ sends each component of $\gamma \cap \Omega$ to a path minimal in its homotopy class relative to its endpoints. Note also that $\tilde{\epsilon}$ and $\tilde{\epsilon}'$ have fundamental domains whose $\tilde{\psi}$-images have bounded diameter. (Since $\phi(\Pi)$ has bounded horizontal size.) Thus, Lemma 3.1 gives us a track, ω, in Ω, with terminal points in $I \cup I'$, and carrying map, $p : \gamma \cap \Omega \longrightarrow \omega$, homotopic to inclusion relative to $I \cup I'$.

If Ω is isolated, let $z, z' \in \partial\Omega$ be terminal points of the carrying graph $\lambda \subseteq \Omega$. We now let $\omega \subseteq \Omega$ be the track given by Lemma 3.3, and a map $f : \omega \longrightarrow M$, sending $\delta(\omega)$ a curve of minimal length in its homotopy class. If $\beta = \gamma \cap \Omega \neq \emptyset$, we can take $\delta(\omega) = \gamma \cap \Omega$, and suppose that $f(\gamma \cap \Omega)$ curve β_M in M.

We now apply the construction of Section 4 to give us a track, τ in Σ. Let $\rho \subseteq \tau$ be the part arising from $S(\Pi)$. The components of $B(\pi)$ give rise to a special set, Σ, of loops in τ.

Note that ρ carries a metric, d_ρ, so that the carrying map, $(\gamma, d_\gamma) \longrightarrow (\rho, d_\rho)$ is locally ρ-bilipschitz, where μ depends only on the constants arising in Propostion 2.2, and hence ultimately, only on the hyperbolicity constant, and the toplogical type of Σ. Lemma 3.1 allows us to extend the path-metric to a path-metric, d_τ, on τ.

Each vertex of $\tau \cap S(\Pi)$ arose from a vertical leaf. Lifting to $\tilde{\Sigma}$, we can map each such vertex to any point in the $\tilde{\phi}$-images of the corresponding vertical leaf. We can now extend over τ, sending each edge to a geodesic segment. Descending to Σ, we get a map $f : \rho \longrightarrow M$. We have already have f defined on $\tau \cap B(\Pi)$. We set f to be equal Ψ on λ. This gives us our map, $f : \tau \longrightarrow M$.

In summary, we have constructed a track, τ, with path-metric d_τ and a special set of loops $\mathcal{S}$, a carrying map $p : \delta \longrightarrow \tau$ and a map $f : \tau \longrightarrow M$. We now verify that these satisfy (T1)–(T8) required by Theorem 0.1. The various statement in relation to the thin part follow from Lemma 3.1 and Lemma 3.3. The statements about the thick part follow from the construction of ρ.

6. Hyperbolic surfaces.

We now proceed with the proof of Proposition 2.2. In this section we discuss triangulations on Σ. We will use complete finite area structures on the complement of a finite set, $V \subseteq \Sigma$, which will be the vertex set of a triangulation. This will be conventient to formulate later arguments. We can use this to give a riemannian metric on Σ, as in Lemma 1.1, as disussed at the end of Section 8.

Let θ be an ideal hyperbolic triangle. We write $\omega(\theta)$ for its centre. If a is an edge of θ, we write $p(\theta, a)$ for the orthogonal projection of $\omega(\theta)$ to a, which we refer to as the *midpoint* of a. The three edge midpoints are connected by three horocyclic arcs, each of length $\sqrt{2}$. Each of these arcs cuts off a *spike* of θ. If ζ is one of the spikes, and $t \in (0, \sqrt{2}]$, we write $\delta(\zeta, t)$ for the horocyclic arc of length t in ζ. These arcs foliate ζ. We write $\Delta(\zeta, t) = \bigcup_{u \le t} \delta(\zeta, u)$ for the spike cut off by $\delta(\zeta, t)$. We write $\Delta(\theta, t)$ for the union of the three $\Delta(\zeta, t)$ as e ranges over the three spikes. Given an edge a, we write $I(\theta, a, t) = a \setminus \Delta(\theta, t)$. This is an open interval of length $2 \log(\sqrt{2}/t)$ centred at $p(\theta, a)$.

Now each spike of θ has two orthogonal foliations: one by horocycles, and one by geodesics rays. We refer to these as *vertical* and *horizontal* respectively. We can extend these smoothly and symmetrically to give two orthogonal foliations of θ with a tripod singularity at the centre, $\omega(\theta)$. We also refer to these extended foliations as *vertical* and *horizontal*.

Suppose for the moment, that Σ is a closed surface. (We discuss the case with boundary at the end of this section.) Let χ be a triangulation of Σ with vertex set $V = V(\chi) \subseteq \Sigma$. We write $E(\chi)$ for the set of edges and $T(\chi)$ for the set of triangles. Given $v \in V$, let $E(\chi, v)$ be the set of edges incident on v. This gives an ideal triangulation of the punctured surface, $S = \Sigma \setminus V$ and we give S a hyperbolic structure in which every triangle is an ideal hyperbolic triangle. To determine this uniquely, we need to specify shift parameters along the edges.

We suppose that Σ is oriented, so that any orientation on an edge a of χ also determines a transverse orientation. We write $\theta_- = \theta_-(a)$ and $\theta_+ = \theta_+(a)$ for the adjacent triangles on the negative and positive side of a respectively. We write $p_\pm(a) = p(\theta_\pm(a), a)$ for the respective edge midpoints. We write $\Lambda(a)$ for the signed distance between $p_-(a)$ and $p_+(a)$ along a. Note that the sign of $\Lambda(a)$ depends only on the orientation of Σ. This gives us a map $\Lambda : E(\chi) \longrightarrow \mathbf{R}$. Indeed any such map determines a hyperbolic structure on S. This structure will be complete if and only if for each $v \in V$, $\sum_{a \in E(\chi,v)} \Lambda(a) = 0$ (counting an edge twice if both endpoints are at v).

Suppose now that $\lambda : E(\chi) \longrightarrow \mathbf{R}$ is another map. Suppose that $a \in E(\chi)$ is oriented with adjacent triangles $\theta_{\pm}$. Let b_+, c_+ be the other edges of θ_+ adjacent to a in the positive and negative directions respectively, and let b_-, c_- be the other edges of θ_- adjacent to a in the negative and positive directions respectively. Let $\Lambda(a) = \frac{1}{2}(\lambda(b_+) + \lambda(b_-) - \lambda(c_+) - \lambda(c_-))$. Note that this does not depend on the orientation of a. Now $\sum_{a \in E(\chi,v)} \Lambda(a) = 0$ for each $v \in V$, and so any such λ will determine a complete hyperbolic structure, σ, on S.

Suppose that $\lambda(a) \geq 0$ for all a, and that λ satisfies a triangle inequality for each triangle, that is, if $\theta \in T(\chi)$ has edges a_1, a_2, a_3, then $\lambda(a_1) \leq \lambda(a_2) + \lambda(a_3)$. If $a \in E(\chi)$, then, using the above notation, we can define a point, $q_+(a) \in a$, a distance $\frac{1}{2}(\lambda(a) + \lambda(c_+) - \lambda(b_+))$ from $p_+(a)$ in the positive direction along a. We have chosen our structure such that this is the same as the point a positive distance $\frac{1}{2}(\lambda(a) + \lambda(b_-) - \lambda(c_-))$ from $p_-(a)$. We similarly define a point $q_-(a)$ a distance $\frac{1}{2}(\lambda(a) + \lambda(b_+) - \lambda(c_+))$ from $p_+(a)$ in the negative direction. We also note that if we orient c_+ so so that a and c_+ meet in the positive direction of c_+, then $q_+(a)$ and $q_+(c_+)$ are the endpoints of a horocyclic arc in θ_+.

Continuing around the vertex, we see that there is closed horocycle about v, meeting the edges of χ precisely in the points $q_{\pm}$. Let $D(v)$ be the horodisc bounded by this horocycle and set $D = \bigcup_{v \in V} D(v)$. In summary, we have:

Lemma 6.1 : *Suppose that $\lambda : E(\chi) \longrightarrow [0, \infty)$ satisfies the triangle inequalities for all triangles of χ. Then there is a complete hyperbolic structure on $S = \Sigma \setminus V$ and a horodisc, $D(v)$, about v for each $v \in V$ such that* $\mathrm{length}(a \setminus D) = \lambda(a)$ *for all $a \in E(\chi)$.* $\diamondsuit$

Note that, by construction, the points $p_{\pm}(a)$ all lie outside $\tilde{D}$.

We write $\tilde{\Sigma}$ for the universal cover of Σ, and $\tilde{V} \subseteq \tilde{\Sigma}$ for the preimage of v. Thus χ lifts to a triangulation $\tilde{\chi}$, of $\tilde{\Sigma}$ with vertex set $\tilde{V}$. Suppose that $h : \tilde{V} \longrightarrow H$ is any Γ-equivariant map. Given $a \in E(\tilde{\chi})$, with endpoints $x, y \in \tilde{V}$, we set $\lambda(a) = d(h(x), h(y))$. Thus, Lemma 6.1 gives us a hyperbolic structure, σ, on S, and a union of horodiscs $D \subseteq \Sigma$. Let $\tilde{D}$ be the preimage of D in $\tilde{\Sigma}$. If $a \in E(\tilde{\chi})$, then $\mathrm{length}(a \setminus \tilde{D}) = d(h(x), h(y))$, so we can define map, $\phi : a \setminus \tilde{D} \longrightarrow H$, sending $a \setminus \tilde{D}$ isometrically to some geodesic segment (of length $\lambda(a)$ in H from $h(x)$ to $h(y)$. Since Γ acts freely on Σ, this can be done Γ-equivariantly. If $\theta \in T(\tilde{\chi})$, we can extend over $\theta \setminus \tilde{D}$ so that each vertical leaf gets sent to a set of diameter at most H, where H depends only on k. (This is a simple consequence of the "thin triangles" characterisation of hyperbolicity.) We now get an equivariant map, $\tilde{\phi} : \theta \setminus \tilde{D} \longrightarrow H$. (Of course, $\tilde{\phi}$ need not be continuous.) We can now extend $\tilde{\phi}$ to $\tilde{\Sigma} \setminus \tilde{V}$ simply by mapping each horodisc to the corresponding point in $h(\tilde{V}) \subseteq H$.

To apply the above, let $\gamma \subseteq \Sigma$ be a multicurve, and let χ be a compatible triangulation (as in Section 2). Let γ_M be a realisation in M. We now take any map $V \longrightarrow M$ sending each $v(\alpha)$ to some point of its realisation, α_M, and let $h : \tilde{V} \longrightarrow \tilde{M} = H$ be an equivariant lift. This gives rise to a map $\tilde{\phi} : \tilde{\Sigma} \setminus V$ in the manner described above. If $\tilde{\alpha}$ is a component of $\tilde{\gamma}$ corresponding to the edge $a(\tilde{\alpha})$ of $\tilde{\chi}$ then $\tilde{\phi}(a(\tilde{\alpha}) \setminus \tilde{D})$ is an interval in $\tilde{\alpha}_M$. This is a fundamental domain for the $G(\tilde{\alpha})$-action, where $G(\tilde{\alpha})$ is the setwise stabliser of $\tilde{\alpha}$ in $\pi_1(\Sigma \setminus V)$.

We extend this discussion to the case where Σ has non-empty boundary. Here the map λ is still defined on $E(\chi)$, but Λ is not defined only on the set $E_I(\chi) \subseteq E(\chi)$ of interior edges of χ, i.e. those that don't constitute a boundary component of Σ. (The condition on Λ at the vertices only applies to those in the interior of Σ.) In this case, we obtain a hyperbolic structure on $\Sigma \setminus V$, where each boundary edge is a bi-infinite geodesic, with a spike at the vertex. Each such spike is a component of the set D. the corresponding vertex of χ.

In summary, we have shown:

Lemma 6.2 : *Suppose that $\gamma \subseteq \Sigma$ is a multicurve, and that χ is a compatible triangulation. There is a complete finite-area metric on $\Sigma \setminus V$ of the type described above, and an Γ-equivariant map, $\tilde{\phi} : \tilde{\Sigma} \setminus \tilde{V} \longrightarrow \tilde{M}$, sending each component of $\tilde{D}$ to be point in $\tilde{M}$. For each $a \in E(\chi)$, $a \setminus \tilde{D}$ gets mapped isometrically to a geodesic segment. Each vertical leaf in each triangle of χ gets mapped to a geodesic segment whose length in $\tilde{M}$ in bounded by some universal constant, $F \geq 0$. In the quotient, $\phi : \Sigma \setminus V \longrightarrow M$ extends to a map $\phi : \Sigma \longrightarrow M$, with $\phi(\gamma) = \gamma_M$.*

Topologically, ϕ collapses a regular neighbourhood of each component of each vertex to a point. It is continuous on the 1-skeleton of χ and on each vertical leaf.

7. The visual distance.

We recall the standard construction of a visual pseudometric based at some point, $p \in H$ (see [GhH]). This is usually described in relation to the boundary, ∂H, of H (where it is a metric) though it applies equally well to H itself.

Given $x, y, p \in H$, write $\langle x, y \rangle_p = \frac{1}{2}(d(p,x) + d(p,y) - d(x,y))$ for the "Gromov product". One variant of the definition of k-hyperbolicity asserts that for all $x, y, z, p \in H$, we have $\langle x, y \rangle_p \geq \min\{\langle x, z \rangle_p, \langle z, y \rangle_p\} - k$. Writing $s_p(x,y) = e^{-\langle x,y \rangle_p}$, this translates to the "quasiultrametric" condition: $s_p(x,y) \leq e^k \max\{s_p(x,z), s_p(z,y)\}$. We define $r_p(x,y) = \inf\{\sum_{i=1}^n s_p(x_{i-1}, x_i)\}$, where the infimum is taken over all sequences $x_0, x_1, \ldots, x_n$ in H with $x_0 = x$ and $x_n = y$. Clearly r_p is pseudometric on H, and $r_p(x,y) \leq s_p(x,y)$. Moreover, if $e^k \leq \sqrt{2}$, one can show [GhH] that $r_p(x,y) \geq (3 - 2\sqrt{2}) s_p(x,y)$. We shall therefore assume that $k \leq \frac{1}{2} \log 2$. Now the pseudometric r_p extends continuously to ∂H, where it is a metric.

Given two oriented geodesics a and b in H, we write $\rho_p = r_p(e_-(a), e_-(b)) + r_p(e_+(a), e_+(b))$, where e_- and e_+ denote the positive and negative endpoints (possibly in ∂H). This gives a pseudometric on the set of all geodesics.

We note that, from the construction, if $p, q \in H$, then for any geodesics, a, b, we have $\rho_q(a,b) \leq e^{d(p,q)} \rho_p(a,b)$.

We will need:

Lemma 7.1 : *There are universal constants $r > 0$ and $R \geq 0$ such that if a, b are geodesiscs in H and $p \in a$ with $\rho_p(a, b) \leq r$, then $d(p, b) \leq R$.*

Proof : It is a simple exercise in hyperbolic spaces to show that if $\langle e_-(a), e_-(b) \rangle_p$ and $\langle e_+(a), e_+(b) \rangle_p$ are both at least some universal constant K, then b must pass within some bounded distance, say, R, of p.

We now set $r = (3-2\sqrt{2})e^{-R_0}$. Thus if $\rho_p(a, b) \leq r$ then $e^{-\langle e_+(a), e_+(b) \rangle_p} = s_p(e_+(a), e_+(b)) \leq (3 + 2\sqrt{2}) r_p(e_+(a), e_+(b)) \leq (3 + 2\sqrt{2})r = e^{-R_0}$, and so $\langle e_+(a), e_+(b) \rangle_p \geq K$. Similarly, $\langle e_-(a), e_-(b) \rangle_p \geq K$, and so $d(p, b) \leq R$ as required. ♢

8. The proof of Propostion 2.2.

To simplify notation, we will rescale the metric by a factor of $2k/\log 2$, so that H is k_0-hyperbolic, where $k_0 = \frac{1}{2}\log 2$.

Let $\phi : \tilde{\Sigma} \setminus \tilde{V} \longrightarrow H$ be the Γ-equivariant map constructed in Section 6. In this section, it will be convenient to pass to the universal over, $\tilde{S}$, of $S = \Sigma \setminus V$, which is also the universal cover of $\tilde{\Sigma} \setminus \tilde{V}$. Let $G = \pi_1(S)$.

We write $\hat{\chi}$ for the lift of the triangulation, χ, to $\tilde{S}$. This is dual to a trivalent tree. Note that an orientation on an edge $a \in E(\hat{\chi})$ determines an orientation on all other edges in such a way that there are no cycles, and every maximal flow line starts at the initial vertex $e_-(a)$ and terminates at $e_+(a)$. We also write $\hat{D}$ for the preimage of D in $\subseteq \Sigma \setminus V$ in $\tilde{\Sigma}$. We partition $E(\chi)$ as $E_I(\chi) \sqcup E_\partial(\chi)$ into the interior and boundary edges.

Suppose that θ is a triangle of $\hat{\chi}$. Recall from Section 6 that $\Delta(\zeta, t)$ is the subset of a spike e, of θ bounded by a horocyclic arc, $\delta(\zeta, t) = \partial\Delta(\zeta, t)$ of length $t \leq \sqrt{2}$. This has area t. We write $\Delta(t)$ for the union of all $\Delta(\zeta, t)$ as ζ varies over all spikes of all triangles. Note that $\Delta(t)/G \subseteq S$ has area at most $3Nt$, where N is bounded in terms of the topological type of Σ.

If $a \in E(\hat{\chi})$ is oriented, we write $I_\pm(a, t)$ for the closure of $a \setminus \Delta(\theta_\pm(a), t)$. This is an interval of length $2\log(\sqrt{2}/t)$ centred on $p_\pm(a) = p(\theta_\pm(a), a)$.

By a *straight arc*, ϵ, in $\Delta(\zeta, t)$, we mean a segment, ϵ, connecting two points, $p \in a$ and $q \in b$, where $a, b \in E(\hat{\chi})$ are the edges bounding ζ, and such that ϵ is transverse to both the horizontal and vertical foliations, or possibly a vertical leaf. If this is oriented from p to q, we write $\delta(\epsilon)$ for the horocyclic arc with endpoint at p. We write $L(\epsilon)$ for the length of $\delta(\epsilon)$, and let $K(\epsilon) = \sigma(p', q)$, where $b \cap \delta(\epsilon) = \{p'\}$. These are the "vertical" and "horizontal" lengths of ϵ, respectively.

By a *zigzag path* in $\Delta(t)$, we mean a path, $\beta = \bigcup_{i=1}^n \epsilon_i$, where each ϵ_i is a straight arc, and such that adjacent straight arcs lie in different triangles. Thus β never enters the same triangle twice. It crosses some sequence of edges of $E(\hat{\chi})$, and, unless otherwise stated, we assume that these edges are oriented so that they are all crossed in the positive (transverse) direction. We write $L(\beta) = \sum_{i=1}^n L(\epsilon_i)$ for the vertical length, and $K(\beta) = \max\{K(\epsilon_i) \mid i \in 1, \ldots, n\}$ for the "horizontal shift".

We write $\hat{\phi} : \tilde{S} \longrightarrow H$ for the G-equivariant lift of ϕ to $\tilde{S}$. If $x \in \tilde{S}$ we write $\bar{x} = \hat{\phi}(x) \in H$. Similarly, we write $\bar{\beta} = \hat{\phi}(\beta)$ for any path, β, in $\tilde{S}$.

We describe a process for constructing zigzag paths in $\Delta(t)$, with bounded horizontal shift, and whose $\hat{\phi}$-images have bounded diameter. The following gives us something specific to aim for:

Lemma 8.1 : *There are constants, t_0, K_0, R_0, such that if $a \in E(\hat{\chi})$ is any oriented edge, then given any $t \in (0, t_0]$, there is a zigzag path, β, in $\Delta(t)$ emanating from $p_-(a)$ and terminating either in $I_+(b,t)$ for some $b \in E_I(\hat{\chi})$, or in b for some $b \in E_\partial(\hat{\chi})$, and with $K(\beta) < K_0$ and $\operatorname{diam}(\bar{\beta}) \le R_0$. Moroever, we can assume that $\beta \cap \hat{D} = \emptyset$.*

(Recall we are orienting β so that β never enters $\theta_+(b)$.)

For this statement, we need to allow the possibility that β just consists of a single point. (This is exceptional, in that this point need not lie in $\Delta(t)$.)

We begin with the following observation.

Lemma 8.2 : *There is some universal constant $\mu > 0$ such that if δ is a horocyclic arc of length $t \le 1$ in a triangle θ, with endpoints in edges a, b, then $\rho_{\bar{p}}(\bar{a}, \bar{b}) \le \mu t$, where $a \cap \delta = \{p\}$.*

Proof : Let $\epsilon, \zeta \subseteq \theta$ be horocyclic arcs meeting a and b respectively, far out the other spikes of θ. We can assume (from the construction of ϕ) that $\bar{\epsilon}$ and $\bar{\zeta}$ are just points of $\bar{a}$ and $\bar{b}$ respectively. By choosing them far enough away, we will have $\frac{1}{2}(\sigma(p,\epsilon) + \sigma(q,\zeta) - \sigma(\epsilon,\zeta))$ arbitrarily close to $\log(\sqrt{2}/t)$. Since $\hat{\phi}$ sends edges isometrically to geodesics, we also have $\frac{1}{2}(d(\bar{p},\bar{\epsilon}) + d(\bar{q},\bar{\zeta}) - d(\bar{\epsilon},\bar{\zeta}))$ arbitrarily close to $\log(\sqrt{2}/t)$. But, by construction of $\tilde{f}$, we have $d(\bar{p},\bar{q}) \le \operatorname{diam}(\bar{\delta}) \le F$ (the universal constant of Lemma 6.2). Thus $\langle \bar{\epsilon}, \bar{\zeta} \rangle_{\bar{p}} = \frac{1}{2}(d(\bar{p},\bar{\epsilon}) + d(\bar{p},\bar{\zeta}) - d(\bar{\epsilon},\bar{\zeta}))$ is equal to $\log(\sqrt{2}/t)$ up to an additive constant, F'.

But now $\rho_{\bar{p}}(\bar{a},\bar{b}) = r_{\bar{p}}(\bar{\epsilon},\bar{\zeta}) \le (3+2\sqrt{2}) s_{\bar{p}}(\bar{\epsilon},\bar{\zeta}) \le (3+2\sqrt{2}) e^{-\log(\sqrt{2}/t)+F'} = \mu t$, where $\mu = (3+2\sqrt{2})e^{F'}/\sqrt{2}$. $\diamondsuit$

Lemma 8.3 : *There are universal constants K_1, R_1, r_1, with the following property. Suppose that $t \in (0,1]$ and that $a \in E(\hat{\chi})$ is an oriented edge. Then there is a zigzag path, $\beta \subseteq \Delta(t)$, emanating from $p_-(a)$ with $K(\beta) \le K_1$ and $\operatorname{diam}(\bar{\beta}) \le R_1$, and satisfying either:*

(1) β terminates in $I_+(b,t)$ for some $b \in E_I(\hat{\chi})$, or in b for some $b \in E_\partial(\hat{\chi})$. or:

(2) $L(\beta) \ge r_1$.

Moreover, we have $\beta \cap \hat{D} = \emptyset$.

Proof : Let r, R be the constants of Lemma 7.1, let F be the constant of Lemma 6.2, and let μ be the constant of Lemma 8.2. We set $\lambda = \mu e^R$, $r_1 = r/\lambda$, $K_1 = 4R + 2$ and $R_1 = 2F + 4R + 2$.

We shall construct β inductively. Let us assume we have constructed a zigzag path $\beta = \epsilon_1 \cup \cdots \cup \epsilon_n \subseteq \Delta(t) \setminus \hat{B}$, where ϵ_i is a straight arc connecting $p_{i-1} \in a_{i-1} \in E(\tilde{\chi})$ to $p_i \in a_i \in E(\tilde{\chi})$, and with $p_0 = p_-(a)$. Thus $a_0 = a$ and β crosses, in turn, the edges $a_1, \ldots, a_{n-1}$, and terninates at $p_n \in a_n$.

We suppose that $K(\beta_n) \leq 2R+1$, that $d(\bar{p}_0, \bar{p}_i) \leq R$ for all i, that $\rho_{\bar{p}_0}(\bar{a}_0, \bar{a}_n) \leq \lambda L(\beta_n)$ and that $L(\beta_n) \leq r_1$. We need to decide what to do next. We distinguish two cases.

(1) $\sigma(p_n, I_+(a_n, t)) \leq 2R+1$.

Let q be the nearest point of $I_+(a_n, t)$ to p_n. (Possibly, $q = p_n$.) We let ϵ' be a straight are from p_{n-1} to q, and set $\beta = \epsilon_1 \cup \cdots \cup \epsilon_{n-1} \cup \epsilon'$, and stop.

(2) $\sigma(p_n, I_+(a_n, t)) > 2R+1$.

Let $\theta = \theta_+(a_n)$ and let δ be the horocyclic arc in γ with endpoint at p_n. Let u be the length of δ. Thus $u \leq t$ (since $\delta \subseteq \Delta(t)$). We again split into two cases.

(2a) $L(\beta_n) + u > r_1$.

We set $\beta = \beta_n \cup \delta$ and stop. Note that $L(\beta) = L(\beta_n) + u > r_1$.

(2b) $L(\beta_n) + u \leq r_1$.

We are assuming that $d(\bar{p}_0, \bar{p}_n) \leq R$. By Lemma 7.1, $\rho_{\bar{p}_n}(\bar{a}_n, \bar{a}_{n+1}) \leq \mu u$ and so $\rho_{\bar{p}_0}(\bar{a}_n, \bar{a}_{n+1}) \leq e^{d(\bar{p}_0, \bar{p}_n)} \rho_{\bar{p}_n}(\bar{a}_n, \bar{a}_{n+1}) \leq e^R \mu \lambda = \lambda u$. We are also assuming that $\rho_{\bar{p}_0}(\bar{a}_0, \bar{a}_n) \leq \lambda L(\beta_n)$ and so $\rho_{\bar{p}_0}(\bar{a}_0, \bar{a}_{n+1}) \leq \lambda L(\beta_n) + \lambda u = \lambda(L(\beta_n) + u) \leq \lambda r_1 \leq r$. Thus, by Lemma 7.1, we have $d(\bar{p}_0, \bar{a}_n) \leq R$. Choose $p_{n+1} \in a_{n+1} \setminus \hat{D}$ so that $d(\bar{p}_0, \bar{p}_{n+1}) \leq R$. Then $\sigma(q, p_{n+1}) = d(\bar{q}, \bar{p}_{n+1}) \leq d(\bar{q}, \bar{p}_{n+1}) + d(\bar{p}_n, \bar{p}_0) + d(\bar{p}_0, \bar{p}_{n+1}) \leq t + 2R \leq 2R+1$.

Let ϵ_{n+1} be a straight arc form p_n to p_{n+1}. Since $\sigma(p_n, I_+(a_n, t)) \geq R$ we see that $\epsilon_{n+1} \subseteq \Delta(t)$. Thus $\beta_{n+1} = \beta_n \cup \epsilon_{n+1}$ is a zigzag path in $\Delta(t)$. Moreover, $L(\beta_{n+1}) = L(\beta_n) + u \leq r_1$, $\rho_{\bar{p}_0}(\bar{a}_0, \bar{a}_{n+1}) \leq \lambda(L(\beta_n) + u) = \lambda L(\beta_{n+1})$ and $K(\beta_{n+1}) \leq 2R+1$.

In this case, we have verified the inductive hypotheses, and proceed to the next step.

We claim that the process must terminate. For if it continued indefinitely, we would get $L(\epsilon_n) \to 0$ as $n \to \infty$. But then the projection of β_n to S must go out the cusp of S, and so $\bar{\beta}_n$ goes out one of the rays, $\xi(x)$, in H, which we have disallowed.

We can thus assume we are in case (1) or case (2a), which correspond respectively to cases (1) and (2) of the statement of the lemma. We know that in both cases, $K(\beta) \leq 4R+2 = K_1$ and that $d(\bar{p}_0, \bar{p}_n) \leq F$ for all i. For each i, $\operatorname{diam} \bar{\delta}(\epsilon_i) \leq F$, the constant of Lemma 1.2, and so $\operatorname{diam}(\epsilon_i) \leq F + K(\epsilon_i) \leq F + (4R+2)$. It follows that $\operatorname{diam}(\bar{\beta}) \leq F + (F + 4R + 2) = 2F + 4R + 2 = R_1$. $\diamondsuit$

To deduce Lemma 8.1 from Lemma 8.3, we bring the G-action on $\tilde{S}$ into play. This will rule out case (2) of Lemma 8.1.

First we need another definition. If $\delta = \delta(e, t)$ is a horocyclic arc in $\theta \in T(\hat{\chi})$, and $t \leq e^{-1}$, we define the collar of δ to be $P(\delta) = \bigcup\{\delta(e, u) \mid t \leq u \leq t+1\}$, i.e. the 1-neighbourhood of δ in $\theta \setminus \Delta(\theta, t)$. This has area $(e-1)t$. (The argument would work for any uniform collar — we have chosen 1 just for notational convenience.) Given a zigzag path $\beta = \bigcup_i \epsilon_i \subseteq \Delta(t)$ with $t \leq e^{-1}$, we write $P(\beta) = \bigcup_i P(\delta(\epsilon_i))$. (Recall that $\delta(\epsilon_i)$ is the horcyclic arc with the same initial point as ϵ_i.) Thus, $P(\beta) \subseteq \Delta(et)$. Note that $\operatorname{area}(P(\beta)) = (e-1)L(\beta)$.

We now set $t_0 = \min\{(e-1)r_1/9\nu e, \sqrt{2}e^{-(2R_1+2)}\}$, where R_1, r_1 are the constants of Lemma 8.3. Recall that $\operatorname{area}(\Delta(et)/G) \leq 3\nu et \leq 3\nu r_1 e(e-1)/9\nu e = (e-1)r_1/3$.

Suppose now that $a \in E(\hat{\chi})$ and that β is a zigzag path emanating from $p_-(a)$ as given by Lemma 8.3, and suppose we are in case (2), i.e. $L(\beta) \geq r_1$. We want to derive a contradiction. Note that $P(\beta) \subseteq \Delta(et)$.

Now area $P(\beta) = (e-1)L(\beta) \geq (e-1)r_1 \geq 3\,\mathrm{area}(\Delta(t)/G)$. Thus, there must a point of $\tilde{S}$, where at least three G-images of $P(\beta)$ all meet. Now the corresponding G-images of β must all cross the horizontal leaf through that point, and so, at least two of them do so in the same direction. We can call these β and $\gamma = g\beta$ for some $g \in G \setminus \{1\}$. Since $P(\beta) \cap P(\gamma) \neq \emptyset$, we see that if $b \in E(\hat{\chi})$ is one of the neighbouring edges of $\hat{\chi}$, then $\sigma(x, y) \leq 1$, where $b \cap \beta = \{x\}$ and $b \cap \gamma = \{y\}$.

We now follow β and γ back to their respective initial points. In doing so, β and γ both cross some sequence of edges of $\hat{\chi}$. Let $c \in E(\hat{\chi})$ be the last edge for which these sequences coincide. Let $c \cap \beta = \{z\}$ and $c \cap \gamma = \{w\}$, and let $\theta \in T(\hat{\chi})$ be the triangle on the far side of c from b.

There are three possibilities. It may be that z is the initial point of β, so that $z = p(\theta, c)$. Since $g \neq 1$, in this case, w cannot be the initial point of w, and so $w \notin I(\theta, c, t)$. This means that $\sigma(z, w) \geq \log(\sqrt{2}/t)$. We may have the same situation with the roles of z and w interchanged, so again $\sigma(z, w) \geq \log(\sqrt{2}/t)$. Finally, it may be that neither z nor w is an initial point. Thus β and γ must diverge after c. In other words, z and w must lie on opposite sides of the interval $I(\theta, c, t)$. In this case, $\sigma(z, w) \geq 2\log(\sqrt{2}/t)$. In all cases, $\sigma(z, w) \geq \log(\sqrt{2}/t)$.

However, by Lemma 8.3, we have $d(\bar{x}, \bar{z}) \leq \mathrm{diam}(\bar{\beta}) \leq R_1$, and $d(\bar{y}, \bar{w}) \leq \mathrm{diam}(\bar{\gamma}) \leq R_1$. Also, $d(\bar{x}, \bar{y}) \leq 1$, and so $\sigma(z, w) = d(\bar{z}, \bar{w}) \leq 2R_1 + 1$. Thus $\log(\sqrt{2}/t) \leq 2R_1 + 1$, so $t \geq \sqrt{2}e^{-(2R_1+1)} \geq t_0$. But we chose $t \leq t_0$ giving a contradiction.

Thus, case (2) of Lemma 8.3 cannot arise in this situation, and so we are in case (1). In other words β terminates in $I_+(b, t)$ for some edge $b \in E(\hat{\chi})$.

This proves Lemma 8.1.

We now apply Lemma 8.1 G-equivariantly to each pair (θ, a), where $\theta \in T(\hat{\chi})$ and $a \in E(\hat{\chi})$ is and edge of θ. This gives us a G-invariant set, $B(t)$, of zigzag arcs. We can assume that all intersections are transverse crossings. If $\beta, \gamma \in B(t)$ intersect at some point x, we say the intersection is *positive* if β, γ cross the horizontal leaf through x in the same direction, and *negative* if they cross it in opposite directions.

We want to eliminate all intersections. For this, we need to assume that we chose t even smaller. Let $t_1 = \min\{t_0, \sqrt{2}e^{-(2R_0+2K_0)}\}$ and $t_2 = \min\{t_1, \sqrt{2}e^{-(4R_0+4K_0)}\}$.

Lemma 8.4 : *If $t \leq t_1$, then the arcs of $B(t)$ have no positive intersections.*

Proof : Suppose β, γ intersect positively at some point of $\delta(e, t)$ for a spike e. Then β and γ both intersect b at points x and y respectively. Now $\sigma(x, y) \leq K(\beta) + K(\gamma) \leq 2K_0$.

We now follow β and γ backwards to two points z and w, as in the earlier argument. This time we get $\log(\sqrt{2}/t) \leq \sigma(z, w)$, so $t > \sqrt{2}e^{-(2R_0+2K_0)} \geq t_1$. ♢

Lemma 8.5 : *If $t \leq t_2$, then each arc of $B(t)$ meets at most one other.*

Proof : Suppose $\alpha, \beta, \gamma \in B(t)$ are distinct, and that α meets β at x and γ at y. We can assume that x precedes y along α.

We follow β backwards and γ forwards from x, again considering the sequence of edges that they both cross. Let a be the last edge of this sequence. Now, since $t \leq t_1$, we get a contradiction, exactly as in Lemma 8.4, unless α terminates in a. But this means that β must cross the horizontal leaf through y, Let b be the first edge after this leaf. Then b meets α, β, γ in points, u, z, w respectively.

Now $\sigma(u, w) \leq K(\alpha) + K(\gamma) \leq 2K_0$, and (as in the proof of Lemma 8.4) we get $\sigma(z, w) \leq 2K_0 + 2R_0$. Thus, $\sigma(z, w) \leq 4K_0 + 2R_0$.

We now follow β and γ backwards from z and w respectively. As in Lemma 8.4 (with z, w now playing the roles of x, y) we get a contradicton if $\log(\sqrt{2}/t) \leq (4K_0+2R_0)+2R_0 = 4K_0 + 4R_0$, in particular, if $t \leq t_2 \leq \sqrt{2}e^{-(4K_0+4R_0)}$. $\diamondsuit$

We now set $B = B(t_2)$.

We now modify B to give a new set, A, of transverse arcs which are all disjoint. We need to do this equivariantly, and without eliminating any of the initial points. Since, by Lemma 8.5, the elements of B intersect at most in pairs, it is enough to deal with these arcs pairwise.

Suppose $\beta, \gamma \in B$ intersect at some point x. By Lemma 8.4, this is a negative intersection. First note that β and γ lie in different G-orbits. (For if $\gamma = g\beta$, then we must have $g^2\beta = \beta$, giving the contradiction that g has order 2.) Let α be the concatenation of the segments of β and γ preceding x. Note that $\operatorname{diam}\bar{\alpha} \leq \operatorname{diam}\bar{\beta} + \operatorname{diam}\bar{\gamma} \leq 2R_0$. We now replace $\{\beta, \gamma\}$ by the single arc, α, carrying out this construction equivariantly for all such pairs.

Now A projects to a set of disjoint embedded transverse arcs in S. We extend these transverse arcs to a transverse graph, Υ. Suppose the endpoint of an arc lies in an interior edge of χ. We connect this endpoint to centre of the triangle containing the endpoint, and on the far side of the arc. If θ is such a triangle, and the arc terminates as some point, q in the edge, a, of θ, then by construction, $a \in I(\theta, a, t_2)$. Thus we can take the extension to lie in $(\theta \setminus \Delta(t_2))/G$ to S. Note that (back in the cover, $\tilde{S}$, the $\hat{\phi}$-image of $\hat{\theta} \setminus \hat{\Delta}(t_2)$ has bounded diameter bounded by $2\log(\sqrt{2}/t_2) + F$, where F is the bound given by Lemma 6.2. Thus, the diameters of the images of the extensions are bounded.

The map ϕ lifts to a map $\tilde{\phi} : \tilde{\Sigma} \longrightarrow \tilde{M}$. The map $\hat{\phi}$ factors through $\tilde{\phi}$, via the natural maps $\tilde{S} \longrightarrow S \hookrightarrow \Sigma$. The graph Υ lifts to a graph $\tilde{\Upsilon} \subseteq \tilde{\Sigma}$. The $\tilde{\phi}$-image of any edge of $\tilde{\Upsilon}$ in $\tilde{M}$ had diameter bounded above (by $2\log(\sqrt{2}/t_2) + 2F + 2R_0$).

We also note that the number of edges of Υ is bounded above (by some number n_0) depending only on the topological type of Σ.

For Proposition 2.2, we need another observation:

Lemma 8.6 : *Let $\eta > 0$. In the above, we can construct Υ so that if $a \in E(\chi)$ with* $\operatorname{length}(a \setminus D) \geq 2\eta$, *then each component of $a \setminus (\Upsilon \cup D)$ has length at least η.*

Proof : We need to adjust the constructions at a few points, and modify the relevant constants.

In the proof of Lemma 8.5, we can always choose $p_{n+1} \in a_{n+1}$ so that $d(p_{n+1}, \partial\hat{D}) > \eta$ (at the cost of increasing the constant R by η). This ensures that no arc of $B(t)$ can cross any such edge within distance η of $\hat{D}$.

We can strengthen Lemma 8.5 to say that no two arcs of $B(t)$ can cross any adge of $\hat{\chi}$ in the same direction at points closer than η. We can also strengthen Lemma 8.5 to say that any arc of $B(t)$ can come within distance η of at most one other arc of $B(t)$. The arguments are essentially unchanged.

After projecting to Σ, we can arrange that no two arcs cross the same edge within distance η of each other. If they do, we we can replace them by a single arc consisting of two initial segments and a short arc close to that edge.

We now construct the transverse graph Υ from the resulting set of arcs as before. $\diamondsuit$

To derive Proposition 2.2 in the form stated, we can collapse each component of D to a point, and take some riemannian approximation to the resulting metric.

References.

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

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

[Bow1] B.H.Bowditch, *Surface group actions and length bounds* : preprint, Warwick 2010.

[Bow2] B.H.Bowditch, *Coarse hyperbolic models for 3-manifolds* : preprint, Warwick 2010.

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

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

[PH] R.C.Penner, J.L.Harer, *Combinatorics of train tracks* : Annals of Mathematics Studies, No. 125, Princeton University Press (1992).

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