Surface group actions and length bounds

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

0. Introduction.

Let Σ be compact orientable surface, and write $\Gamma = \pi_1(\Sigma)$. Let H be a hyperbolic space (in the sense of Gromov [Gr]) admitting a discrete action by Γ. and write $M = H/\Gamma$. Let $X(\Sigma)$ be the set of homotopy classes of essential non-peripheral closed curves in Σ. Let $\mathcal{G}(\Sigma)$ be the curve graph with vertex set $V(\mathcal{G}) = X(\Sigma)$ (i.e. the 1-skeleton of the curve complex [H]). To each $\gamma \in X(\Sigma)$ we can associate a "stable length" $l^S_M(\gamma)$. In [MaM2], the authors define the notion of a "tight geodesic" in $\mathcal{G}(\Sigma)$. The aim of this paper is to describe certain bounds on the stable lengths of curves arising in a tight geodesic. These are a direct generalisation of the "a-priori bounds" described by Minsky in [Mi], where H is hyperbolic 3-space. Here we generalise the approach in [Bo2]. One of the main motivations of this work is its use in describing "coarse models" of hyperbolic 3-manifolds, see [Bo4].

If H is simply connected, we can think of this in terms of the quotient $M = H/\Gamma$. Then $l^S_M(\gamma)$ is, up to additive constant, equal to the length of a shortest (or nearly shortest) representative of the free homotopy class of γ in M. In fact, there is no essential loss of generality in assuming H to be simply connected, as we will explain later.

To be more precise, we make an assumption on the space H, namely (H1) which corresponds to coarse bounded geometry, and a hypothesis on the action, namely (H2) which substitutes for the Margulis lemma. (It is effectively the conclusion of the Margulis lemma if H happens to be a simply connected manifold of pinched negative curvature.)

Recall that a subset, $Q \subseteq H$ is *r-separated* if $d(x, y) \geq r$ for all distinct $x, y \in Q$. We assume:

(H1) $(\exists s_0)(\forall r \geq s_0)(\exists n)$ if $Q \subseteq H$ is r_0-separated and $\operatorname{diam}(Q) \leq r$, then $|Q| \leq n$.

(For example, this holds if H is a manifold of bounded Ricci curvature.)

Given $x \in H$ and $r \geq 0$, write $\Gamma_r(x) = \{g \in \Gamma \mid d(x, gx) \leq r\}$. We assume:

(H2) There exist $r_0 \geq s_0$ and ν_0 such that for all $x \in H$, either $|\Gamma_{r_0}(x)| \leq \nu_0$ or $\langle \Gamma_{r_0}(x) \rangle$ is infinite cyclic.

Clearly (H1) then holds for r_0 also, so we may as well set $s_0 = r_0$.

Let $Y(\Sigma)$ be the set of homotopy classes essential closed curves in Σ. Thus $X(\Sigma) \subseteq Y(\Sigma)$. We write $Y(\partial\Sigma) \subseteq Y(\Sigma)$ for the (possibly empty) set of boundary components. Each $\gamma \in Y(\Sigma)$ determines an element $g \in \Gamma$ defined up to conjugacy. Its *stable length* is defined as $||g|| = \lim_{n \to \infty} \frac{1}{n} d(x, g^n x)$ for some, hence any, $x \in H$. We set $l^S_M(\gamma) = ||g||$. We will $l^S_M(\partial\Sigma) = \sum\{l^S_M(\gamma) \mid \gamma \in Y(\partial\Sigma)\}$. (This is 0 if $\partial\Sigma = \emptyset$.)

It will be convenient to view the following as another hypothesis on the action. We fix some $L \geq 0$ and suppose:

(H3) $l^S_M(\partial\Sigma) \leq L$.

Definition : By the *parameters of the action* we mean collectively the hyperbolicity constant of H, the topological type of Σ, the constant r_0 and function $[r \mapsto n(r)]$ featuring in (H1), the constants r_0 and ν_0 featuring in (H2) and the constant L featuring in (H3).

(Recall that we have set $s_0 = r_0$ for convenience of notation.)

Given $l \geq 0$, let $Y(M,l) = \{\gamma \in Y \mid l^S_M(\gamma) \leq l\}$. We write $X(M,l) = X(\Sigma) \cap Y(M,l)$.

The notion of a "tight geodesic" we use here is that of [Bo2] which is a slight generalisation of the original in [MaM]. See Section 7 for further discussion.

Theorem 0.1 : *Suppose $\Gamma = \pi_1(\Sigma)$ acts on a hyperbolic space H satisfying (H1), (H2), (H3) above. Suppose that $\gamma_0\gamma_1\cdots\gamma_n$ is a tight geodesic in $\mathcal{G}(\Sigma)$ with $\gamma_0, \gamma_p \in X(M,l)$. Then $\gamma_i \in X(M,l')$ for all $i \in \{0,\ldots,n\}$, where l' depends on l and the parameters of the action.*

If $M \cong \mathrm{int}(\Sigma) \times \mathbf{R}$ is a complete hyperbolic 3-manifold, then l^S_M is just the length of the corresponding closed geodesic, γ_M, in M (or 0 if this is parabolic). Thus, in this case, Theorem 0.1 is just a variant of Minsky's a-priori bounds theorem [Mi]. The formulation we give here ties in with that in [Bo2] (cf. Theorem 1.3 of that paper.)

Theorem 0.1 divides into two cases, namely when Σ is exceptional (that is, a one-holed torus or four-holed sphere) or non-exceptional. The proofs are quite different, and the latter occupies most of this paper. The former is largely independent and described in Section 8. In this case, $\mathcal{G}(\Sigma)$ is a Farey graph, and every geodesic is deemed tight.

In the course of the proof, we show:

Theorem 0.2 : *Suppose $\Gamma = \pi_1(\Sigma)$ acts on a hyperbolic space, H satisfying (H1), (H2) and (H3). Then there exist $l_0, h_0 \geq 0$, depending only on the parameters of the action such that $X(M,l_0)$ is non-empty and h_0-quasiconvex in $\mathcal{G}(\Sigma)$. Moreover, if $l \geq l_0$, then there is some $h_1 \geq 0$ depending only in l and the parameters of the action such that $X(M,l) \subseteq N(X(M,l_0),h_1)$.*

Here $N(.,r)$ denotes the r-neighbourhood.

The following variant of Theorem 0.1 will also prove useful in applications (cf. Theorem 1.4 of [Bo2].)

Theorem 0.3 : *Suppose Γ acts on H satisfying (H1), (H2), (H3). There is some r_0 depending only or the parameters of the action such that $\gamma_0,\ldots,\gamma_p$ is a tight geodesic in $\mathcal{G}(\Sigma)$ and $r + r_0 \leq i \leq p - r - r_0$, where $r = \max\{d(\gamma_0, X(M,l)), d(\gamma_p, X(M,l))\}$, then $\gamma_i \in X(M,l')$, where l' depends only on the parameters of the action and l.*

(Note we can always take $l \geq l_0$, so that $X(M,l_0) \neq \emptyset$.)

In combination with the above, this gives "a-priori bounds" for "hierarchies" in $\mathcal{G}(\Sigma)$ in the same manner as described in [Bo2]. Such bounds for 3-manifolds were originally described in [Mi], and were central to the proof of the Ending Lamination Conjecture in [Mi,BrCM]. The application to hierarchies will be discussed in Section 9.

In order to simplify the exposition, we will make a few additional assumptions on M, namely:

$(*)$ H is a simply connected geodesic space, and Γ acts freely on H. Moreover, for each $g \in \Gamma$, the minimum $|g| = \min\{d_H(x, gx) \mid x \in H\}$ is attained.

We can do this without any loss of generality as far as the main results are concerned. It is sufficient to note that any hyperbolic space admitting any Γ-action is equivariantly quasi-isometric to one satisfying $(*)$. (This is described in [Bo3] using a variation on the Rips complex construction.) Moreover the hypotheses and conclusions are essentially quasi-isometry invariant. Alternatively one could readily reinterpret the arguments presented here without using $(*)$, though they become more clumsy to formulate.

Assuming $(*)$, we see that $\pi_1(M) \equiv \Gamma$, and we can associate to each $\gamma \in Y(\Sigma)$ a free homotopy class of curves in M. Such a homotopy class admits a minimal length representative, γ_M, (not necessarily unique). We write $l_M(\gamma) = \mathrm{length}(\gamma_M) = |g|$, where $g \in \Gamma$ is corresponding element (or conjugacy class) in Γ.

We note that $l^S_M(\gamma) \leq l_M(\gamma)$, and that $l_M(\gamma) - l^S_M(\gamma)$ is bounded above by some fixed multiple of the hyperbolicity constant. Thus, in the above theorem we could equally well replace l^S_M by l_M.

Definition : We refer to a space M arising in the manner above (for an action of $\Gamma = \pi_1(\Sigma)$ on a hyperbolic space H satisfying (H1), (H2), (H3) and $(*)$ as a *coarse hyperbolic manifold* We refer to the parameters of the action also as the *parameters of* M.

The proofs of the main results follow a similar strategy to that in [Bo2]. This was originally inspired by an idea of Bestvina and Fujiwara, which in turn was inspired by an argument of Luo. The basic idea is as follows.

Suppose we have a tight geodesic, $(\gamma_i)_{i=0}^p$ in $\mathcal{G}(\Sigma)$ whose elements are realised as shortest curves $(\gamma_i)_M$, in M. Suppose that the realisations of γ_0 and γ_p have bounded length, but those of some intermediate γ_i are very long. There is a sense in which these γ_i "fill out" certain subsurfaces of Σ. Tightness implies that consecutive surfaces overlap is such a way that we can find curves in them that shortcut the geodesic $(\gamma_i)_i$, thereby giving a contradiction. In [Bo2], these subsurfaces were described using pleating surfaces. Here, we will use a coarse analogue of that construction.

This construction is motivated by the following observation. Suppose that σ is a (constant curvature) hyperbolic metric on Σ, and that γ is a simple closed curve realised as a closed geodesic, γ_σ, in (Σ, σ). One can give a topological description of a metric neighbourhood, N, of γ_σ in (Σ, σ) along the following lines. Suppose that ϵ is an arc in Σ meeting γ exactly at its endpoints — that is a "bridge arc". If ϵ can be realised as a short arc in (Σ, σ), then at each of its endpoints there are long arcs in γ_σ which follow each either within a bounded distance. Properly quantified, the neighbourhood, N, is determined by the set of such arcs, and can thus be described without explicit reference to the small scale structure of (Σ, σ). We only need to have things defined up to a bounded distance.

We can perform an analogous construction with realisations of curves in M, instead. Arcs in $(\gamma_i)_M$ which remain close over a large distance determine a certain class of bridge arcs, which in turn determine a subsurface of Σ, defined up to homotopy in Σ. If γ_i and γ_{i+1} are consecutive elements of a tight geodesics whose realisations in M are sufficently long, then these subsurfaces overlap. This gives rise to a sequence of curves that shortcut the geodesic $(\gamma_i)_i$.

Much of the work in this paper is involved in describing this process more formally, and checking that it gives rise to something sensible. In [Bo2] this used pleating surfaces, and the uniform injectivity theorem. Here we will need a coarse analogue of this (see Section 6). This in turn uses the fact, from [Bo3], that simple closed curves realised in M follow close to a train track in M.

The above discussion applies to the case where Σ is "non-execptional". The exceptional cases of the one-holed torus and four-holed sphere will be dealt with by in independent argument in Section 8.

1. The thick-thin decomposition.

We describe how (H1) and (H2) give rise to a thick-thin decomposition of M analogous to that of a hyperbolic manifold.

If γ is a closed curve Σ, we write $\tilde{\gamma}$ for the entire preimage of γ in $\tilde{\Sigma}$. We will use $\breve{\gamma}$ for a component of $\tilde{\gamma}$, that is a bi-infinite arc. This is stabilised by some infinite cyclic subgroup, $G(\breve{\gamma}) = \langle g \rangle \subseteq \Gamma$, where g represents (the conjugacy class of) γ in $\Gamma = \pi_1(\Sigma)$. Note that every cyclic subgroup of Γ has this form.

We write $Y_0(\Sigma) \subseteq Y(\Sigma)$ for the set of primitive closed curves (so that $X(\Sigma) \subseteq Y_0(\Sigma)$) and write $\tilde{Y}_0(\Sigma) \subseteq \tilde{Y}(\Sigma)$ for the set of lifts of these to $\tilde{\Sigma}$. Note that $\breve{\gamma} \in \tilde{Y}_0$ if and only if $G(\breve{\gamma})$ is a maximal cyclic subgroup of Γ.

Now suppose Γ acts on hyperbolic space H satisfying (H1) and (H2) as described in Section 0. In particular, $\pi_1(M) \equiv \Gamma$, where $M = H/\Gamma$. We describe the thin part of M.

Recall that $\Gamma_r(x) = \{g \in \Gamma \mid d(x, gx) \leq r\}$. Given $r \geq 0$, $\nu \in \mathbf{N}$ and $G \leq \Gamma$, we set $\tilde{\Delta}(G; r, \nu) = \{x \in H \mid |G \cap \Gamma_r(x)| \geq \nu\}$. We abbreviate $\tilde{\Delta}(r, \nu) = \tilde{\Delta}(\Gamma; r, \nu)$. Thus $\tilde{\Delta}(r, \nu)$ is Γ-invariant and we set $\Delta(r, \nu) = \tilde{\Delta}(r, \nu)/\Gamma \subseteq M$.

If $G \leq \Gamma$ is any subgroup, then for all r, ν, we have $N(\tilde{\Delta}(G; r, \nu), s) \subseteq \tilde{\Delta}(G; r + 2s, \nu)$. This follows directly from the triangle inequalities. Moreover, applying (H1), we see that if $r \geq r_0$ and $\nu \in \mathbf{N}$ there is some $\nu' \in \mathbf{N}$ such that for any G, $\tilde{\Delta}(G; r, \nu') \subseteq \tilde{\Delta}(G; r_0, \nu)$. Putting these together, we obtain:

Lemma 1.1 : *For all $r \geq r_0$, $\nu \in \mathbf{N}$ and $s \geq 0$, then there is some $\nu' \in \mathbf{N}$ such that for all $G \leq \Gamma$ we have $N(\tilde{\Delta}(G; r, \nu'), s) \subseteq \tilde{\Delta}(G; r_0, \nu)$.* $\diamondsuit$

Now given any $\breve{\gamma} \in \tilde{Y}_0(\Sigma)$, we set $\tilde{\Delta}(\breve{\gamma}; r, \nu) = \tilde{\Delta}(G(\breve{\gamma}); r, \nu)$. Given any $\gamma \in Y_0$, we set $\Delta(\breve{\gamma}; r, \nu)$ for the image of $\tilde{\Delta}(G; r, \nu)$ in M, where $\breve{\gamma}$ is any lift of γ.

Lemma 1.2 : *If $\nu \geq \nu_0$, then $\tilde{\Delta}(r_0, \nu)$ is a disjoint union of $\tilde{\Delta}(\breve{\gamma}; r_0, \nu)$ as $\breve{\gamma}$ ranges over $\tilde{Y}$.*

Note, in particular, that this implies that $\Delta(r_0, \nu) = \bigsqcup_{\gamma \in Y_0} \Delta(\gamma; r_0, \nu)$.

Proof : Clearly $\bigcup_{\breve{\gamma} \in \tilde{Y}_0} \tilde{\Delta}(\breve{\gamma}; r_0, \nu) \subseteq \tilde{\Delta}(r_0, \nu)$. For the reverse inclusion, suppose $x \in \Delta(r_0, \nu)$. Thus $|\Gamma_{r_0}(x)| \geq \nu \geq \nu_0$, and so, by (H2), $\langle \Gamma_{r_0}(x) \rangle$ is infinite cyclic, and hence contained in a maximal infinite cyclic subgroup, $G(\breve{\gamma})$ for some $\breve{\gamma} \in \tilde{Y}_0$. Thus $x \in \tilde{\Delta}(\breve{\gamma}; r_0, \nu)$ as required.

For disjointness, suppose $\breve{\gamma}, \breve{\delta} \in \tilde{Y}_0$, with $x \in \tilde{\Delta}(\breve{\gamma}; r_0, \nu) \cap \tilde{\Delta}(\breve{\delta}; r_0, \nu)$. Now $\langle \Gamma_{r_0}(x) \rangle$ is contained in a maximal infinite cyclic subgroup $G \leq \Gamma$ and meets both $G(\breve{\gamma})$ and $G(\breve{\delta})$ non-trivially. It follows that $G(\breve{\gamma}) = G(\breve{\delta}) = G$, and so $\breve{\gamma} = \breve{\delta}$. $\diamondsuit$

We shall abbreviate $\Delta_0 = \Delta(r_0, \nu_0)$, $\tilde{\Delta}_0 = \tilde{\Delta}(r_0, \nu_0)$, $\Delta_0(\gamma) = \Delta(\gamma; r_0, \nu_0)$ and $\tilde{\Delta}_0(\breve{\gamma}) = \tilde{\Delta}_0(\breve{\gamma}; r_0, \nu_0)$.

Note that if $\nu \geq \nu_0$ and $\gamma \in Y_0$, then $\Delta(\gamma; r_0, \nu) = \Delta_0(\gamma) \cap \Delta(r_0, \nu)$ and $\tilde{\Delta}(\breve{\gamma}; r_0, \nu) = \Delta_0(\breve{\gamma}) \cap \tilde{\Delta}(r_0, \nu)$.

We remark that one can show that the sets $\tilde{\Delta}(\breve{\gamma}; r_0, \nu)$ are uniformly quasiconvex in $\tilde{M}$, though we will not be explicitly using that in this paper.

2. The track construction.

We recall some consequences of the track construction described in [Bo3].

We fix some ν_0 as given by (H1), (H2), and let $\Delta_0 \subseteq M$ be the corresponding thin part as described in Section 1.

Now let $\gamma \subseteq \Sigma$ be a multicurve, and let $q : \gamma \longrightarrow M$ be a realisation, so that $q(\gamma)$ has minimal length in its homotopy class. We write d_γ for the induced path-metric on γ. We say that a component, α, of γ is *short* if its length, $l_M(\alpha) = \mathrm{length}(q(\alpha))$ is less than some fixed constant, which for convenience we can set to be the hyperbolicity constant of $\tilde{M}$. This means that if α is not short, then each lift $\tilde{q} : \breve{\alpha} \longrightarrow \tilde{M}$ is uniformly quasigeodesic (with constants depending only on the hyperbolicity constant).

By a *multitrack* we mean a train track, τ, in the sense of Thurston, though dropping the condition that it be connected. In other words, it is a graph embdded in the interior of Σ, such that there are no nullgons, monogons, digons or smooth annuli in the complement. (For terminology, see for example, [PH].)

We recall the following definitions from [Bo3]. 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.

Now the construction of [Bo3] gives us a multitrack, $\tau \subseteq \Sigma$, a path-metric d_τ on τ, a carrying map $p : \gamma \longrightarrow \tau$, a map $f : \tau \longrightarrow M$ in the correct homotopy class, a (possibly empty) collection, $\mathcal{S}_0 = \mathcal{S}$, of simple loops in τ and associated simple annuli, $A_0(\delta) = A_0(\delta, \tau)$ for $\delta \in \mathcal{S}_0$, satifying various geometric conditions. Writing $A_0(\Sigma) = \bigcup_{\delta \in \mathcal{S}_0} A_0(\delta)$, we have:

(A1) $p : (\gamma, d_\gamma) \longrightarrow (\delta, d_\tau)$ is ξ_0-lipschitz.

(A2) For all $x \in \tilde{\gamma}$, $d(\tilde{q}(x), \tilde{f}\tilde{p}(x)) \leq K_1$.

(A3) $\tau \setminus A_0(\Sigma)$ has length at most K_2.

(A4) $f|\tau \setminus \bigcup \mathcal{S}_0 \longrightarrow (M, d)$ is ξ_1-lipschitz.

(A5) For each $\delta \in \mathcal{S}_0$, $f(A_0(\delta)) \subseteq \Delta_0(\delta)$.

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

For suitably chosen $\nu_2 \geq \nu_1 \geq \nu_0$ we set $\Delta_1 = \Delta(r_0, \nu_1)$ and $\Delta_2 = \Delta(r_0, \nu_2)$ so that $\Delta_2 \subseteq \Delta_1 \subseteq \Delta_0$.

In addition, we can find $\mathcal{S}_2 \subseteq \mathcal{S}_1 \subseteq \mathcal{S}_0$, and special annuli $A_1(\delta) = A_1(\delta, \tau) \subseteq A_0(\delta)$ for each $\delta \in \mathcal{S}_1$, and $A_2(\delta) = A_2(\delta, \tau) \subseteq A_1(\delta)$ for each $\delta \in \mathcal{S}_2$, with the following properties:

(A6) $\tau \setminus A_2(\Sigma)$ has length at most K_2'.

(A7) $f|\tau \setminus \bigcup \mathcal{S}_2 \longrightarrow M$ is ξ_1'-lipschitz.

(A8) For each $\delta \in \mathcal{S}_2$, $f(A_2(\delta)) \subseteq \Delta_2(\delta)$.

(A9) For each $\breve{\delta} \in \tilde{\mathcal{S}}_1$, $f^{-1}(\Delta_1(\breve{\delta})) \subseteq A_1(\breve{\delta})$.

Here the new constants depend also on the choice of ν_1 and ν_2.

In fact, we can arrange that $N(\Delta_2, s) \subseteq \Delta_1$ and $N(\Delta_1, s_0) \subseteq \Delta_0$, where $s_0, s \geq 0$, are chosen sufficiently large as described later. (In practice, we will be using two different choices of Δ_1 and Δ_2, in Sections 3 and 7 respectively. Ultimately, however, these choices will depend only the parameters of M.)

Note that a consequence of (A5) is that if $\breve{\delta} \in \tilde{\mathcal{S}}_0$, then $\tilde{f}(A_0(\delta)) \subseteq \tilde{\Delta}_0(\breve{\delta})$.

The above statements, (A1)–(A8), are all consequences of the corresponding statements, (T1)–(T8), listed in [Bo3]. (We set $K - 1 = h_0$, $K_2 = h_1(\nu_0)$, $\xi_1 = \xi_1(\nu_1)$, $\xi_1' = \xi_1(\nu_2)$ in the notation of that paper. Note that (A5) and (A8) are slightly weaker than the corresponding statements of [Bo3].) Indeed, for these we could take any $\nu_2 \geq \nu_0$. (In that paper, τ was a track filling Σ, and we were taking $\partial\Sigma \subseteq \tau$. To get a multitrack in the sense defined above, we can simply replace τ by the subtrack, $p(\gamma)$, pushed into the interior of Σ, if necessary, so that $\tau \cap \partial\Sigma = \emptyset$. Property (A9) is specific to the set up in this paper, and calls for further comment.

We can suppose that $\mathcal{S}_0 \neq \emptyset$, otherwise nothing remains to be proven. We set $\mathcal{S}_1 = \mathcal{S}_0$ and $A_1(\delta) = A_0(\delta)$. Let $s_0 \geq \xi_0 K_2$. We suppose that $\nu_1 \geq \nu_0$ is large enough so that $N(\Delta_1, s_0) \subseteq \Delta_0$. Let $\breve{\delta} \in \tilde{\mathcal{S}}_1$. We claim that $f^{-1}(\Delta_1(\breve{\delta})) \subseteq A_0(\breve{\delta})$. Suppose that $x \in \tilde{\tau}$, with $f(x) \in \Delta_1(\breve{\delta})$, and that $x \notin A_0(\breve{\delta})$. Then we can connect x to some point $y \in A_0(\breve{\epsilon})$ by a path in $\tilde{\tau}$ of d_τ-length at most K_2, where $\breve{\epsilon} \in \tilde{\mathcal{S}}_0 \setminus \{\breve{\delta}\}$. Thus $d(\tilde{f}(x), \tilde{f}(y)) \leq \xi_0 K_2 \leq s_0$.

But since $\tilde{f}(A_0(\epsilon)) \subseteq \tilde{\Delta}_0(\breve{\epsilon})$, it now follows that $\breve{\epsilon} = \breve{\delta}$ giving a contradiction. This now justifies the statement (A9).

Suppose that γ has no short components, and that $\breve{\gamma}$ is a component of $\tilde{\gamma} \subseteq \tilde{\Sigma}$. Now $\tilde{p} : \breve{\gamma} \longrightarrow \tilde{\tau}$ is injective, and $\tilde{p}(\breve{\gamma})$ is a bi-infinite trainpath. Also $\tilde{f}\tilde{p}$ and τ are a bounded distance apart. Now $\tilde{f}$ and $\tilde{p}$ are uniformly lipschitz, and $\tilde{q}$ is uniformly quasigeodesic, so it follows that both $\tilde{p} : \breve{\gamma} \longrightarrow \tilde{\tau}$ and $\tilde{f}|\tilde{p}(\breve{\gamma}) \longrightarrow \tilde{M}$ are both uniformly quasigeodesic.

We also want to relate this to combinatorial distance in τ, where each branch of τ has unit length. Suppose that $\breve{\zeta} \subseteq \breve{\gamma}$ is an arc such that $fp(\zeta) \cap \Delta_2 = \emptyset$. In particular, this means that $p(\zeta) \subseteq \tau \setminus A_2(\delta)$. Since each component of $\tau \setminus A_2(\Sigma)$ has bounded diameter, there is a linear bound on the length of $\tilde{p}(\breve{\zeta})$ in terms of the number of branches it passes through. Since $\tilde{p} : \breve{\zeta} \longrightarrow \tilde{\tau}$ is quasigeodesic, this also places a linear bound on the length of $\breve{\zeta}$.

Conversely, if ζ is an arc in γ of bounded length, with $fp(\zeta) \cap \Delta_2 = \emptyset$, then there is a bound on the number of branches of τ that $p\zeta$ can pass through, since each time $p\zeta$ passes through a given branch, we get an element of bounded displacement in $\tilde{M}$, and the number of these is bounded by the assumption that $fp(\zeta) \cap \Delta_2 = \emptyset$. Thus, if ζ is any path with $fp(\zeta) \cap \Delta_2 = \emptyset$, then the combinatorial length of $p\zeta$ is linearly bounded above by the d_γ-length of ζ.

3. The uniform injectivity theorem.

In this section, we prove a result will serve as a substitute for the uniform injectivity theorem for pleated surfaces in 3-manifolds. The argument we give here is based on a construction of Namazi and Souto [NS], though some reintepretation is necessary to deal with coarse hyperbolic spaces. I am grateful to Hossein Namazi for explaining the main idea of their proof to me.

Suppose γ is a multicurve with no short components (i.e. all components have length at least the hyperbolicity constant in M). Let $q : \gamma \longrightarrow M$ be a realisation, and let $\tilde{q} : \tilde{\gamma} \longrightarrow \tilde{M}$ be its lift to the preimage, $\tilde{\gamma}$, to $\tilde{\Sigma}$. We have observed that each component is uniformly quasigeodesic.

Suppose that $\zeta \subseteq \gamma$ is any arc, and $\breve{\zeta} \subseteq \tilde{\gamma}$ is a lift. We can always find a homeomorphism, $\lambda : I \longrightarrow \zeta$, where $I \subseteq \mathbf{R}$ is some interval, such that for all $t, u \in I$ we have $|d(\tilde{q}\tilde{\lambda}(t), \tilde{q}\tilde{\lambda}(u)) - |t - u||$ bounded above in terms of the hyperbolicity constant of $\tilde{M}$. We refer to such a homeomorphism as an *ambient parameterisation* of ζ. (The appropriate bound chosen will be apparent in what follows.)

Suppose that $\breve{\zeta}, \breve{\zeta}' \subseteq \tilde{\gamma}$ are two oriented arcs with positive and negative endpoints, x, y and x', y' respectively. We write $d_{\tilde{M}}(\tilde{q}\breve{\zeta}, \tilde{q}\breve{\zeta}') = \max\{d_{\tilde{M}}(\tilde{q}x, \tilde{q}x'), d_{\tilde{M}}(\tilde{q}y, \tilde{q}y')\}$.

We note the following:

Lemma 3.1 : *There is some h_0, depending only on the hyperbolicity constant with the following property. Suppose that $\breve{\zeta}, \breve{\zeta}'$ are as above, and let $h = d_{\tilde{M}}(\tilde{q}\breve{\zeta}, \tilde{q}\breve{\zeta}')$. Then there are intervals, $I, J, J' \subseteq \mathbf{R}$, with $I \subseteq J \cap J'$, with each component of $J \setminus I$ and of $J' \setminus I$ having length at most h, and positively oriented ambient parameterisations $\lambda : J \longrightarrow \zeta$*

and $\lambda' : J' \longrightarrow \zeta'$ *such that for all* $t \in I$*, we have* $d_{\tilde{M}}(\tilde{q}\tilde{\lambda}(t), \tilde{q}\tilde{\lambda}'(t)) \leq h_0$. $\diamondsuit$

This is just an expression of the "fellow travelling" property of uniform quasigeodesics in a hyperbolic space: if two quasigeodesics are a bounded distance apart at their endpoints, then they remain a uniformly bounded distance apart over most of their length.

Given maps, $v, v' : I \longrightarrow \tilde{M}$, we write $v \sim_h v'$ to mean that $d_{\tilde{M}}(v(t), v'(t)) \leq h$ for all $t \in I$. We will abbreviate $\sim_{h_0}$ to $\sim$, where h_0 is the constant of Lemma 3.1.

Definition : Suppose $\breve{\zeta}, \breve{\zeta}' \subseteq \tilde{\gamma}$ are oriented arcs. We say that $\tilde{q}\breve{\zeta}$ and $\tilde{q}\breve{\zeta}'$ are *parallel* if they admit positively oriented ambient parameterisations $\lambda : I \longrightarrow \zeta$ and $\lambda' : I' \longrightarrow \zeta'$ such that $\tilde{q}\tilde{\lambda} \sim \tilde{q}\tilde{\lambda}'$.

Thus, modulo reducing the lengths of I and I' by at most h_0, this is equivalent to saying that $d_{\tilde{M}}(\tilde{q}\breve{\zeta}, \tilde{q}'\breve{\zeta}') \leq h_0$.

Recall that in Section 2 we defined three nested thin parts of M, namely $\Delta_2 \subseteq \Delta_1 \subseteq \Delta_0$, with $N(\Delta_1, s_0) \subseteq \Delta_0$ and $N(\Delta_2, s) \subseteq \Delta_1$. We will postpone the explicit specification of s_0 and s. We can certainly assume that s is sufficiently large so that there is some ν with $\nu_1 \leq \nu \leq \nu_2$, giving an intermediate thin part, $\Delta_2 \subseteq \Delta_\nu \subseteq \Delta_1$, and with $N(\Delta_2, s_2) \subseteq \Delta_\nu$ and $N(\Delta_\nu, s_1) \subseteq \Delta_1$. We will specify required bounds on s_1 and s_2, depending on the parameters of M, later.

Now let τ be an associated multitrack as discussed in Section 2. With s_1 is sufficiently large, we have the following:

Lemma 3.2 : $q(\gamma) \cap \Delta_\nu \subseteq \bigcup_{\delta \in \mathcal{S}_1} \Delta_\nu(\delta)$.

Proof : We know that Δ_ν is a disjoint union of $\Delta_\nu(\epsilon)$ as ϵ varies over $Y_0 = Y_0(\Sigma)$. Let $x \in \gamma$ with $q(x) \in \Delta$. Thus $q(x) \in \Delta_\nu(\epsilon)$ for some $\epsilon \in Y_0$. Now $d(q(x), fp(x)) \leq K_3$ and so $fp(x) \in N(\Delta_\nu(\epsilon), K_3) \subseteq \Delta_1(\epsilon)$ (provided $s_1 \geq K_3$). By the property of Δ_1 discussed in Section 2, we have $p(x) \in f^{-1}(\Delta_1) \subseteq A_1(\Sigma)$. That is, $p(x) \in A_1(\delta)$ for some $\delta \in \mathcal{S}_1$, and so $fp(x) \in \Delta_0(\delta)$ (again from Section 2). But we already have $fp(x) \in \Delta_1(\epsilon) \subseteq \Delta_0(\epsilon)$ and so $\Delta_0(\delta) \cap \Delta_0(\epsilon) \neq \emptyset$. Thus $\epsilon = \delta$. $\diamondsuit$

The idea behind the uniform injectivity theorem is that if two intervals in γ follow each other in M then they also follow each other in Σ. We shall express this in terms of the associated multitrack, τ, described in Section 2. In certain cases, it can be reformulated in terms of a hyperbolic structure on Σ as we describe in Section 5.

Definition : We say that two oriented trainpaths, $\breve{\zeta}$, $\breve{\zeta}'$ in $\tilde{\tau}$ *intersect positively* if $\breve{\zeta} \cap \breve{\zeta}'$ is a non-trivial interval, and $\breve{\zeta}$ and $\breve{\zeta}'$ are oriented consistently on this interval.

(Note that the intersection of any two trainpaths in $\tilde{\tau}$ is connected.)

If $\tilde{\pi}, \tilde{\pi}' : I \longrightarrow \tilde{\pi}$ are parameterisations of $\breve{\zeta}$ and $\breve{\zeta}'$, this can be expressed by saying that there exist $t < u \in I$ and $t' < u' \in I$ with $\tilde{\pi}(t) = \tilde{\pi}'(t')$ and $\tilde{\pi}(u) = \tilde{\pi}'(u')$.

In the following formulation of the uniform injectivity theorem, $\Delta_\nu \subseteq M$ is the thin part introduced above. For the proof, we can assume without loss of generality, that ν is sufficiently large that it has the properties given above.

Theorem 3.3 : *Given $\nu > 0$, there is some constant, $T_0 > 0$, depending only on the parameters of M with the following property. Suppose that γ is a multicurve with no short components. Let $q : \gamma \longrightarrow M$ be a realisation, and let τ be an associated multitrack with carrying map, $p : \gamma \longrightarrow \tau$. Suppose that ζ, ζ' are arcs in γ of d_γ-length at least T_0 with $(q\zeta \cup q\zeta') \cap \Delta_\nu = \emptyset$. Let $\breve{\zeta}, \breve{\zeta}'$ be lifts to $\tilde{\gamma}$. If $\tilde{q}\breve{\zeta}$ and $\tilde{q}\breve{\zeta}'$ are parallel, then $\tilde{p}\breve{\zeta}$ and $\tilde{p}\breve{\zeta}'$ intersect positively.*

Proof : From the earlier discussion, we can assume that $\zeta = \lambda(I)$, $\zeta' = \lambda'(I)$, where $I = [-v_0, v_0]$, and $\lambda, \lambda' : I \longrightarrow \breve{\gamma}$, are ambient parameterisations, and that $\tilde{q}\tilde{\lambda} \sim \tilde{q}\tilde{\lambda}'$.

Recall that Δ_2 is chosen so that $N(\Delta_2, s_2) \subseteq \Delta_\nu$ and that $f(\tau \cap A_2(\Sigma)) \subseteq \Delta_2$ (by the construction in Section 2). For all $t \in I$, $d(q\lambda(t), fp\lambda(t)) \leq K_2 \leq s_2$ provided we assume that $s_2 \geq K_2$. It follows that $fp\lambda(I) \cap \Delta_2 = \emptyset$, and so $p\lambda(I) \cap A_2(\Sigma) = \emptyset$. Similarly, $p\lambda'(I) \cap A_2(\Sigma) = \emptyset$.

Now $\tau \setminus A_2(\Sigma)$ has bounded length, an wo we can cover it by a set, $\mathcal{I}$, of intervals of bounded length, each lying in some branch of τ, and with $|\mathcal{I}|$ bounded by some number, N, depending only on ν andthe parameters of M. From the above discussion, we have $p\lambda(I) \cap p\lambda'(I) \subseteq \bigcup \mathcal{I}$. We write $\tilde{\mathcal{I}}$ for the set of lifts of these intervals to $\tilde{\tau}$.

Now write $\tilde{\pi} = \tilde{p}\tilde{\lambda} : I \longrightarrow \tilde{\tau}$ and $\tilde{\pi}' = \tilde{p}\tilde{\lambda}' : I \longrightarrow \tilde{\tau}$, and write π and π' for their projections to M. We write π_+ and π_- for the trainpaths $\pi|[0, v_0]$ and $\pi|[-v_0, 0]$ respectively. We fix some $u_0 \geq 0$, to be determined later. We consider the pairs $(\pi(t), \pi'(t))$ as t varies from 0 to v_0. Each such pair lies in $\theta \times \theta'$ for some $\theta, \theta' \in \mathcal{I}$. There are at most N^2 possibilities. Thus, if v_0 is sufficiently large, we find some fixed $\theta, \theta' \in \mathcal{I}$ and t_i with $0 < t_0 < t_1 < \cdots < t_m < v_0$ with $t_i + u_0 < t_{i+1}$ for all i, such that $\pi(t_i)$ and $\pi'(t_i)$ lie in the interiors of θ and θ' respectively, such that π and π' pass through these intervals in the same direction at these points, and where $m \in \mathbf{N}$ is some fixed constant depending only on the parameters, as determined below. Now if u_0 is sufficiently large, we see that for each i, π must leave the interval θ, and π' must leave the inerval θ', between t_i and t_{i+1}.

Now if $i < j$, we can represent $\pi|[0, t_j]$ as a union of a path $\alpha_i = \pi|[0, t_i]$ and a path $\beta_{ij} = \pi|[t_i, t_j]$ whose image is a closed trainpath in τ (at least if we allow ourselves to slide t_i along the interval θ). This represents an element $g_{ij} \in \pi_1(\Sigma, \pi(0))$. Writing $g_i = g_{0i}$, we see that $g_{ij} = g_j g_i^{-1}$. Similarly, we can write $\pi'|[0, t_j] = \alpha'_i \cup \beta'_{ij}$. This is represented by $g'_j (g'_i)^{-1} \in \pi_1(\Sigma, \pi'(0))$. We can interpret this in terms of $\tilde{\tau} \subseteq \tilde{\Sigma}$ as follows.

Let $\tilde{\pi}(t_i) \in \tilde{\theta}_i \in \tilde{\mathcal{I}}$ and $\tilde{\pi}'(t_i) \in \tilde{\theta}'_i \in \tilde{\mathcal{I}}$. We connect $\tilde{\pi}(0)$ to $\tilde{\pi}'(0)$ by a path $\tilde{\epsilon}$. This projects to a path ϵ in Σ from $\pi(0)$ to $\pi'(0)$, thereby allowing us to identify $\pi_1(\Sigma, \pi(0)) \equiv \pi_1(\Sigma, \pi'(0)) \equiv \Gamma$. Moreover, we see that $\tilde{\theta}_i = g_i \tilde{\theta}_0$ and $\tilde{\theta}'_i = g'_i \tilde{\theta}'_0$.

Let $\tilde{x}_i = \tilde{q}\tilde{\lambda}(t_i) \in \tilde{M}$ and $\tilde{y}_i = \tilde{f}\tilde{p}\tilde{\lambda}(t_i) \in \tilde{M}$. Thus, $d_{\tilde{M}}(\tilde{x}_i, \tilde{y}_i) \leq K_3$, by the property of τ described in Section 2. We similarly define $\tilde{x}'_i$ and $\tilde{y}'_i$, and again have $d_{\tilde{M}}(\tilde{x}'_i, \tilde{y}'_i) \leq K_3$. Moreover, since $\tilde{q}\tilde{\lambda} \sim \tilde{q}\tilde{\lambda}'$, we have $d(\tilde{x}_i, \tilde{x}'_i) \leq h_0$ for all i.

Now $\tilde{\pi}(t_i), g_i\tilde{\pi}(t_0) \in \tilde{\theta}_i$, and so $d_{\tilde{\tau}}(\tilde{\pi}, g_i\tilde{\pi}(t_0)$ is bounded since $\tilde{f}$ is lipschitz, and therefore is $d_{\tilde{M}}(\tilde{y}_i, g_i\tilde{y}_0) = d_{\tilde{M}}(\tilde{f}\tilde{\pi}(t_i), g_i\tilde{f}\tilde{\pi}(t_0))$ is bounded. It follows that $d_{\tilde{M}}(\tilde{x}_i, g_i\tilde{x}_0)$ is bounded. Similarly, $d_{\tilde{M}}(\tilde{x}'_i, g'_i\tilde{x}'_0)$ is bounded. This therefore also bounds $d_{\tilde{M}}(\tilde{x}'_0, g_i^{-1}g'_i\tilde{x}'_0) = d_{\tilde{M}}(g_i\tilde{x}'_0, g'_i\tilde{x}'_0) \leq d_{\tilde{M}}(\tilde{x}_0, \tilde{x}'_0) + d_{\tilde{M}}(g_i\tilde{x}_0, g'_i\tilde{x}'_0)$. That is, for all i, $g_i^{-1}g'_i \in \Gamma_r(\tilde{x}_0)$, where r

depends only on the parameters of the action. But $\tilde{x}_0 \in \tilde{q}\tilde{\lambda}(I) = \tilde{q}\breve{\zeta}$ and so, by assumption, $x_0 \notin \Delta_\nu$. Thus, by (M2), there is some bound, m, on the cardinality of $\Gamma_r(\tilde{x}_0)$, again depending only on the parameters of the action. This retrospectively gives us our constant, m, introduced above. It now follows that there exist $i < j$ so that $g_i^{-1}g_i' = g_j^{-1}g_j'$ and so $g_j g_i^{-1} = g_j'(g_i')^{-1}$.

Now let $\alpha_+ = \alpha_i$, $\beta_+ = \beta_{ij}$, $\alpha_+ = \alpha_i'$ and $\beta_+' = \beta_{ij}'$ be as above. Note that $\sigma_+ = \alpha_+ \cup \beta_+$ and $\sigma_+' = \alpha_+' \cup \beta_+'$ represent the same element of Γ. In particular, β_+ and β_+' are freely homotopic in Σ, and so represent the same closed trainpath in τ. Moreover, since $\tilde{\pi}, \tilde{\pi}' : I \longrightarrow \tilde{\tau}$ are uniform quasigeodesics, we can arrange that this closed trainpath has τ-length at least a given fixed constant (to be chosen below) by choosing the constant u_0 sufficiently large.

We now carry out the same argument with π_- and π_-'. We get paths $\alpha_- \cup \beta_-$ and $\alpha_-' \cup \beta_-'$ representing the same element of Γ, and so β_- and β_-' are the same closed trainpath in τ. We write $\sigma = \sigma_- \cup \sigma_+$ and $\sigma' = \sigma_-' \cup \sigma_+'$. Let $\breve{\sigma}$ and $\breve{\sigma}'$ be lifts of these arcs to $\tilde{\tau}$.

Back in $\tilde{\tau}$, we have two bi-infinite paths $\breve{\beta}_+ = \breve{\beta}_+'$ and $\breve{\beta}_- = \breve{\beta}_-'$ such that the lifted trainpaths $\breve{\sigma}$ and $\breve{\sigma}'$ both run from $\breve{\beta}_-$ in the same sense and into $\breve{\beta}_+$ in the same sense.

We distinguish two possibilities as follows.

Case (1): $\breve{\beta}_- \cap \breve{\beta}_+ = \emptyset$.

In this case $\breve{\sigma}$ and $\breve{\sigma}'$ both cross from $\breve{\beta}_-$ to $\breve{\beta}_+$ along the same trainpath and so, in particular, intersect positively. (We use the fact that a train track has no complementary digons.)

Case (2): $\breve{\beta}_- \cap \breve{\beta}_+ \neq \emptyset$.

For this we will use the following general principle. Suppose that μ, μ' are overlapping trainpaths in $\tilde{\tau}$, with $\mu \cup \mu'$ a trainpath. Suppose that $\tilde{f}\mu$ and $\tilde{f}\mu'$ are quasigeodesic in $\tilde{M}$. Then if the $d_{\tilde{\tau}}$-length of $\mu \cap \mu'$ is sufficiently large in relation to the quasigeodesic constants, then $\tilde{f}(\mu \cup \mu')$ is also quasigeodesic, with constants depending only on the original. This follows from the fact that being quasigeodesic in a Gromov hyperbolic space is a local property. In fact, the above extends to any sequence of intervals, each overlapping the next in a sufficiently long arc.

Now we know that $\tilde{f}\breve{\sigma}$ and $\tilde{f}\breve{\sigma}'$ are uniformly quasigeodesic in $\tilde{M}$ (since they follow, respectively, within a bounded distance of $\tilde{q}\breve{\zeta}$ and $\tilde{q}\breve{\zeta}'$). We can therefore choose some l sufficiently large for the above construction to work, with these quasigeodesic constants, and choose u_0 in the above construction, large enough so that β and β' both have length at least l.

Now let ω be the trainpath $\breve{\beta}_- \cap \breve{\beta}_+$. (This will be all of $\breve{\beta}_-$ and $\breve{\beta}_+$ if these happen to be equal.) Now since σ and σ' both wrap around β_- and β_+ in the same direction, it follows that if $\breve{\sigma}$ and $\breve{\sigma}'$ intersect, then they must intersect positively. Thus, we assume, for contradiction that $\breve{\sigma} \cap \breve{\sigma}' = \emptyset$. Since both cross from $\breve{\beta}_-$ to $\breve{\beta}_+$, we see that, up to interchanging σ and σ', we have $\breve{\sigma} \subseteq \breve{\beta}_-$ and $\breve{\sigma}' \subseteq \breve{\beta}_+$, and that $\breve{\sigma} \cup \omega \cup \breve{\sigma}'$ is a trainpath. We can also assume that β_- is no longer than $\breve{\beta}_-$ (by interchanging the positive and negative directions).

Now σ must wrap at least twice around $\breve{\beta}_-$, and so from the above principle, we see that if length$(\beta_-) \geq l$ is sufficiently large, the $\tilde{f}\breve{\beta}_-$ will be uniformly quasigeodesic (since $\breve{\beta}_-$ is made up of a sequence of overlapping images of $\breve{\sigma}$). In particular, $\tilde{f}\omega$ is uniformly quasigeodesic.

But now σ also wraps at least once around β_-, and so it follows that length$(\sigma \cap \omega) \geq$ length$(\beta_+) \geq l$. Similarly, length$(\breve{\sigma}' \cap \omega) \geq$ length$(\beta_-) \geq l$. So again, since $\tilde{f}\breve{\sigma}$, $\tilde{f}\breve{\sigma}'$ and $\tilde{f}\omega$ are all quasigeodesic, it follows that $\tilde{f}(\breve{\sigma} \cup \omega \cup \breve{\sigma}')$ is uniformly quasigeodesic. But now, from the original construction, $\tilde{f}\breve{\sigma}$ and $\tilde{f}\breve{\sigma}'$ remain a bounded distance apart in $\tilde{M}$. Thus, again if l is sufficiently large, we get a contradiction.

We have shown that, in both cases, $\breve{\sigma}$ and $\breve{\sigma}'$, and hence $\tilde{p}\breve{\zeta}$ and $\tilde{p}\breve{\zeta}'$ intersect positively. This involved choosing v_0 sufficiently large in relation to the parameters of M. This is allowed, since T_0 is linearly bounded above in terms of v_0, so it is enough to choose T_0 sufficiently large. $\diamondsuit$

An immediate consequence of Theorem 3.3 is the following.

Suppose that $I \subseteq \mathbf{R}$ is an interval, and $t \in I$ with $[t-2v_0, t+2v_0] \subseteq I$. Then, both $\tilde{p}\tilde{\lambda}|[t-2v_0, 0]$ and $\tilde{p}\tilde{\lambda}|[0, t+2v_0]$ must intersect $\tilde{p}\tilde{\lambda}'|[t-2v_0, t+2v_0]$. Since trainpaths in $\tilde{\tau}$ have connected intersection, it follows that $\tilde{p}\tilde{\lambda}(t) \in \tilde{p}\tilde{\lambda}'|[t-2v_0, t+2v_0]$. Moreover, $\tilde{p}\tilde{\lambda}(I)$ and $\tilde{p}\tilde{\lambda}'(I)$ intersect positively there. Less formally this is saying that, up to changing the parametrisations by a bounded amount, $\tilde{p}\tilde{\lambda}$ and $\tilde{p}\tilde{\lambda}'$ agree in all but a bounded neighbourhood of the boundary of I.

Reinterpreting in terms arcs in γ, we have shown:

Corollary 3.4 : *There is some T_0 such that, with same hypotheses as Theorem 3.3, if $x \in \zeta$ is at least a distance T_0 from the endpoints of ζ, then there exist $y, z \in \zeta'$ with $d_\gamma(y,z) \leq T_0$, $d(\tilde{q}\breve{x}, \tilde{q}\breve{y}) \leq h_0$, $\tilde{p}(\breve{x}) = \tilde{p}(\breve{z})$, and with $\breve{\zeta}$ and $\breve{\zeta}'$ intersecting positively at $\tilde{p}(\breve{x}) = \tilde{p}(\breve{z})$.* $\diamondsuit$

4. Quasiprojections.

We want to associate to a curve, $\alpha \in X(\Sigma)$, another "derived curve", $\alpha' \in X(\Sigma)$, with $l_M(\alpha')$ bounded. This will lead to a quasiprojection on the curve graph $\mathcal{G}(\Sigma)$, giving us a proof of Theorem 0.2. The derived curve is obtained surgery on the original. We begin with the relevant topological construction.

By a *bridge arc*, we mean an arc, ϵ, embedded in Σ with endpoints, $x, x' \in \alpha$. We regard it as defined up to homotopy relative to its endpoints. We realise it so that $|\alpha \cap \epsilon|$ is minimal. We can think of $\alpha \cup \epsilon$ as the immersion of a theta-curve in Σ. If θ and θ' are the components of $\alpha \setminus \{x, x'\}$, we refer to the closed curves $\epsilon \cup \theta$ and $\epsilon \cup \theta'$ as the *halves* of the theta-curve. Now, if we put an orientation on α, then there is a parity to each intersection of ϵ with α. If the parities at the endpoints, x, x', are the same, then we say that ϵ is *consistently oriented*.

Definition : We say that ϵ, θ is *simple* if it is consistently oriented and $\epsilon \cap \theta = \{x, x'\}$.

In this case, $\epsilon \cup \theta$ is an essential non-peripheral simple closed curve — the *derivative* of ϵ.

Definition : We say that ϵ, θ is *semisimple* if it is consistently oriented, and $\epsilon \cap \theta = \{x, x', y\}$ with $y \notin \{x, x'\}$, and where the intersection at y has opposite partity to that at x (and therefore at x').

In this case, we need a more complicated surgery. Let $\epsilon = \epsilon_1 \cup \epsilon_2$ and $\theta = \theta_1 \cup \theta_2$, where ϵ_1, θ_1 are the segments between x and y and ϵ_2, θ_2 are the segments between y and x'. We now set the *derivative* to be $\theta_1 \cup \epsilon_2 \cup (-\theta_2) \cup (-\epsilon_1)$. In other words, we are performing a surgery at y to eliminate the intersection there. Note that the derivative is again simple, essential and non-peripheral.

We note:

Lemma 4.1 : *Suppose that ϵ is a consistently oriented bridge arc of α, and that θ is a component of $\alpha \setminus \partial\epsilon$. There is a subarc, $\epsilon' \subseteq \epsilon$, that is also a bridge arc of α, and a component, θ', of $\alpha \setminus \partial\epsilon'$, with $\theta' \subseteq \theta$, and such that ϵ, θ is either simple or semisimple.*

Proof : This follows easily considering the intersections of θ along ϵ. $\diamondsuit$

Suppose now that $q : \alpha \longrightarrow M$ is a realisation of α, and that α is not short, so that each component of $\tilde{q} : \tilde{\alpha} \longrightarrow \tilde{M}$ is uniformly quasigeodesic. Suppose that ζ, ζ' are arcs in α, and that ϵ is a bridge arc connecting $x \in \zeta$ to $x' \in \zeta'$. We refer to ζ and ζ' as *terminal arcs*. We say that ζ and ζ' are *U-long* if each component of $\zeta \setminus \{x\}$ and of $\zeta' \setminus \{x'\}$ has d_γ-length at least U. If $\breve{\zeta}$ is any lift of ζ to $\tilde{\alpha}$, then we can also lift ϵ and ζ' to $\breve{\epsilon}$ and $\breve{\zeta}'$ respectively, so that $\breve{\epsilon}$ connects $\breve{x} \in \breve{\zeta}$ to $\breve{x}' \in \breve{\zeta}'$. If we orient ζ and ζ' so that they cross ϵ in the same sense (regardless of their orientations in α), then we say that $q\zeta$ and $q\zeta'$ are parallel if $\tilde{q}\breve{\zeta}, \tilde{q}\breve{\zeta}'$ are parallel in $\tilde{M}$. Note that if we were to allow ourselves to slide x and x' along ζ and ζ', we could always arrange that $d_{\tilde{M}}(\breve{x}, \breve{x}') \leq h_0$. In this case, we say that $(q\zeta, qx), (q\zeta', qx')$ are *parallel*. We can then extend $q : \alpha \longrightarrow M$ to $q : \alpha \cup \epsilon \longrightarrow M$ so that ϵ is a path of length at most h_0 in the appropriate homotopy class.

In general, we say that a bridge, ϵ, with endpoints x, x' is *U-wide* of we can find U-long terminal arcs $\zeta, \zeta' \subseteq \alpha$ (not necessarily disjoint), with $x \in \zeta$, $x' \in \zeta'$ so that $(q\zeta, qx), (q\zeta', qx')$ are parallel.

We can extend this terminology to an intersection, y, of α with ϵ. In this case, saying that ϵ is U-wide at y is interpreted to mean that both ϵ_1 and ϵ_2 are U-wide, where y cuts ϵ into ϵ_1 and ϵ_2. (This means that y lies in a third arc, $\zeta' \subseteq \alpha$, again oriented, so that ζ'' crosses ϵ at y in the same direction as ζ and ζ', and satisfying the same conditions).

In practice, we will be dealing with the following two situations. Suppose that ϵ is a bridge, and that θ is a component of $\alpha \setminus \partial\epsilon$. Suppose ϵ, θ is simple and U-wide. Then $\alpha' = \epsilon \cup \theta$ lies in $X(\Sigma)$, and $q : \alpha \cup \epsilon \longrightarrow M$ restricts to a map $q : \alpha' \longrightarrow M$. Thus $l_M(\alpha') \leq \text{length}(q(\alpha')) \leq \text{length}(\theta) + h_0$. (Recall that the metric d_α is defined as the path-metric induced from M.)

Suppose that ϵ, θ is semisimple, with θ crossing ϵ at y, and that y cuts ϵ into ϵ_1 and ϵ_2. Suppose that ϵ is U-wide at y. We can extend q to a map $q : \alpha \cup \epsilon_1 \cup \epsilon_2 \longrightarrow M$ so that $q(\epsilon_1)$ and $q(\epsilon_2)$ both have length at most h_0. This also gives us a map of the derived curve α' into M, and we see in this case that $l_M(\alpha') \leq \mathrm{length}(\theta) + 2h_0$.

Definition : We say that ϵ is a *U-surgerable* bridge if it is one of the above two types.

To describe our procedure, we need to indroduce some constants depending only on the parameters of the action, to be specified later. We will fix a constant, $\nu_3 \leq \nu_1$, and set $\Delta_3 = \Delta_{\nu_3}$. We will also choose two constants, U_0 and U_1, and set $U_2 = 2U_0 + 2U_1$.

Suppose now that $\alpha \in X(\Sigma)$, and that $q : \alpha \longrightarrow M$ is a realisation. We proceed as follows:

(1) If $l_M(\alpha) \leq U_2$, set $\alpha' = \alpha$.

(2) If $l_M(\alpha) > U_2$ and there is some $\delta \in X(\Sigma)$ with $q(\alpha) \cap \Delta_3(\delta) \neq \emptyset$, set $\alpha' = \delta$.

(3) Suppose that $l_M(\alpha) > U_2$ and $\alpha \cap \Delta_3 = \emptyset$. Suppose that ϵ is a bridge arc, and θ is a component of $\alpha \setminus \partial\epsilon$ with $\mathrm{length}(\theta) \leq U_1$. Suppose that ϵ, θ is U_0-surgerable. Then we set α' to be the derived curve of ϵ, θ, as described above (depending on whether ϵ, θ is simple or semisimple).

The above cases are clearly mutually exclusive, though it is less clear that they are inclusive, and we postpone the issue of existence until later (see Lemma 4.4).

We begin by showing that a choice of derived curve in the above proceedure is well defined up to bounded intersection in Σ. Indeed, any respective derived curves of two disjoint curves have bounded intersection. In what follows we can suppose that we have fixed a realisation in M for each $\alpha \in X(\Sigma)$ (though this choice makes no essential difference).

We begin by setting U_0 to be the constant T_0 of Theorem 3.3, given $\nu = \nu_3$ (to be specified later).

Now if $\alpha \cap \Delta_3 = \emptyset$, we can apply Theorem 3.3 (the uniform injectivity theorem). Let $p : \alpha \longrightarrow \tau$ be a carrying map to an associated track, τ. If ζ, ζ', ϵ are as above, then the trainpaths $\tilde{p}\breve{\sigma}$ and $\tilde{p}\breve{\zeta}'$ intersect positively. Back down in τ, this means that the trainpaths $p\zeta$ and $p\zeta'$ intersect positively, and moreover, we can extend p to a map of $\alpha \cup \epsilon$ homotopic to the original map of $\alpha \cup \epsilon$ into Σ, such that $p\epsilon$ is a point of this intersection.

We can now apply this in case (3) of the above construction. If ϵ, θ is simple, we slide ϵ as above, and see that the derived curve, α' is represented (up to free homotopy in Σ) as a closed trainpath in τ. Similarly, if ϵ, θ is semisimple, then we can slide ϵ_1 and ϵ_2 separately, and again see that α' is represented by a closed trainpath. We denote such a closed trainpath by $p\alpha'$. Since p is lipschitz, we see that such a closed trainpath has bounded d_τ-length. It must therefore also have bounded combinatorial length, following the discussion at the end of Section 3. That is, it passes through a bounded number of branches of τ.

We are now ready for:

Proposition 4.2 : *Suppose that $\alpha, \beta \in X(\Sigma)$ are equal or adjacent in $\mathcal{G}(\Sigma)$, and that α', β' are respective derived curves (obtained by the above proceedure). Then the intersection number, $\iota(\alpha', \beta')$ is bounded above in terms of the parameters of the action.*

Proof : Let γ be the multicurve $\alpha \cup \beta$. The realisations $q : \alpha \longrightarrow M$ and $q : \beta \longrightarrow M$ together give us a realisation $q : \gamma \longrightarrow M$. Let $p : \gamma \longrightarrow \tau$ be a carrying map to an associated multitrack, and let $\mathcal{S}$ be the set of special loops.

Consider first α. Note that by Lemma 3.2, if $\delta \in X(\Sigma)$ with $q(\alpha) \cap \Delta_3(\delta) \neq \emptyset$, then $\delta \in \mathcal{S}$, and $p\alpha' = p\delta$ has combinatorial length 1. Thus in all three cases, $p\alpha'$ has bounded combinatorial length. The same goes for $p\beta'$. It now follows that $\iota(\alpha', \beta')$ is bounded. $\diamondsuit$

In fact, we have also effectively shown:

Lemma 4.3 : *If α' is any derived curve of α, then $\iota(\alpha, \alpha')$ is bounded above in terms of $l_M(\alpha)$ and the parameters of the action.*

Proof : This follows from the above argument, since the combinatorial length of $p\alpha$ in τ is bounded above in terms $l_M(\alpha)$. $\diamondsuit$

Note that all the above works for an arbitrarily chosen constant, $U_1 \geq 0$. We now move on to consider the existence of derived curves. For that, we will need to choose U_1 sufficiently large in relation to the parameters of the action.

First we make some general observations regarding bridge arcs.

Suppose that α is a curve and ϵ a bridge arc. Suppose that $p : \alpha \longrightarrow \tau$ is a carrying map that extends to a map $p : \alpha \cup \epsilon \longrightarrow \tau$, homotopic to the standard map of $\alpha \cup \epsilon$ into Σ, and sending ϵ to a point. Suppose that the endpoints x, x' of ϵ lie in intervals μ, μ' in α (not necessarily disjoint). We say that μ and μ' *follow* each other in τ if $p\mu$ and $p\mu'$ are equal as trainpaths in τ. This means that the lifts, $\breve{\mu}, \breve{\mu}'$, to $\tilde{\tau}$ get mapped to identical paths, $\tilde{p}\breve{\mu} = \tilde{p}\breve{\mu}'$, in $\tilde{\tau}$. Now suppose that α crosses ϵ at a third point, y. Then there must be an arc, $\mu'' \subseteq \alpha$ containing y, so that μ'' follows both μ and μ' in τ. (This is with respect to the bridge arcs ϵ_1 and ϵ_2 inside ϵ.) This is due the fact that μ'' gets trapped between μ and μ'. (More formally, we can think of α as a leaf in a foliated neighbourhood of τ in Σ (cf. [PH]).

In summary, using the earlier discussion of simple and semisimple bridges, this entails the following. Suppose we can find a bridge arc, ϵ, whose endpoints lie at the centre point of an arc in α, of τ-length at least some constant, V_0, say and that these follow each other in τ and in the same direction (taking induced orientations from an orientation of α). Let θ be a component of $\alpha \setminus \partial\epsilon$. Then we can assume that ϵ, θ is either simple or semisimple. In the second case, we can find another arc, ζ'', centred on the third intersection of ϵ with θ, oriented in the opposite direction along α, which also follows ζ and ζ' in τ, and so also has τ-length at least V_0.

Note that we have seen that the map $\tilde{p} : \tilde{\alpha} \longrightarrow \tilde{\tau}$ is a quasi-isometry on each component of $\tilde{\alpha}$. Thus, if V_0 is sufficiently large, we can assume that the d_α-lengths of μ, μ' (and μ'') are as large as we want. Moreover, since they follow each other in τ, the paths $q\mu, q\mu'$ (and $q\mu'$) must be parallel in M over most of their length. More precisely, there are subarcs

$\zeta \subseteq \mu$, $\zeta' \subseteq \mu'$ (and $\zeta'' \subseteq \mu''$) so that $q\zeta$, $q\zeta'$ (and $q\zeta''$) are parallel in M, and where each component of $\mu \setminus \zeta$, $\mu' \setminus \zeta'$ (and $\mu'' \setminus \zeta''$) have bounded d_α-length. If we choose V_0 sufficiently large in relation to U_0, we see that, by definition, the bridge arc, ϵ (as well as ϵ_1 and ϵ_2) is U_0-wide, and so ϵ is, by definition U_0-surgerable.

To find such a bridge arc, we proceed as follows. Suppose $q(\alpha) \cap \Delta_3 = \emptyset$. Suppose ω is an arc in α of length U_1 to be determined below. As discussed at the end of Section 3, this places a lower bound on the combinatorial length of $p\omega$ in τ. Now there are only a bounded number of combinatorial possibilities for the image in τ of any subarc of ω of length U_0. Thus, provided U_1 is large enough, we can always find two disjoint subarcs, μ, μ', of ω of length $2U_0$ which follow each other in τ.

With the above choice of U_1, we have:

Lemma 4.4 : *Each curve in $X(\Sigma)$ has a derived curve.*

Proof : In other words, at least one of the cases (1), (2) or (3) of the construction must always arise.

Let $\alpha \in X(\Sigma)$. By Lemma 3.2, if $q(\alpha) \cap \Delta_3 = \emptyset$, $q(\alpha)$ meet $\Delta_3(\delta)$ for some $\delta \in \mathcal{S}$. In particular $\delta \in X(\Sigma)$.

Thus, if neither case (1) or (2) arise, $l_M(\alpha) > U_2$ and $q(\alpha) \cap \Delta_3 = \emptyset$. We now let τ be and associated track, and let $\omega \subseteq \alpha$ be any arc of length U_1. We can now apply the argument above to find a U_0-surgerable bridge arc, and so we are in case (3). $\diamondsuit$

Now for any $\alpha \in X(\Sigma)$, we choose any derived curve $\alpha' \in X(\Sigma)$ and set $\mathrm{proj}(\alpha) = \alpha'$.

Proposition 4.5 :
(1) There is some $l_0 \geq 0$ such that $\mathrm{proj}(X(\Sigma)) \subseteq X(M, l_0)$.
(2) There is some r_0, such that if $\alpha, \beta \in X(\Sigma)$ are adjacent in $\mathcal{G}(\Sigma)$, then $d_{\mathcal{G}}(\mathrm{proj}\,\alpha, \mathrm{proj}\,\beta) \leq r_0$.
(3) Given any $l \geq l_0$, then there is some $r(l)$ such that if $\alpha \in X(M, l)$, then $d_{\mathcal{G}}(\alpha, \mathrm{proj}\,\alpha) \leq r(l)$.

Proof : We know that for any $\gamma \in X(\Sigma)$, the stable length, $l_M^S(\gamma)$ agrees with the shortest length in M, namely $l_M(\gamma)$, up to an additive constant, so it doesn't really matter which we deal with.

The fact that the length of a derived curve is bounded is an immediate consequence of the construction, thereby proving (1). Parts (2) and (3) follow respectively by Proposition 4.2 and Lemma 4.3. $\diamondsuit$

Proof of Theorem 0.2 : This now follows from the fact that $\mathcal{G}(\Sigma)$ is hyperbolic, and the existence of a quasiprojection with the properties described by Proposition 4.6, cf. [Bo2]. $\diamondsuit$

Before leaving this section we elaborate on the discussion of bridge arcs for reference in Section 6. Here we shall deal will a multicurve γ, and we allow a bridge arc to connect different components of γ.

Let x, x' be the endpoints of ϵ. Let ζ, ζ' be the be terminal arcs in γ, i.e. containing x, x' in their interiors. Let $\breve{x} \in \breve{\zeta}$, $\breve{x}' \in \breve{\zeta}'$ be the respective lifts to $\tilde{\gamma}$. Recall that ζ and ζ' are U-long if if each component of $\zeta \setminus \{x\}$ and $\zeta' \setminus \{x'\}$ has d_γ-length at least U. We say that ϵ is *U-wide* if ζ, ζ' are U-long, if $(q\zeta, qx)$, $(q\zeta', qx')$ are parallel in M. In particular, this means that $d(\tilde{q}\breve{x}, \tilde{q}\breve{x}') \leq h_0$. We could weaken this last assumption to placing some other bound on $d(\tilde{q}\breve{x}, \tilde{q}\breve{x}')$. Provided U is large enough in relation to this bound, then this is equivalent to saying that we can we can slide the endpoint x' along ζ to some other point $y \in \zeta$ so that $d(\tilde{q}\breve{x}, \tilde{q}\breve{y}) \leq h_0$. In other words, $(q\zeta, qx)$, $(q\zeta, qy)$ are parallel. In so doing, we reduce the width of the bridge arc by at most a bounded amount. This leads to:

Definition : The bridge arc ϵ is h-almost U-wide if we can slide its endpoints a d_γ-distance at most h along γ so that the resulting bridge arc is U-wide.

Suppose now that $q(\gamma) \cap \Delta_\nu = \emptyset$. Let τ be the associated track. There is an upper bound on the total length of (τ, d_τ). Moreover, since $f(\tau) \cap \Delta_2 = \emptyset$, we see that there is a positive lower bound on the length of any closed curve in τ (not necessarily a trainpath) that is essential in Σ. Now, after splitting τ along a set of short arcs, we can assume that there is a positive lower bound on the length of each branch of τ. These bounds depend ultimately on ν and the parameters of the action. Since we are only interested in d_τ up to uniform bilipschitz equivalence, we may as well assume that each edge of τ has unit length. In other words, we are taking the combinatorial distance on d_τ.

Suppose that ϵ is bridge arc of γ.

Definition : We say that ϵ is *W-wide in τ* if there are W-long terminal arcs ζ, ζ' such that $p\zeta$ and $p\zeta'$ follow each other in τ, and if we can extend p to a map $\gamma \cup \epsilon$ in the given homotopy class so that ϵ gets mapped to a point. We say that ϵ is *h-almost* W-wide if we can slide one of its endpoints a d_γ-distance at most h in γ such that the resulting bridge arc is W-wide. We say that it is *almost* W-wide if it is h-almost W-wide for some h.

We note:

Lemma 4.6 : *There is some W_0 depending only on the parameters of the action such that the following holds. Suppose that $q : \gamma \longrightarrow M$ is a realisation and that $q(\gamma) \cap \Delta_3 = \emptyset$. Let τ be an associated track with combinatorial metric, and carrying map $p : \gamma \longrightarrow \tau$. Let ϵ be a bridge arc, and let $W \geq 0$. If ϵ is $(W + W_0)$-wide in τ, then it is W_0-almost W-wide in M. If ϵ is $(W + W_0)$-wide in M, then it is W_0-almost W-wide in τ.*

Proof : This follows from the earlier discussion, now using Corollary 3.4 in place of Theorem 3.3. $\diamondsuit$

5. Surfaces of bounded geometry.

In this section we describe some basic facts regarding curves on hyperbolic surfaces of bounded geometry, that is, with injectivity radius bounded below. In Section 6, we will apply this to give a formulation of the Uniform Injectivity Theorem in the case where the realisation of a multicurve in M lies in the thick part.

We begin with a purely topological discussion. We say that a subsurface $\Phi \subseteq \Sigma$ is *incompressible* if no component of Φ is a disc, no two annular components are freely homotopic, and no component of $\Sigma \setminus \Phi$ is either a disc or a peripheral annulus. (We are not assuming that Φ to be connected.) We write $\mathcal{F}$ to the set of incompressible surfaces up to homotopy in Σ. We write $\partial\Phi$ for the intrinsic manifold boundary of Φ, and $\partial_\Sigma \Phi = \partial\Sigma \setminus \Phi$ for the relative boundary in Σ.

We write $X(\Sigma, \Phi)$ for the set of curves of $X(\Sigma)$ that can be homotoped into Φ. We can identify this with $X(\Phi) \cup X(\partial_\Sigma \Phi)$, where $X(\partial_\Sigma \Phi)$ is the set of componenents of $\partial_\Sigma \Phi$ defined up to homotopy in Σ. We can partially order Φ by inclusion up to homotopy.

The following is a fairly simple observation:

Lemma 5.1 : *If $Z \subseteq X(\Sigma)$ is any set, then there is a unique minimal $\Phi \in \mathcal{F}$ with $Z \subseteq X(\Sigma, \Phi)$.* $\diamondsuit$

In this case, we write $\Phi(Z) = \Phi$.

Suppose $Q, Q' \subseteq \Sigma$.

Definition : By an *ambient map* from Q to Q', we mean a continuous map, ϕ, from Q to Q', equipped with a continuous extension from Σ to itself, defined up to homotopy relative to Q, and freely homotopic to the identity on Σ.

This is essentially equivalent to saying that ϕ has a preferred lift, $\tilde{\phi} : \tilde{Q} \longrightarrow \tilde{Q}'$, between the preimages of Q and Q' in Σ.

We now fix, once and for all, some $\eta_0 > 0$. By a *thick metric* on Σ, we mean a hyperbolic metric, σ, on Σ, with geodesic boundary, and with injectivity radius at least η_0. This is taken to imply that there is no essential arc of length at most η_0 between boundary components. All we require of η_0 is that such a metric should always exist. (It can be defined independently of type(Σ).) Note that the length of the boundary of Σ is necessarily bounded. If $x, y \in \Sigma$ with $\sigma(x, y) \leq \eta_0$, we write $[x, y]$ for the unique shortest geodesic between them.

Suppose $\eta \leq \eta_0$. By an *η-sequence* in Q, we mean a sequence, $x_0, x_1, \ldots, x_n$, with $x_i \in Q$ and $\sigma(x_i, x_{i+1}) \leq \eta$ for all i. This determines a broken geodesic *η-path*, $[x_0, x_1] \cup [x_1, x_2] \cup \cdots \cup [x_{n-1}, x_n]$. We similarly define a cyclic η-sequence and associated closed path. We write $Z(Q, \eta) \subseteq X(\Sigma)$ for the set of homotopy classes of all closed simple η-paths with vertices in Q. We set $\Phi(Q, \eta) = \Phi(Z(Q, \eta))$. In fact, it turns out that $X(\Sigma, \Phi(Q, \eta)) = Z(Q, \eta)$.

We can give a more intuitive description of $\Phi(Q, \eta)$ as follows. Take a small regular neigbourhood of $N(Q, \eta/2)$, and throw away any disc components. (In any case, there will be no disc components when we actually apply this construction.) Now add in all complementary discs and peripheral annuli. Then identify any pair of homotopic annular components. (Again this last step will not be necessary in applications.) The resulting surface agrees with $\Phi(Q, \eta)$ up to homotopy.

Suppose $h : [0, \infty) \longrightarrow [0, \infty)$ is a homeomorpism, and $\theta : (Q, \sigma) \longrightarrow (Q', \sigma')$ is a map between metric spaces.

Definition : We say that θ is h-continuous if for all $x, y \in Q$, $\sigma'(\theta(x), \theta(y)) \leq h(\sigma(x, y))$. It is h-bicontinuous if it is bijective, and both θ and θ^{-1} are both h-continuous.

Suppose that $Q, Q' \subseteq \Sigma$ and σ, σ' are both thick metics on Σ. Suppose that $\phi : Q \longrightarrow Q'$ is h-continuous. If $(x_i)_i$ is any η-sequence in Q, then $(\phi(x_i))_i$ is an $h(\eta)$ sequence in Q'. If $h(\eta) \leq \eta_0$, then $Y(Q, \eta) \subseteq Y(Q', h(\eta))$ and so $\Phi(Q, \eta) \subseteq \Phi(Q', \eta')$. In particular, if ϕ is h-bicontinuous, we have

$$Y(Q', h^{-1}(\eta)) \subseteq Y(Q, \eta) \subseteq Y(Q', h(\eta))$$

and

$$\Phi(Q', h^{-1}(\eta)) \subseteq \Phi(Q, \eta) \subseteq \Phi(Q', h(\eta)).$$

Definition : If $I \subseteq \mathbf{R}$, we say that Q is I-*stable* if $\Phi(Q, t) = \Phi(Q, u)$ for all $t, u \in I$.

In this case, we write $\Phi(Q, I) = \Phi(Q, t)$ for any $t \in I$.

Thus, in the above discussion, if Q' happens to be $[h^{-1}\eta, h\eta]$-stable with respect to the metric σ, then $\Phi(Q, \eta) = \Phi(Q', [h^{-1}\eta, h\eta])$.

Suppose that γ is a multicurve. We can realise γ as a disjoint union, γ_σ, of closed geodesics in Σ, and we write $l_\sigma(\gamma)$ for its total length.

We will want to compare the realisation of a multicurve, γ, in two different thick structures, σ and σ', on Σ. We can speak of the "quasi-isometric" distance between σ and σ' as the minimal quasi-isometric constant of an equivariant quasi-isometry between the covers $(\tilde{\Sigma}, \tilde{\sigma})$ and $(\tilde{\Sigma}, \tilde{\sigma}')$. Since we are only interested in bounding this quanitity, we could equivalently speak about bilipschitz or quasi-conformal maps etc., though quasi-isometries are natural to deal with in this context.

Note that if $\theta : \gamma_\sigma \longrightarrow \gamma_{\sigma'}$ is an ambient map, then the lift $\tilde{\theta} : \tilde{\gamma}_\sigma \longrightarrow \tilde{\gamma}_{\sigma'}$ is a quasi-isometry with respect to the metrics $\tilde{\sigma}$ and $\tilde{\sigma}'$. Thus, we can also view $\tilde{\theta}$ as a quasi-isometry between $(\tilde{\Sigma}, \tilde{\sigma})$ and $(\tilde{\Sigma}, \tilde{\sigma}')$. Conversely, we can always find such an ambient map, θ, such that the quasi-isometry constants of $\tilde{\theta}$ are bounded in terms of the quasi-isometric distance between σ and σ'.

To see this, start with an equivariant quasi-isometry, $\phi : (\tilde{\Sigma}, \tilde{\sigma}) \longrightarrow (\tilde{\Sigma}, \tilde{\sigma}')$ which we can assume continuous. Any point, $x \in \tilde{\gamma}_\sigma$ lies in a leaf, $\mu \subseteq \tilde{\sigma}$. Now $\phi(\mu)$ lies a bounded distance from μ' for some leaf, μ' of $\tilde{\gamma}_{\sigma'}$. Let $\tilde{\theta}(x)$ be the nearest point to x in μ'. We see that $\tilde{\theta} : \tilde{\gamma}_\sigma \longrightarrow \tilde{\gamma}_{\sigma'}$ is continuous and Γ-equivariant, and that $d(x, \tilde{\theta}(x))$ is bounded above in terms of the original quasi-isometry constants. Thus $\tilde{\theta}$ is a uniformly quasi-isometry. It descends to the required map $\theta : \gamma_\sigma \longrightarrow \gamma_{\sigma'}$.

Another observation is that we can assume θ to be uniformly bilipschitz on each leaf. This follows from the fact that any self-quasi-isometry of the real line is a bounded distance from a uniformly bilipschitz map, and moreover, this construction can be made equivariant with respect to infinite cyclic actions. We also need the fact that the ratio of lengths of

γ_σ and $\gamma_{\sigma'}$ are bounded in terms of the quasi-isometric distance. These facts are all well known. In fact, one can do better:

Lemma 5.2 : *Suppose that γ_σ and $\gamma_{\sigma'}$ are realisations of the multicurve γ with respect to two thick structures, σ and σ' on Σ. Then there is an ambient map $\theta : \gamma_\sigma \longrightarrow \gamma_{\sigma'}$, which is h-bicontinuous with respect to the metrics σ and σ', and ξ-bilipschitz with respect to the induced path metrics, where h and ξ depend only on the quasi-isometric distance between σ and σ'.*

We begin with a few remarks.

The bilipschitz statement is a local property intrinsic to the curves γ_σ, whereas h-bicontinuity refers to the ambient metrics, and therefore makes reference to the "transverse structure" of the curves. In fact, we can take h to have the form $[t \mapsto at^b]$, where $a, b > 0$ depend only on the quasi-isometric distance. In other words, θ is uniformly biholder.

In particular, the map $\tilde{\theta}$ is a uniform quasi-isometry. Note that any two equivariant quasi-isometries remain a bounded distance apart. So the statement can be interpeted as saying that, if we start with an ambient map, then we can homotope it a bounded amount so that it becomes biholder and intrinsically bilipschitz, where the constants depend only on the quasi-isometry constants of the lift of the original map.

Proof of Lemma 5.2 : We split the proof into two stages. We first construct a map that is uniformly biholder, dropping the bilipshitz requirement on γ. For this, we use a fairly standard construction.

If σ and σ' are two thick structures on Σ then there is a uniformly biholder map between the respective projectives tangent bundles, $P\Sigma$ and $P\Sigma'$. First, lift to $\mathbf{H}^2$. Suppose that λ is a tangent line at a point $x \in \mathbf{H}^2$. Consider the endpoints, $y_1, y_2 \in \partial \mathbf{H}^2$, of the bi-infinite geodesics tangent to λ at x, and those, y_3, y_4, of the othogonal bi-infinite geodesics. There is a unique homeomorphism $\partial \mathbf{H}^2 \longrightarrow \partial \mathbf{H}^2$, equivariant with respect to the covering transformations of (Σ, σ) and (Σ, σ'). Let y_1', y_2', y_3', y_4' be the respective images of y_1, y_2, y_3, y_4 under this homeomorphism. The geodesics $[y_1', y_2']$ and $[y_3', y_4']$ intersect at some point, $x' \in \mathbf{H}^2$. The angle between them is bounded below by some positive constant depending only on the type of Σ and the quasi-isometric distance between σ and σ'. Let λ' be the tangent line to $[y_1', y_2']$ at x'. The map between projectivised tangent bundles that sends λ to λ' is biholder. It descends to a biholder homeomorphism from $P\Sigma$ to $P\Sigma'$. The constants depend only on Σ and the quasi-isometric distance. Note that it is invariant under tangential flow. In particular, it gives a biholder map from γ_σ to $\gamma_{\sigma'}$.

We need to modify this so that it is intrinsically bilipschitz on γ_σ. To this end, we find constants $0 < \eta_2 < \eta_1 < \eta_0$, $\eta_3 > 0$, $h_0 \geq 0$ and $m \in \mathbf{N}$, depending only on the topological type of Σ, such that there is some set, $\mathcal{I}$, of at most m geodesic segments in (Σ, σ), each of length at most η_1, with the following properties. If $I \in \mathcal{I}$, then it meets γ_σ transversely at angles at least η_3, and each component of $I \setminus \gamma$ containing and endpoint of I has length at least η_2. Moreover, any two distinct elements of $\mathcal{I}$ are a distance at least η_2 apart, and each component of $\gamma_\sigma \setminus \bigcup \mathcal{I}$ has length at most h_0. The existence of such an $\mathcal{I}$ follows by considering a small metric neighbourhood, N, of γ_σ in (Σ, σ). This will be foliated by geodesics segments of this type, transverse to γ_σ. Since the topology of N is bounded, we restrict to a bounded number of such segments. This is all fairly standard technology developed in relation to train tracks (cf. [PH]). We now modify the map from γ_σ to $\gamma_{\sigma'}$ so that it agrees with the original on $\gamma_\sigma \cap \bigcup \mathcal{I}$, and interpolate linearly on the components of $\gamma_\sigma \setminus \bigcup \mathcal{I}$. It is easily seen to remain biholder and it becomes intrinsically bilipschitz on γ_σ. This then proves Lemma 5.2. $\diamondsuit$

Now suppose that we have an intrinsic path-metric, d_γ, on γ, and an ambient map $j : \gamma \longrightarrow \gamma_\sigma \subseteq (\Sigma, \sigma)$, that is ξ_0-bilipschitz for some fixed constant ξ_0 (in applications depending only on the parameters of the action). On γ_σ, we are taking induced path-metric from σ.

Suppose that ϵ is a bridge arc, with endpoints, x, x' lying in terminal arcs ζ, ζ' (as usual, consistently oriented with respect to ϵ, and irrespective of their relative orientation in γ). Let $\breve{x}, \breve{x}', \breve{\zeta}, \breve{\zeta}'$ be their respective lifts to $\tilde{\gamma}$. Following the terminology of Section 4, we say that $j\zeta$ and $j\zeta'$ are *parallel* if $\tilde{j}\breve{\zeta}, \tilde{j}\breve{\zeta}'$ are at most a unit distance apart in $\tilde{\Sigma}$ (in the sense that their respective initial and final points are unit distance apart). Of course, the choice of unity to bound distance is arbitrary here.

Definition : Given $V \geq 0$, we say that ϵ is V-*wide* in (Σ, σ) if we can find V-long terminal arcs, ζ, ζ' containing the endpoints x, x' so that $j\zeta$ and $j\zeta'$ are parallel and $\tilde{\sigma}(\tilde{j}\breve{x}, \tilde{j}\breve{x}') \leq 1$. We say that ϵ is h-almost V-wide if we can slide its endpoints a distance at most h along γ so that it becomes V-wide. We say that it is almost V-wide if it is h-almost V-wide for some h.

Note that being almost V-wide depends only on the homotopy class of ϵ relative to γ.

Now suppose that $Q \subseteq \gamma$ is a closed subset.

Definition : By a *bridge path* (with respect to Q) we mean a closed path in Σ that is homotopically simple and non-trivial in Σ, and has the form $\alpha_1 \cup \epsilon_1 \cup \alpha_2 \cup \epsilon_2 \cdots \cup \alpha_n \cup \epsilon_n$, where each α is an interval in γ, and each ϵ_i is a bridge arc (with endpoints in Q).
We say that it is V-wide if each ϵ_i is V-wide. We say that it is h-almost V-wide if each ϵ_i is h-almost V-wide.

Let $Z_\sigma(Q, V, h)$ be the set of h-almost V-wide bridge paths, and $Z_\sigma(Q, V)$ be the set of all almost V-wide bridge paths. We regard this as defined up to homotopy in Σ, so that $Z_\sigma(Q, V, h), Z(Q, V) \subseteq X(\Sigma)$. Let $\Phi_\sigma(Q, V, h) = \Phi(Z_\sigma(Q, V, h))$ and $\Phi_\sigma(Q, V) = \Phi(Z_\sigma(Q, V))$. Recall the notation $\Phi(Q, \eta)$ defined for $\eta > 0$, defined earlier in this Section.

Lemma 5.3 : *$(\forall V \geq 0)(\exists \eta(V) > 0)$ and $(\forall \eta > 0)(\exists V(\eta) \geq 0)$ such that if $j : \gamma \longrightarrow \gamma_\sigma \subseteq \Sigma$ is a uniformly bilipschitz embedding, then*

$$\Phi_\sigma(\gamma, V(\eta)) \subseteq \Phi(\gamma_\sigma, \eta)$$

and

$$\Phi(\gamma_\sigma, \eta(V)) \subseteq \Phi_\sigma(\gamma, V).$$

Here the functions $[\eta \mapsto V(\eta)]$ and $[V \mapsto \eta(V)]$ depend only on $\mathrm{type}(\Sigma)$ *and the bilipschitz constant.*

Proof : This follows from the standard fact of hyperbolic geometry that two disjoint geodesic segments in the plane remain close over a long distance if and only if they are exponentially close near their midpoints. $\diamondsuit$

There is a similar statement for a subset $Q \subseteq \gamma$.

Lemma 5.4 : *Suppose that $Q \subseteq \gamma$ and $h \geq 0$, then there is some $r(h) \geq 0$ such that if each component of Q has length at least $r(h)$, then for all $\eta > 0$, $V \geq 0$, we have*

$$\Phi_\sigma(Q, V(\eta), r(h)) \subseteq \Phi(j(Q), \eta)$$

and

$$\Phi(j(Q), \eta(V)) \subseteq \Phi_\sigma(Q, V, h),$$

where the functions $[\eta \mapsto V(\eta)]$ and $[V \mapsto \eta(V)]$ are as in Lemma 5.3. $\diamondsuit$

We want to relate this to multitracks. We shall assume here that each multitrack, τ, in Σ is given the combinatorial metric, so that each branch has unit length. We note:

Lemma 5.5 : *Given any multitrack τ in Σ, there is a thick metric, σ, on Σ and a realisation, $\tau \hookrightarrow \Sigma$, such that the lift, $\tilde{\tau} \hookrightarrow \tilde{\Sigma}$, is a unformly quasi-isometric embedding.*

Here "uniform" means that the constants of quasi-isometry depend only on $\mathrm{type}(\Sigma)$.

Proof : This is certainly true of any given multitrack, and the result follows from the observation that there are only finitely many combinatorial possibilities for a multitrack, modulo the action of the mapping class group of Σ. $\diamondsuit$

We will refer to an embedding of τ in Σ arising in this way as a *uniform embedding*.

Lemma 5.6 : *Suppose that $p : \gamma \longrightarrow \tau$ is a uniformly locally bilipschitz carrying map, and that $\tau \hookrightarrow \Sigma$ is a uniform embedding with respect to a thick metric, σ, on Σ. Then there is a bilipschitz map, $j : \gamma \longrightarrow \gamma_\sigma \subseteq \Sigma$, such that for all $x \in \tilde{\gamma}$, the distance $\tilde{\sigma}(\tilde{p}(x), \tilde{j}(x))$ is bounded above. Here, the bounds and bilipschitz constant depend only on* $\mathrm{type}(\Sigma)$ *and the bilipschitz constant of the carrying map.*

Proof : This is similar to the earlier construction. If μ is a leaf of $\tilde{\gamma}$, then $p(\mu)$ is uniformly quasigeodesic in $(\tilde{\Sigma}, \tilde{\sigma})$, and hence a bounded distance from its realisation as a bi-infinite geodesic μ'. We can thus take nearest point projection, and then approximate by a uniformly bilipschitz map. (Note that we need not worry about the transverse structure here.) $\diamondsuit$

We refer to j as an "associated realisation".

Suppose ϵ is a bridge arc of γ, and $W \geq 0$. Recall the notion of ϵ being W-wide, or h-almost W-wide from the end of Section 4.

Lemma 5.7 : *There is some $V_0 \geq 0$, depending only on* type(Σ) *and ξ_0 with the following property. Suppose that $p : \gamma \longrightarrow \tau$ is a uniform embedding, and $j : \gamma \longrightarrow \gamma_\sigma \subseteq \Sigma$ be an associated realisation. Let ϵ be a bridge arc.*

If ϵ is $(V + V_0)$-wide in τ, then it is V_0-almost V-wide in Σ.

If ϵ is $(V + V_0)$-wide in Σ, then it is V_0-almost V-wide in τ.

Proof : Let ζ, ζ' be terminal arcs of ϵ.

If ϵ is $(V + V_0)$-wide in τ, then $p\zeta$ and $p\zeta'$ follow each other in τ, over a distance linearly bounded below by $V + V_0$. The lifts $\tilde{p}\breve{\zeta}$ and $\tilde{p}\breve{\zeta}'$ to $\tilde{\tau}$ are identical. Since $\tilde{\tau} \hookrightarrow \tilde{\Sigma}$ is a quasi-isometric embedding and since $\tilde{j} : \tilde{\gamma} \longrightarrow \tilde{q}$ is a bounded distance from $\tilde{p}$, we see that $\tilde{j}\breve{\zeta}$ and $\tilde{j}\breve{\zeta}'$ are quasigeodesic and a bounded distance apart. Thus, their images remain a uniformly bounded distance apart over most of their length, say by 1. In particular, modulo sliding one of the endpoints of ϵ a bounded distance we can assume that $\tilde{\sigma}(\tilde{j}\breve{x}, \tilde{j}\breve{x}') \leq 1$. It follows that, provided we take V_0 large enough, then ϵ is V_0-almost V-wide in Σ.

Conversely, suppose that ϵ is $(V + V_0)$-wide in Σ. The same argument shows that $\tilde{p}\breve{\zeta}$ and $\tilde{p}\breve{\zeta}'$ as well as $\tilde{p}\breve{x}$ and $\tilde{p}\breve{x}'$, remain a bounded distance apart. We claim that $\tilde{p}\breve{\zeta}$ and $\tilde{p}\breve{\zeta}'$ follow each other over most of their length. It follows, again choosing V_0 sufficiently large, that ϵ is V_0-almost V-wide in τ. $\diamondsuit$

For the last part of the above argument we need the following observation regarding multitracks.

Lemma 5.8 : *Given $r \geq 0$, there is some $l \geq 0$ depending only on r and* type(Σ) *with the following property. Suppose that $\tau \hookrightarrow \Sigma$ is a uniform embedding and $\tilde{\tau} \subseteq \tilde{\Sigma}$ is its lift to $\tilde{\Sigma}$. Suppose that μ, μ' are oriented trainpaths of length at least l in $\tilde{\tau}$ that remain distance at most r apart in Σ. Then μ and μ' intersect positively.*

Proof : Let e_i and e'_i be the ith directed branches of μ and μ' respectively. For each i, there is some $k(i)$ such that the distance in σ between e_i and $e'_{k(i)}$ is at most r. Thus, there are only boundedly many possibilities for the pair $(e_i, e'_{k(i)})$ modulo the covering transformations. Thus, if l is sufficiently large, we can find $i < j$ and a covering translation g with $e_j = ge_i$ and $e'_{k(j)} = ge'_{k(i)}$. We can also assume that the distance between e_i and e_j is large in relation to r. Now the segment, $\zeta \subseteq \mu$, between e_i and e_j and the segment, ζ', between $e'_{k(i)}$ and $e'_{k(j)}$ project to closed trainpaths in τ that are freely homotopic in Σ and hence equal as closed trainpaths. In particular, $k(j) - k(i) = j - i$, and ζ and ζ' lie in the same bi-infinite axis in $\tilde{\tau}$ and are oriented in the same sense along this axis. Since their length is greater than r, they must intersect positively. $\diamondsuit$

6. Bounded geometry form of the unform injectivity theorem.

In this section, we give an alternative form of the uniform injectivity theorem in the case where the realisation of our multicurve lies in the thick part of M. This statement makes no reference to tracks, and is more in line with the standard version for hyperbolic 3-manifolds.

We return to the coarse hyperbolic manifold M with $\Gamma \cong \pi_1(\Sigma) \cong \pi_1(M)$. As in Section 4, we write Δ_ν for the thin part $\Delta_\nu = \Delta(r_0, \nu)$.

Proposition 6.1 : *Suppose that γ is a multicurve and $q : \gamma \longrightarrow M$ is a realisation. Let d_γ be the induced path-metric on γ. Suppose $q(\gamma) \cap \Delta_\nu = \emptyset$. Then there is a bounded geometry structure, σ, in Σ and a map $j : \gamma \longrightarrow \gamma_\sigma$ to the realisation of γ in Σ which is ξ_0-bilipschitz in the induced path metric and such that the following hold. Let ϵ be a bridge arc of γ. If ϵ is $(U + U_0)$-wide in M, then it is U_0-almost U-wide in (Σ, σ). If ϵ is $(U + U_0)$-wide in (Σ, σ), then it is U_0-almost U-wide in M. Here ξ_0 and U_0 depend only on ν the parameters M.*

Proof : Let τ be an associated multitrack and $p : \gamma \longrightarrow \tau$ be the carrying map. Since $q(\gamma) \cap \Delta_\nu = \emptyset$, as discussed at the end of Section 4, we can assume that τ has the combinatorial metric. Let $\tau \longrightarrow (\Sigma, \sigma)$ be a uniform embedding as given by Lemma 5.6, and let $j : \gamma \longrightarrow \gamma_\sigma \subseteq \Sigma$ be an associated realisation as given by Lemma 5.7.

Now let V_0 be the constant of Lemma 5.8, let W_0 be the constant of Lemma 4.6, and set $U_0 = V_0 + W_0$. Suppose that ϵ is a bridge arc. If it is $(U + U_0)$-wide in (Σ, σ), then it is V_0-almost $(U + W_0)$-wide in τ, and hence U_0-almost U-wide in M. Similarly, if it is $(U+U_0)$-wide in M, it is W_0-almost $(U+V_0)$-wide in τ, hence U_0-almost U-wide in (Σ, σ). $\diamondsuit$

We refer to $j : \gamma \longrightarrow \gamma_\sigma \subseteq \Sigma$ as an *associated realisation*.

Recall from Section 5 that a bridge path has the form $\beta = \alpha_1 \cup \epsilon_1 \cup \alpha_2 \cup \epsilon_2 \cup \cdots \cup \alpha_n \cup \epsilon_n$, where each α_i is an arc in γ and each ϵ_i is a bridge. We also assume that it is homotopically simple in Σ (though this makes no essential difference to the discussion). We say that β is *U-wide* in M if each ϵ_i is U-wide.

If $Q \subseteq \Sigma$, we define $\Phi_M(Q, U, h)$ and $\Phi_M(Q, U)$ exactly as with $\Phi_\sigma(Q, U, h)$ and $\Phi_\sigma(Q, U)$. An immediate consequence of Proposition 6.1 is the following.

Lemma 6.2 : *Let U_0 be the constant given by Proposition 6.1. Suppose that γ is a multicurve, and $q : \gamma \longrightarrow M$ is a realisation with $q(\gamma) \cap \Delta_\nu = \emptyset$, and let $q : \gamma \longrightarrow (\Sigma, \sigma)$ be an associated realisation. Then for all $U \geq 0$, we have*

$$\Phi_\sigma(\gamma, U + U_0) \subseteq \Phi_M(\gamma, U)$$
$$\Phi_M(\gamma, U + U_0) \subseteq \Phi_\sigma(\gamma, U).$$

Putting this together with Lemma 5.4, we get:

Lemma 6.3 : $(\forall U \geq 0)(\exists \eta(U) > 0)$ *and* $(\forall \eta > 0)(\exists U(\eta) \geq 0)$ *such that if $q : \gamma \longrightarrow M$ and $j : \gamma \longrightarrow (\Sigma, \sigma)$ are as in Lemma 6.2, then we have*

$$\Phi(\gamma_\sigma, \eta(U)) \subseteq \Phi_M(\gamma, U)$$
$$\Phi_M(\gamma, U(\eta)) \subseteq \Phi(\gamma_\sigma, \eta(U)).$$

$\diamondsuit$

In particular, suppose $\eta \in (0, \eta_0)$ and set $I_\eta = [\eta(U(\eta)), \eta] \subseteq \mathbf{R}$. Suppose that γ_σ is I_η-stable in (Σ, σ), then $\Phi_M(\gamma, U(\eta)) = \Phi(\gamma_\sigma, I_\eta)$.

If ϵ is a bridge with endpoints x to x'. We denote by $j\epsilon$ its realisation in (Σ, σ) as a geodesic with the endpoints fixed at jx and jx'. We write $l_\sigma(\epsilon)$ for the length of $j\epsilon$.

Suppose that $q : \gamma \longrightarrow M$ is a realisation with $q(\gamma) \cap \Delta_\nu = \emptyset$. Suppose that $j : \gamma \longrightarrow (\Sigma, \sigma)$ and $j' : \gamma \longrightarrow (\Sigma, \sigma')$ are two associated realisations. We want to consider how they are related.

Lemma 6.4 : *There are constants $\eta > 0, l \geq 0$, depending only on the parameters of the action, such that if ϵ a bridge arc with $l_\sigma(\epsilon) \leq \eta$, then $l_{\sigma'}(\epsilon) \leq l$.*

Proof : As in the proof of Lemma 5.5, we see that if η is small enough, then η is $(2V_0)$-wide in (Σ, σ). By Proposition 6.1, it is V_0-almost V_0-wide in M, hence $(2V_0)$-almost 0-wide in (Σ, σ'). The last statement means that we can slide the endpoint of ϵ a bounded distance (with respect to the metric d_γ, hence also with respect to σ') so that the geodesic realisation of ϵ in (Σ, σ') has length at most 1. This implies that its original realisation, $j'\epsilon$, is also of bounded length. $\diamondsuit$

We can rephrase this in terms of the universal cover as follows. Let $\phi : j' \circ j^{-1} : \gamma_\sigma \longrightarrow \gamma_{\sigma'}$. This has a lift $\tilde{\phi} : \tilde{\gamma}_\sigma \longrightarrow \tilde{\gamma}_{\sigma'}$. Lemma 6.4 says that if $x, y \in \tilde{\gamma}_\sigma$ with $\tilde{\sigma}(x, y) \leq \eta$, then $\tilde{\sigma}'(\tilde{\phi}x, \tilde{\phi}y) \leq l$.

Our aim is to prove the following:

Lemma 6.5 : *Suppose $q : \gamma \longrightarrow M$ is a realisation of the multicurve and $j : \gamma \longrightarrow \gamma_\sigma \subseteq (\Sigma, \sigma)$ and $j' : \gamma \longrightarrow \gamma_{\sigma'} \subseteq (\Sigma, \sigma')$ are associated realisations of γ. Then there is a homeomorphism $\psi : \gamma_\sigma \longrightarrow \gamma_{\sigma'}$ which is h-bicontinuous and ξ_0-bilipschitz with respect to the induced path metric, where the function h and the constant ξ_0 depend only on the parameters of M.* $\diamondsuit$

Note that we have not assumed here that ψ is an ambient map. In fact, the construction will give an equivariant h-bicontinuous lift $\tilde{\psi} : \tilde{\gamma}_\sigma \longrightarrow \tilde{\gamma}_{\sigma'}$. though the associated mapping class of Σ need not be trivial.

For the proof, we need the following construction. If $\beta = \alpha_1 \cup \epsilon_1 \cup \alpha_2 \cup \epsilon_2 \cup \cdots \cup \alpha_n \cup \epsilon_n$, and $j : \gamma \longrightarrow (\Sigma, \sigma)$ is realisation of γ, then we obtain a map, $j : \beta \longrightarrow \Sigma$, by sending each bridge ϵ_i to the corresponding geodesic in (Σ, σ). We say that β is a (η, h)-*bridge path* if the σ-length of each $j\epsilon_i$ is at most η and the σ-length of each $j\alpha_i$ is at most h.

Lemma 6.6 : *Let $j : \gamma \longrightarrow \gamma_\sigma \subseteq (\Sigma, \sigma)$ be a realisation of γ, and $\eta > 0$. Then there exist $n \in \mathbf{N}$, depending only in* $\mathrm{type}(\Sigma)$ *and $h > 0$, depending only on η and* $\mathrm{type}(\Sigma)$ *such that each component of $\partial_\Sigma \Phi(\gamma_\sigma, \eta)$ is homotopic in Σ to a (η, h)-bridge path with at most n-bridges.*

Proof : This follows from the description of $\Phi(\gamma_\sigma, \eta)$ in terms of $N(\gamma_\sigma, \eta/2)$. $\diamondsuit$

Corollary 6.7 : *Suppose that j, j' are realisations of the same multicurve, γ and $\eta \le \eta_0$. Then there is some $l \ge 0$, depending only on η and* $\mathrm{type}(\Sigma)$ *such that $l_\sigma(\partial_\Sigma \Phi(\gamma_\sigma, \eta)) \le l$ and $l_{\sigma'}(\partial_\Sigma \Phi(\gamma_\sigma, \eta)) \le l$.*

Here $\Phi(\gamma_\sigma, \eta)$ is constructed with respect to the metric σ and is defined up to homotopy in Σ.

Proof : The first statement follows directly from Lemma 6.6. For the second, we realise each component of $\partial_\Sigma \Phi(\gamma_\sigma, \eta)$ as a bridge path β using Lemma 6.6. Applying Lemma 6.4 and the fact that realisations are assumed uniformly bilipschitz with respect to the induced path metrics, we see that the length of the realisation of any such β in (Σ, σ') has bounded length. Thus, $l_{\sigma'}(\partial_\Sigma \Phi(\gamma_\sigma, \eta))$ is bounded. $\diamondsuit$

Note that we can also carry out the above construction in a subsurface, Ψ, of Σ. Suppose that $\gamma \subseteq \Psi$. Then $\gamma_\sigma \subseteq \Psi_\sigma$, where Ψ_σ is realised so that it has geodesic boundary. Given $\eta > 0$, we can define $\Phi_\Psi(\gamma_\sigma, \eta)$ intrinsically to Ψ_σ. Note that Lemma 6.6 and 6.7 still hold, though the constants will now depend also on $l_\sigma(\partial\Psi)$.

Lemma 6.8 : *Suppose that j, j' are realisations of γ, satifying the conclusion of Lemma 6.4 (with constants, η and k). Then there is some $l \ge 0$, depending only on η, k, and* $\mathrm{type}(\Sigma)$, *such that there is a subsurface $\Psi \subseteq \Sigma$, with $\gamma \subseteq \Psi$, satisfying $l_\sigma(\partial\Psi) \le l$, $l_{\sigma'}(\partial\Psi) \le l$, $\Phi_\Psi(\gamma_\sigma, \eta) = \Psi$ and $\Phi_\Psi(\gamma_{\sigma'}, \eta) = \Psi$.* $\diamondsuit$

Proof : We go back and fore between (Σ, σ) and (Σ, σ'). Let $\Phi_1 = \Phi(\gamma_\sigma, \eta)$, $\Phi_2 = \Phi_{\Phi_1}(\gamma_{\sigma'}, \eta)$, $\Phi_3 = \Phi_{\Phi_2}(\gamma'_\sigma, \eta)$, etc. Thus, Φ_i is a descending sequence of subsurfaces of Σ. After a bounded number of steps, there must be three consecutive subsurfaces, all equal to a fixed subsurface Ψ. By the above observation, $l_\sigma(\partial\Phi_i)$ and $l_{\sigma'}(\partial\Phi_i)$ are bounded inductively, and so in particular, $l_\sigma(\partial\Psi)$ and $l_{\sigma'}(\partial\Psi)$ are bounded. By construction, $\Phi_\Psi(\gamma_\sigma, \eta) = \Phi_\Psi(\gamma_{\sigma'}, \eta) = \Psi$. $\diamondsuit$

Proof of Lemma 6.5 : Let η, h be the constants of Lemma 6.4. We now apply Lemma 6.8 to give us the subsurface Ψ. Let Ω be a component of Ψ, and let $\alpha = \gamma \cap \Omega$. Thus, α has a realisations $\alpha_\sigma \subseteq \Omega_\sigma$ and $\alpha_{\sigma'} \subseteq \Omega_{\sigma'}$. Let $\tilde\alpha_\sigma \subseteq \tilde\Omega_\sigma$ and $\tilde\alpha_{\sigma'} \subseteq \tilde\Omega_{\sigma'}$ be the lifts to the universal covers. By construction, any two points of $\tilde\alpha_\sigma$ and $\tilde\alpha_\sigma$ are connected by an η-path. Moreover, by Lemma 6.4, if $x, y \in \tilde\alpha_\sigma$ with $\tilde\sigma(x, y) \le \eta$, then $\tilde\sigma'(\tilde\phi x, \tilde\phi y) \le k$, and similarly interchanging $\tilde\alpha_\sigma$ and $\tilde\alpha_{\sigma'}$. Thus, $\tilde\phi$ is a quasi-isometry. Applying Lemma 5.3 (with Ω in place of Σ) we we get another map, $\psi : \tilde\alpha_\sigma \longrightarrow \tilde\alpha_{\sigma'}$, which is uniformly bicontinuous, and bilipschitz in the induced path metrics. We now do this for every component, Ω, of Ψ, and piece them together to give $\psi : \tilde\gamma_\sigma \longrightarrow \tilde\gamma'_\sigma$. Note that if Ω and Ω' are distinct components of Ψ, then there is a positive lower bound $\sigma(\gamma_\sigma \cap \Omega'_\sigma, \gamma_\sigma \cap \Omega_\sigma)$ and on $\sigma'(\gamma_{\sigma'} \cap \Omega_{\sigma'}, \gamma_{\sigma'} \cap \Omega'_{\sigma'})$

(arising from the upper bounds on $l_\sigma(\Psi)$ and $l_{\sigma'}(\Psi)$). Thus ψ has the required properties. $\diamondsuit$

7. Tight geodesics.

In this section we give proofs of Theorems 0.1 and 0.3 for non-exceptional surfaces. More exactly, we describe how the proofs of the corresponding statements in [Bo2] go through in this more general situation. The statements are essentially identical except that in [Bo2] we assumed that M was a hyperbolic 3-manifold. We made much use of the existence of a supply of "pleating surfaces" in M, in terms of which the uniform injectivity theorem was expressed. The main purpose of this is to relate the geometry in some region of M to the geometry in Σ with some hyperbolic metric. Here we have already established the essential requirements in terms of the uniform injectivity theorem that we have formulated. Indeed that represents most of the work in adapting these arguments. What remains is largely an exercise in translating the terminology.

Let M be a coarse hyperbolic manifold (as defined in Section 0). Recall that $X(M,l) = \{\alpha \in X(\Sigma) \mid l_M(\alpha) \leq l\}$. We have so far shown that $X(M,l_0)$ is uniformly quasiconvex in $\mathcal{G}(\Sigma)$ by defining a quasiprojection, proj : $X(M) \longrightarrow X(M,l_0)$. We next recall the definition of a *tight geodesic* in $\mathcal{G}(\Sigma)$. (This is slight generalisation of that in [MaM] as used in [Bo2].)

If γ is a multicurve, we write $X(\gamma) \subseteq X(\Sigma)$ for the set of components of γ. If γ and δ are multicurves, we say that they are *exactly distance* r *apart* if $d(\alpha,\beta) = r$ for all $\alpha \in X(\gamma)$ and $\beta \in X(\delta)$. A sequence of multicurves $(\gamma_i)_i$ is an *multigeodesic* if for all $i < j$, γ_i and γ_j are at exactly distance $j - i$ apart. We say that it is *tight* at index i if each curve that crosses some curve of $X(\gamma_i)$ also crosses some curve of $X(\gamma_{i-1}) \cup X(\gamma_{i+1})$. We say that the multigeodesic is "tight" if it is tight for all indices.

A geodesic, $(\alpha_i)_i$ in $\mathcal{G}(\Sigma)$ is defined in the usual way as a multigeodesic whose elements, α_i, are all curves. A *tight geodesic* if there exists a tight multigeodesic, $(\gamma_i)_i$, such that $\alpha_i \in X(\gamma_i)$ for all i. (A complication in this terminology is that a tight geodesic need not be tight when considered as a multigeodesic.)

Let M be a space satisfying (M1)–(M3). We shall fix a realisation, $j : \alpha \longrightarrow \alpha_M \subseteq M$ for each $\alpha \in X(\Sigma)$. By combining these, we get a fixed realisation of each multicurve.

The first thing we need to do is establish analogues of the "tube penetration" lemmas in [Bo2] (these being rephrasings of similar results in [Mi]). In particular, we need versions of Lemmas 5.1, 5.2 and 5.3. A "Margulis tube" (denoted $T(\delta,\eta)$ in [Bo2]) is replaced by a region of the form $\Delta(\delta; r_0, \nu)$. Here $\delta \in Y_0(\Sigma)$ is a primative closed curve in Σ. The number $\nu \in \mathbf{N}$ replaces the "Margulis constant", η, of $T(\delta,\eta)$. (In the case where M is a hyperbolic 3-manifold, η would be roughly exponential in $-\nu$.) It is no longer clear, a-priori, that the curve δ is simple in Σ, though we shall see this to be so in the cases of interest to us. In the statement of the lemma, the various constants arising will depend on the parameters of the action (and not just on type(Σ) as in [Bo2]).

The strongest form of the tube penetration lemma in [Bo2] was Lemma 5.2, so we give an explicit reformulation of this as follows. Let l_0 be the constant of Theorem 0.2, so that $X(M, l_0)$ is uniformly quasiconvex.

Lemma 7.1 : *Given $h \geq 0$ and $\nu \in \mathbf{N}$, there is some $\nu' \geq \nu$ depending only on the parameters of the action and ν, h with the following property. Suppose that $\delta \in Y_0(\Sigma)$. Suppose that $\gamma_0, \ldots, \gamma_p$ is a multigeodesic which is tight at index i, and with $d(\gamma_0, X(M, l_0)) \leq h$ and $d(\gamma_p, X(M, l_0)) \leq h$. If $q(\gamma_0) \cap \Delta(\delta; r_0, \nu) \subseteq q(\delta)$ and $q(\gamma_p) \cap \Delta(\delta; r_0, \nu) \subseteq q(\delta)$, then $q(\gamma_i) \cap \Delta(\delta; r_0, \nu') \subseteq q(\delta)$.*

(Here the statement $q(\gamma) \cap \Delta(\delta; r_0, \nu) \subseteq q(\gamma)$ is interpreted to imply that $q(\alpha) \cap \Delta(\delta; r_0, \nu) = \emptyset$ for all $\alpha \in X(\gamma) \setminus \{\delta\}$.)

Proof : First, note that we can assume that $\nu \geq \nu_3$, so that $\Delta(\delta; r_0, \nu) \subseteq \Delta_3(\delta)$. Here, ν_3 is the constant used in defining quasiprojection in Section 4. Thus, if $\alpha \in X(\Sigma)$, and $q(\alpha) \cap \Delta(\delta; r_0, \nu) = \emptyset$, then $\delta \in X(\Sigma)$ and $d(\delta, \mathrm{proj}(\alpha))$ is bounded. (This follows directly from the definition of quasiprojection.) Since $X(M, l_0)$ is uniformly quasiconvex, $\bigcup_i X(\gamma_i)$ lies in a uniform neighbourhood of $X(M, l_0)$ depending on h. From standard properties of quasiprojection, it then follows that if $\alpha \in X(\Sigma)$, then $d(\alpha, \mathrm{proj}(\alpha))$ is bounded. In summary, we see that if $\alpha \in X(\gamma_i)$ for any i, and if $q(\alpha) \cap \Delta(\delta; r_0, \nu) \neq \emptyset$, then $\delta \in X(\Sigma)$ and $d(\alpha, \delta) \leq R$, where R depends only on ν, h and the parameters of M.

We now find ν' so that $N(\Delta(\delta; r_0, \nu'), 2K_1(2R+1)) \subseteq \Delta(\delta; r_0, \nu)$, where K_1 is the constant in property (2) of an associated track. We claim that this has the desired properties.

For notational convenience, we shift indices to give us a multigeodesic $\gamma_{-m}, \ldots, \gamma_0, \ldots, \gamma_r$, which is tight at index 0, and where $(q(\gamma_{-m}) \cup q(\gamma_r)) \cap \Delta(\delta; r_0, \nu) \subseteq q(\delta)$. Suppose, for contradiction, that there is some $\alpha_0 \in X(\gamma_0) \setminus \{\delta\}$ and some $y_0 \in \alpha_0$, with $q(y_0) \in \Delta(\delta; r_0, \nu')$.

Now certainly $q(y_0) \in \Delta_3(\delta)$. Let τ_0 be an associated track to α_0. We have maps $p_0 : \alpha_0 \longrightarrow \tau_0$ and $f_0 : \tau_0 \longrightarrow M$. Now, in Section 3, ν_1 was chosen so that necessarily $f_0 p_0(y_0) \in \Delta_1(\delta)$ and $f_0^{-1}(\Delta_1(\delta)) \subseteq A_1(\delta)$, where $A_1(\delta)$ is the special annular neighbourhood of $\delta \in \mathcal{S}_1 \subseteq \mathcal{S}_0$. From the description of $\tau_0 \cap A_1(\delta)$, (since $A_1(\delta)$ is a special annulus), it follows that α_0 crosses δ. We write $A_1(\delta, \tau_0) = A_1(\delta)$.

Since $(\gamma_i)_i$ is tight at 0, it follows that either γ_{-1} for γ_1 also crosses δ. Let $\sigma_{-s}, \ldots, \gamma_0, \ldots, \gamma_t$ be a maximal sequence of consecutive multicurves, all of which cross δ. Thus, $s + t \geq 1$. We claim that either $s = m$ or $t = r$. For if not, $\delta \cup \gamma_{-s-1}$ and $\delta \cup \gamma_{t+1}$ are both multicurves, and so $d(\gamma_{-s-1}, \gamma_{t+1}) \leq 2$, contradicting the assumption that $(\gamma_i)_i$ is a multigeodesic. We can thus suppose that $t = r$. In particular, δ does not lie in γ_r. Since, by assumption, $q(\gamma_r) \cap \Delta(\delta; r_0, \nu) \subseteq q(\delta)$, we see that, in fact, $q(\gamma_r) \cap \Delta(\delta; r_0, \nu) = \emptyset$.

Now let τ_1 be an associated track for $\gamma_0 \cup \gamma_1$. We have maps $p_1 : \gamma_0 \cup \gamma_1 \longrightarrow \tau_1$ and $f_1 : \tau_1 \longrightarrow M$. Again, $f_1 p_1(y_1) \in \Delta_1(\delta)$ and $f^{-1}(\Delta_1) \subseteq A_1(\Sigma) = A_1(\Sigma, \tau_1)$. Since γ_1 crosses δ, it follows that $\tau_1 \cap A_1(\delta) \subseteq p_1(\gamma_1)$. In particular, there is some $y_1 \in \gamma_1$ with $f_1 p_1(y_1) = f_1 p_1(y_0)$. Now $d_M(q(y_0), f_1 p_1(y_0)) \leq K_1$ and $d_M(q(y_1), f_1 p_1(y_1)) \leq K_1$. In particular, $d_M(q(y_0), q(y_1)) \leq 2K_1$. Thus $q(y_1) \in N(\Delta(\delta; r_0, \nu), 2K_1) \subseteq \Delta_3(\delta)$.

We can now apply the same argument again, with γ_1, γ_2 replacing γ_0, γ_1 respectively. This time, we find $y_2 \in \gamma_2$ with $d_M(q(y_1), q(y_2)) \leq 2K_1$, and so $d_M(q(y_0), q(y_2)) \leq 4K_1$.

We continue this process inductively. We chose ν' so that $N(\Delta(\delta; r_0, \nu'), 2K_1(2R +$

$1)) \subseteq \Delta(\delta; r_0, \nu) \subseteq \Delta_3(\delta)$. Thus, provided $i \leq 2R + 1$, we can find $y_i \in \gamma_i$, with $d_M(q(y_0), q(y_i)) \leq 2K_1 i$. In particular, $q(\gamma_i) \cap \Delta(\delta; r_0, \nu) = \emptyset$. But we know that $q(\gamma_r) \cap \Delta(\delta; r_0, \nu) = \emptyset$, and so it follows that $r \geq 2R + 1$. From the definition of R in the earlier discussion, it follows that $d(\gamma_0, \delta) \leq R$ and $d(\gamma_0, \delta) \leq R$. Thus, $d(\gamma_0, \gamma_{2R+1}) \leq 2R$, contradicting the assumption that $(\gamma_i)_i$ is geodesic. $\diamondsuit$

We now move on to consider the construction of subsurfaces. The idea is that if we have a multicurve in Σ that is very long in M, then it fills up a certain subsurface of Σ. One can then see that the subsurfaces associated to consecutive elements of a tight multigeodesic have nice intersection properties, allowing us to shortcut the geodesic, thereby giving a contradiction.

First, we consider a sequence of multicurves, $(\gamma^n)^n$, in Σ. Suppose that they have realisations $j^n : \gamma^n \longrightarrow \gamma^n_\sigma$ in (Σ, σ) with respect to a fixed hyperbolic metric, σ. Passing to a subsequence, we can write $\alpha^n = \check{\alpha}^n \sqcup \hat{\alpha}^n$, where each component of $\check{\alpha}^n_\sigma$ has bounded length, and the length of each component of $\hat{\alpha}^n_\sigma$ tends to ∞. Again passing to a subsequence, we can suppose that $\check{\alpha}^n = \beta$ is fixed, and that $\hat{\alpha}^n$ converges on a lamination, $\lambda \subseteq \Sigma$. Let $\mu(\lambda)$ be the union of all minimal sublaminations of λ. We construct subsurfaces $F(\underline{\alpha})$, $G(\underline{\alpha})$, $H(\underline{\alpha})$ as follows. The surface $H(\underline{\alpha})$ is just a regular neigbourhood of β, $G(\underline{\alpha})$ is a small neighbourhood of λ with complementary discs and peripheral annuli added in, and $F(\underline{\alpha})$ is the same thing using $\mu(\lambda)$. Note that $F(\underline{\alpha}) \subseteq G(\underline{\alpha})$ and $G(\underline{\alpha}) \cap H(\underline{\alpha}) = \emptyset$. This construction is described in Section 3 of [Bo2].

Now, suppose that for each n we have a sequence $(\alpha^n_i)_i$ of multicurves, and that α^n_i is compatible with α^n_{i+1} for all i and n (i.e. $\alpha^n_i \cup \alpha^n_{i+1}$ is a multicurve). We get sequence, $\underline{\alpha}_i$ of sequences indexed by n, enabling us to construct a sequence of subsurfaces $F_i = F(\underline{\alpha})$, $G_i = G(\underline{\alpha})$ and $H_i = H(\underline{\alpha})$. These have the intersection properties (F1)–(F9) as laid out out in Section 3 of [Bo2].

Suppose now that M is a coarse hyperbolic manifold. Let Δ_ν be the thin part. Suppose that γ is a multicurve and that $q(\gamma) \cap \Delta_\nu = \emptyset$. There is an associated realisation $j : \gamma \longrightarrow (\Sigma, \sigma)$ where σ is a thick hyperbolic metric on Σ. Since the properties of an associated realisation are invariant under uniform bilipschitz homeomorphism of Σ, we can always take σ to be fixed modulo the action of the mapping class group of Σ. Put another way, we can fix once and for all the structure, $S = (\Sigma, \sigma)$. Then given γ, we can find a geodesic multicurve, $\alpha \subseteq S$, a mapping class ω sending α to γ, and a new realisation, $j' : \gamma \longrightarrow \alpha$ sending γ to $\alpha = (\omega^{-1}\gamma)_\sigma$.

In order to maintain consistent notation with [Bo2], we introduce the following maps. We choose a homotopy class $\chi : M \longrightarrow \Sigma$ which induces the identity on $\pi_1(M) \equiv \pi_1(\Sigma) \equiv \Gamma$. If γ, ω are as above, we write $\phi = \chi^{-1} \circ \omega$ for the corresponding homotopy class $S \longrightarrow M$. (Here χ^{-1} is any right inverse homotopy class to χ.) Moreover, we can define $\phi|\alpha$ to be $q \circ (j')^{-1} : \alpha \longrightarrow M$. From the property of associated realisations, this is uniformly bilipschitz with respect to the induced path metrics.

This now provides the essential features of a pleating surface, as used in Section 7 of [Bo2] (see the "modified definition" there). In summary, suppose that γ is a multicurve, and that γ_M is its realisation in M, and suppose that $\gamma_M \cap \Delta_\nu = \emptyset$ for some $\nu \in \mathbf{N}$. Then there is a geodesic multicurve, $\alpha \subseteq S$, and a map $\phi : \alpha \longrightarrow M$, which is uniformly bilipschitz onto γ_M in the induced path metric. This extends to a map $\phi : S \longrightarrow M$, defined up to homotopy, so that $\omega = \chi \circ \phi : S \longrightarrow \Sigma$ is a homotopy equivalence. The map, ϕ, satisfies the uniform injectivity theorem as laid out in Section 6, where $j = (\phi|\alpha)^{-1} \circ q$ is the associated realisation sending γ to α, and $q : \gamma \longrightarrow \gamma_M$ is the realisation of γ in M.

The construction of Section 7 of [Bo2] now goes through with little change, as we now describe.

Suppose we have a sequence, M^n of spaces with thin parts Δ^n_ν and maps $\chi^n : M^n \longrightarrow \Sigma$, all defined with respect to the same set of parameters and a fixed value of ν. Suppose that γ^n is a sequence of realisations, $\gamma^n_{M^n}$ in M^n and with $\gamma^n_{M_n} \cap \Delta^n_3 = \emptyset$. We let $\alpha^n, \phi^n : \Sigma \longrightarrow S$ be the pleating surface maps described above, and set $\omega^n = \chi^n \circ \phi^n$. Passing to a subsequence we obtain subsurfaces, $F(\underline{\alpha})$, $G(\underline{\alpha})$, $H(\underline{\alpha})$ in S, and set $F^n = \omega^n F(\underline{\alpha})$, $G^n = \omega^n G(\underline{\alpha})$ $H^n = \omega^n H(\underline{\alpha})$ in Σ.

We want to recognise F^n, G^n, H^n directly in terms of the realisations $\gamma^n_{M^n}$ in M^n, and without reference to the choice of pleating surfaces. (This is where we made essential use of the uniform injectivity theorem in [Bo2].) To this end, we employ Lemma 6.3. Given $\eta > 0$, let $I_\eta = [\eta(U(\eta)), \eta]$. Now by the process described in [Bo2], we can find some $\eta > 0$ (depending on $\underline{\alpha}$) so that for all sufficiently large n, the realisations α^n are I_η-stable in S (in the sense described in Section 6 of [Bo2]). Moreover, $\Phi(\alpha^n, I_\eta) = G(\underline{\alpha})$ (defined up homotopy). Setting $U = U(\eta)$, Lemma 6.3 tells us that $\Phi_M(\gamma, U) = \omega^n \Phi(\alpha^n, I_\eta)$. (Since $\phi^n \Phi(\alpha^n, I_\eta) = \Phi(\gamma_{\phi^n\sigma}, I_\eta)$, where $\gamma_{\phi^n\sigma}$ denotes the geodesic realisation of γ in the induced structure, $\phi^n\sigma$, on Σ. That is we have just introduced a non-trivial mapping class on both sides.) In other words, $G^n = \Phi_M(\gamma^n, U)$. (In [Bo2], this was described in terms of the projectivised tangent bundle. Instead of tangent vectors being close, we now have long segments in $\gamma^n_{M^n}$ running parallel. This is merely a reinterpretation of the projectivised tangent bundle to a more general setting.)

To define F^n, we proceed similarly. This time, we choose subsets $a^n \subseteq \alpha^n$, so that $a^n \longrightarrow \mu(\lambda)$, and let $c^n = \phi^n \subseteq \gamma^n_{M^n}$, We again get some η, which we can assume to be the same as the original so that c^n is also I_η-stable. Moreover, $\Phi(a^n, I_\eta) = F(\underline{\alpha})$ and with $U = U(\eta)$, we get $F^n = \Phi_M(c^n, U)$, similarly as above.

We still need to check that this is independent of the choice of c^n. But this follows as in Lemma 7.5 of [Bo2], with Proposition 6.5 providing us with the appropriate uniformly bicontinuous homomorphism between different realisations.

We can now apply this construction to a sequence of multicurves. As a result, we can now formulate the analogue of Lemma 7.6 of [Bo2]. Here "(F1)–(F9)" refer to the properties of a sequence of subsurfaces laid out in Section 3 of that paper.

Lemma 7.2 : *Suppose we have a sequence, M^n, of coarse hyperbolic manifolds with thin parts $\Delta^n_\nu \subseteq M^n$, and with $\chi^n : M^n \longrightarrow \Sigma$, all defined with respect to fixed parameters and fixed ν. Suppose, for each n, we have a sequence $(\gamma^n_i)^n_{i=0}$ of geodesic multicurves in M^n with γ^n_i and γ^n_{i+1} compatible for all i and n. Suppose that for all i and n, the realisations $(\gamma^n_i)_{M^n}$ satisfy $(\gamma^n_i)_{M^n} = \emptyset$. Then, for an infinite subsequence of n, we can construct subsurfaces F^n_i, G^n_i, H^n_i of Σ satisfying (F1)–(F9).* $\diamondsuit$

We now move on to Section 8 of [Bo2]. The key lemma is Lemma 8.1 of that paper, which can be rephrased as follows. Let M be a coarse hyperbolic manifold.

Lemma 7.3 : *Given $k \geq 0$ and $\nu, p \in \mathbf{N}$, there is some $l \geq 0$, depending only on k, p, ν, and the parameters of M such that if $(\gamma_i)_{i=0}^p$ is a tight multigeodesic in M with $(q\gamma_0 \cup q\gamma_p) \cap \Delta_\nu = \emptyset$, and $p \geq 12r + 19$, where $r = \max\{d(\gamma_0, X(M, l_0)), d(\gamma_p, X(M, l_0))\}$, then there is some $i \in \{0, \ldots, p\}$ such that $L(M, \gamma_i) \leq l$.*

Proof : We follow the argument in [Bo2]. Assume that the statement fails. We get a sequence of spaces, M^n, and sequences (γ_i^n), of multigeodesics, whose realisations in M^n are getting longer and longer. The argument of [Bo2] involves finding suitable Margulis constants with respect to which all the relevant realisations are "non-penetrating". Here we achieve the analogous statements using Lemma 7.1. We end up with a thin part, $\Delta_{\nu'}$, possibly smaller than the original Δ_ν. Now Lemma 7.2 (applied with ν') enables us to construct surfaces (F_i^n), (G_i^n), (H_i^n) and arrive at a contradiction exactly as in [Bo2]. $\diamondsuit$

The same reasoning also applies to the "interpolation lemma", namely Lemma 8.2 of [Bo2], where "hyperbolic 3-manifold" is replaced by "coarse hyperbolic manifold".

Theorems 0.1 and 0.3 of this paper are, respectively, reinterpretations of Theorems 1.3 and 1.4 of [Bo2]. What remains of their proofs, as described in Section 8 of [Bo2], now goes through in exactly the same way.

8. Exceptional surfaces.

In this section, we deal with the case where Σ is a one-holed torus or four-holed sphere (abbreviated to 1HT and 4HS respectively). In this case, $\mathcal{G}(\Sigma)$ is interpreted as the modified curve graph with vertex set $V(\mathcal{G}) = X(\Sigma)$, and where $\alpha, \beta \in X(\Sigma)$ are adjacent if they have minimal possible intersection number (1 or 2 respectively). Thus $\mathcal{G}(\Sigma)$ is a Farey graph. Here every geodesic is deemed tight. Suppose that Σ is a 1HT or 4HS, and that $\Gamma = \pi_1(\Sigma)$ acts on a k-hyperbolic space, H. Given $\gamma \in X(\Sigma)$, we write $l(\gamma)$ for its stable length (denoted $l_M^S(\gamma)$ previously). We suppose that $l(\delta) \leq L_0$ for each component δ of $\partial\Sigma$. We write $X(l) = \{\gamma \in X(\Sigma) \mid l(\gamma) \leq l\}$.

We claim that Theorems 0.1, 0.2 and 0.3 hold. Indeed, the assumptions (M1) and (M2) are not needed here, and we will restate (M3) explicitly. Thus, for example, Theorem 0.1 is a consequence of the following analogue of Theorem 9.3 of [Bo2].

Theorem 8.1 : *Given $l \geq 0$, there is some l' depending only on l, L_0 and the hyperbolicity constant k with the following property. Let $\gamma_0, \ldots, \gamma_p$ be a geodesic in $\mathcal{G}(\Sigma)$ with $l(\gamma_0) \leq l$, $l(\gamma_p) \leq l$, then $l(\gamma_i) \leq l'$ for all i.*

Under the same hypotheses we also have the following, from which Theorem 0.2 follows:

Theorem 8.2 : *There is some l_0, depending only on k and L_0 such that for all $l \geq l_0$, the full subgraph on $X(l)$ is non-empty and connected. Moreover, for all $l \geq l_0$, $X(l)$ lies in an r-neighbourhood of $X(l_0)$, where r depends only k, L_0 and h.*

To relate the above to the case of non-exceptional surfaces, we note that any connected subgraph of a Farey graph is 1-quasiconvex.

There is also a similar variation on Theorem 0.3, which is an immediate consequence of the above, the geometry of the Farey graph, and Lemma 8.3 below.

The logic of the proof will follow that of Section 9 of [Bo2]. For this, we need to reinterpret the essential features of the trace identities used there in terms of inequalities involving distances. We recall from that paper, the following terminology.

Three distinct curves, $\alpha, \beta, \gamma \in V(\mathcal{G})$ form a *triangle* if they are pairwise adjacent. Four distinct curves, $\alpha, \beta, \gamma, \delta$ form a *rhombus* if α, β, γ and α, β, δ are both triangles.

The key lemma will be:

Lemma 8.3 :
(1) Suppose $t \geq 0$ and α, β, γ and form a triangle with $l(\alpha) \leq t$ and $l(\beta) \leq t$, then $l(\gamma) \leq t'$, where t' depends only on t, k and L_0.
(2) There exists $h_0 \geq 0$ depending only on k such that $\alpha, \beta, \gamma, \delta$ form a rhombus and $l(\alpha) \geq h_0$, $l(\beta) \geq h_0$, then $\max\{l(\gamma), l(\delta)\} \geq l(\alpha) + l(\beta) - 2L_0 - h_0$.

We have stated this in unified form, whether Σ is a 1HT or 4HS, though the proofs in the two cases are different. (In fact one can give slightly stronger statements in each of the two cases — for example in (2) for a 1HT the term in $2l_0$ can be omitted whereas for a 4HS the assumptions on $l(\alpha)$ and $l(\beta)$ can be omitted.)

To simpify notation in what follows, we shall introduce the following conventions (cf. [Bo1]). Given $x, y \in \mathbf{R}$ write $x \simeq_t y$, $x \preceq_t y$ and $x \ll_t y$ to mean respectively $|x - y| \leq t$, $x \leq y + t$ and $x + t \leq y$. The constant, t, will change during the course of an argument, but at any given stage will be a fixed multiple of the hyperbolicity constant, h, which could in principle, be determined by following through the argument step by step. For this reason we omit the subscript t. We will behave as though the relations $\simeq$ and $\preceq$ were transitive, and $\preceq$ and $\gg$ were mutually exclusive. This entails choosing a sufficiently large initial constant for $\ll$ at the outset.

Note that, in these terms, Lemma 8.3(2) could be interpreted as saying that if $l(\alpha), l(\beta) \gg 0$, then $\max\{l(\gamma), l(\delta)\} \succeq l(\alpha) + l(\beta) - 2l_0$.

To get from Lemma 8.3 to Theorems 8.1 and 8.2 we proceed as in [Bo2]. Given a rhombus $\alpha, \beta, \gamma, \delta$ we put a transverse orientation on the edge $\alpha\beta$ from γ to δ if $l(\gamma) > l(\delta)$. (If $l(\gamma) = l(\delta)$ we orient it arbitrarily.) Note that if $l(\alpha), l(\beta) \gg 0$, this implies that $l(\gamma) \gg l(\alpha) + l(\beta) - 2l_0$. We note:

Lemma 8.3 : *If α, β, γ is a triangle, with $\alpha\beta$ and $\alpha\gamma$ transversely oriented outwards, then $\min\{l(\alpha), l(\beta), l(\gamma)\} \preceq 2l_0$.*

Proof : If not, then $l(\alpha), l(\beta), l(\gamma) \gg 0$, and so $l(\alpha) + l(\beta) \preceq l(\gamma) + 2l_0$ and $l(\alpha) + l(\gamma) \preceq l(\beta) + 2l_0$, so $l(\alpha) \preceq 2l_0$ a contradiction. $\diamondsuit$

This is enough to follow through the logic of [Bo2] with $l(\alpha)$ playing the role of traces (after taking logarithms) and inequality interpreted in the obvious way.

To prove Lemma 8.4, we need some general observations about a group Γ acting on a hyperbolic space H.

Given any $g \in \Gamma$ and $r \geq 0$, write $F(g,r) = \{x \in H \mid d(x,gx) \leq r\}$. We note that if $r \gg l(g)$ then $F(g,r)$ is non-empty and uniformly quasiconvex in H. We choose some suitable fixed multiple, h_1, of the hyperbolicity constant h, and set $F(g) = F(g, l(g)+h_1)$, so that $F(g)$ is always non-empty and uniformly quasiconvex in H. Note that $d(x,gx) \preceq l(g)$ for all $x \in F(g)$. Given $g,h \in \Gamma$, set $D(g,h) = d(F(g),F(h))$. If $D(g,h) \gg 0$, set $\sigma(g,h)$ to be a shortest geodesic from $F(g)$ to $F(h)$ (or shortest up to bounded distance). This is uniquely defined up to bounded distance. These statements are all fairly elementary applications of hyperbolicity.

If $l(g) > 0$, then there is a bi-infinite geodesic (at least up to an arbitarily small additive constant), $\pi(g)$ with $g\pi(g)$ a uniformly bounded distance from $\pi(g)$. This is well-defined up to uniformly bounded Hausdorff distance. If $l(g) \gg 0$ then it is a bounded Hausdorff distance from $F(g)$. We refer to it as an *axis* of g. If $l(g) \gg 0$ this it is assumed to be oriented in the direction of displacement. Note that $l(g^{-1})$ is the same as $l(g)$ (but oriented in the opposite sense).

If $Q \subseteq H$ and π is a bi-infinite geodesic, we write $P_\pi(Q) \subseteq \pi$ for the projection of Q to π, i.e. the set of points in π that are nearest to some point in Q. If Q is quasiconvex, then we can assume (up to bounded distance) that $P_\pi(Q)$ is a connected segment of π.

We need one final observation. Suppose that σ is a finite directed segment with $\mathrm{length}(\sigma) \gg 0$. Let x,y be it initial and final vertices. We say that $g \in \Gamma$ *displaces* σ a positive distance $t \geq 0$ if there is a geodesic τ from x to gy so that gx and y are a bounded distance from τ and $d(x,gx) \simeq t$. In this case $l(g) \simeq t$, and $d(y,gy) \simeq t$. If $t \gg 0$, then the segments σ, $g\sigma$ and τ all lie in a bounded neighbourhood of $\pi(g)$. (We can similary define negative displacement, which is equivalent to reversing the orientaion on σ.)

We can use these observations for the following:

Lemma 8.5 : *Let $g,h \in \Gamma$. Then $D(g,h) \gg 0$ if and only if $l(hg) \gg l(g)+l(h)$. In this case, we have:*

(1) $l(hg) \simeq l(g)+l(h)+2D(g,h)$.

(2) $\sigma(g,h)$ lies in a bounded neighbourhood of $\pi(hg)$ and is oriented in the same sense.

(3) $\operatorname{diam} P_{\pi(hg)}(F(g)) \simeq l(g)$ and $\operatorname{diam} P_{\pi(hg)}(F(h)) \simeq l(h)$.

Proof : (Sketch) If $D(g,h) \simeq 0$, then any point x close to both $F(g)$ and $F(h)$ satisfies $d(x,gx) \simeq l(g)$ and $d(x,hx) = d(x,h^{-1}x) \simeq l(h)$. Thus, $d(x,hgx) = d(h^{-1}x,gx) \preceq l(g)+l(h)$, and so $l(hg) \preceq l(g)+l(h)$. In other words, $l(hg) \gg l(g)+l(h)$ implies $D(g,h) \gg 0$.

Suppose now that $D(g,h) \gg 0$. Let $\sigma = \sigma(h,g)$ so that $-\sigma = \sigma(g,h)$, and $\mathrm{length}(\sigma) = D(g,h) \gg 0$. Consider the three geodesic segments, $g^{-1}\sigma, -\sigma, h\sigma$, in turn connecting $F(g^{-1}hg)$ to $F(g)$ to $F(h)$ to $F(hgh^{-1})$. These are all uniformly quasiconvex. From the definition of $F(h)$ we see that σ and $h\sigma$ must diverge after a bounded distance. It follows that $d(F(h),F(hgh^{-1})) \succeq 2\,\mathrm{length}(\sigma) \gg 0$. Similarly, σ and $g^{-1}\sigma$ diverge after a bounded distance. It now follows that $g^{-1}\sigma$, $-\sigma$, $h\sigma$ all lie close to an oriented geodesic. Now hg displaces the long segment $g^{-1}\sigma$ a long distance to $h\sigma$. Thus $l(hg) \gg 0$, and $g^{-1}\sigma$, $-\sigma$,

$h\sigma$ all lie a bounded distance from $\pi(hg)$. Moreover, $d(g^{-1}\sigma,\sigma) \simeq l(g)$ and $d(\sigma,h\sigma) \simeq l(h)$. The pojections of $F(g)$ and $F(h)$ to $\pi(hg)$, up to bounded distance, cover the gaps between the projections of $g^{-1}\sigma$, $-\sigma$ and $h\sigma$ to $\pi(hg)$, and $\operatorname{diam} P_{\pi(hg)}(F(g)) \simeq l(g)$ and $\operatorname{diam} P_{\pi(hg)}(F(h)) \simeq l(h)$. We see that the translation distance of hg is approximately $\operatorname{length}(g^{-1}\sigma) + l(g) + \operatorname{length}(\sigma) + l(h) \simeq l(g) + l(h) + 2D(g,h)$. Thus $l(hg) \simeq l(g) + l(h) + 2D(g,h)$. $\diamondsuit$

Suppose that $g,h \in \Gamma$ with $l(g) \gg 0$, $l(h) \gg 0$ and $D(g,h) \simeq 0$. We write $L(g,h) \in \mathbf{R}$ for the distance over which $\pi(g)$ and $\pi(h)$ remain almost parallel, taking account of orientations. Thus, $L(g,h) = L(h,g) = -L(g^{-1},h)$. This well defined up to an additive constant.

Lemma 8.6 : *Suppose that $g,h \in \Gamma$ with $l(g), l(h) \gg 0$.*

(1) If $D(g,h) \gg 0$ then $l(hg) \simeq l(g) + l(h) + 2D(g,h)$.

(2) If $D(g,h) \simeq 0$ then:

(a) if $L(g,h) \geq 0$ then $l(hg) \simeq l(g) + l(h)$

(b) if $L(g,h) \leq 0$, then $l(hg) \simeq \max\{|l(g) - l(h)|,\, l(g) + l(h) + 2L(g,h)\}$.

In particular, we see that either $l(hg) \succeq l(g) + l(h)$, or else $D(g,h) \simeq 0$ and $L(g,h) \ll 0$.

Part (1) of Lemma 8.6 follows from Lemma 8.5.

For part (2) one can first consider the case of a tree. In this case, the equalities and inequalites hold exactly, as can be seen by consider in turn the various combinatorial possibilites for the axes $\pi(g)$, $\pi(h)$ and $h\pi(g) = \pi(hgh^{-1})$. If σ is a directed segment of $\pi(g)$ meeting $\pi(h)$ at its forward endpoints, one considers how this is displaced, in turn, by g and then h. The general case of a hyperbolic space follows by noting that the union of these axes lies a uniformly bounded distance from a tree. In this case, we condider the displacement of a sufficiently long segment σ of $\pi(g)$ with forward endpoint a bounded distance from $\pi(h)$ and whose projection to $\pi(h)$ has bounded diameter. Again, we see that σ is displaced a large positive distance by hg. The displacement agrees with the case of a tree up to an additive constant.

We now split the proof of Lemma 8.3 into two cases.

(A) Σ is a IHT.

Part(1):

Let α, β, γ be a triangle. There are elements $a, b \in \Gamma$ with $aba^{-1}b^{-1}$ peripheral, so that α, β and γ are represented respectively by a, b and ab. We are assuming that $l(aba^{-1}b^{-1}) \leq L_0$. Suppose that $l(a), l(b) \leq t$, and that $l(ab) \gg 2t$. By Lemma 8.5 we have $D(a,b) \gg 0$. Moreover $\sigma(b,a)$ lies close to $\pi(ab)$ and is oriented in the same direction. Also, since $D(b,a) = D(a,b) \gg 0$, we have $l(ba) \succeq l(a) + l(b)$, and $\sigma(a,b) = -\sigma(b,a)$ is a bounded distance from $\pi(ba)$. It follows that $\pi(ab)$ and $\pi(ba)$ lie parallel and oriented in opposite directions over a distance at least $\operatorname{length}(\sigma(a,b)) = D(a,b)$. In other words $L(ab,ba) \simeq -D(a,b)$, and so $L(ab,(ba)^{-1}) = -L(ab,ba) \succeq D(a,b) \gg 0$. Thus, by Lemma 8.5, we have $l(aba^{-1}b^{-1}) = l(ab(ba)^{-1}) \succeq l(ab) + l(ba) \geq l(ab)$. It follows that $l(ab) \preceq L_0$.

In summary, we have shown that if $l(a), l(b) \leq t$ and $l(ab) \gg 2t$, then $l(ab) \preceq L_0$. In other words, if $l(a), l(b) \leq t$, then $l(ab) \preceq \max\{2t, L_0\}$.

Part (2):

Suppose that $\alpha, \beta, \gamma, \delta$ form a rhombus, They correspond to a, b, ab, ab^{-1}, where $a, b \in \Gamma$ (and $aba^{-1}b^{-1}$ is peripheral). We suppose that $l(a), l(b) \gg 0$. We claim that $\max\{l(ab), l(ab^{-1})\} \succeq l(a) + l(b)$. For if not, $l(ab) \ll l(a) + l(b)$ and so by Lemma 8.6, $D(a, b) \simeq 0$ and $L(a, b) \ll 0$. Also $l(ab^{-1}) \ll l(a) + l(b)$ and so $L(a, b) \ll 0$. Also $l(ab^{-1}) \ll l(a) + l(b)$, so $L(a, b^{-1}) \ll 0$. But $L(a, b^{-1}) = -L(a, b)$ giving a contradicion.

We have shown that there is a constant $k_0 \geq 0$ such that if $l(a), l(b) \geq h_0$, then $\max\{l(ab), l(ab^{-1})\} \geq l(a) + l(b) - h_0$.

(B) Σ is a 4HS.

In this case, if $\alpha, \beta, \gamma, \delta$ forms a rhombus, then we can find elements $p, q, r, s \in \Gamma$, each representing a different boundary component of Σ, with $pqrs = 1$, such that $\alpha, \beta, \gamma, \delta$ are represented respectively by pq, qr, pr, qs. Note that α, β are also represented by rs and sp respectively. We are assuming that $l(p), l(q), l(r), l(s) \leq L_0$.

Part (1):

Suppose that α, β, γ forms a triangle, and that $l(\alpha), l(\beta) \leq t$, and that $l(\gamma) \gg 2t + 4L_0$. Let δ be the curve opposite α across the edge $\beta\gamma$, so that $\gamma, \beta, \alpha, \delta$ is a rhombus. We can find $p, q, r, s \in \Gamma$ with $pqrs = 1$ so that $\gamma, \beta, \alpha, \delta$ are represented respectively by pq, qr, pr, qs. In particular, $l(p), l(q), l(r), l(pqr) \leq L_0$. Since $l(pr), l(qr) \leq t$ and $l(pq) \gg 2t + 2L_0$, it follows that $D(p, r) \preceq t/2$ and $D(q, r) \preceq t/2$. Let $\pi = \pi(pq)$, and set $P(q) = P_\pi(F(p))$, $P(q) = P_\pi(F(q))$ and $P(r) = P_\pi(F(r))$. (Note $\sigma(p, q)$ is a bounded distance from $\pi(p, q)$.) Now $d(P(p), P(r)) \simeq D(p, r) \preceq t/2$, $d(P(q), P(r)) \simeq D(q, r) \preceq t/2$, and $d(P(p), P(q)) \gg t + 2L_0$. It follows that $\operatorname{diam} P(r) \gg 2L_0$. Let $\tau \subseteq P(r)$ be a segment with $\operatorname{length}(\tau) \gg 2L_0$. This must lie a bounded distance from $P(r)$ and so $d(x, rx) \simeq l(r) \leq L_0$ for all $x \in \tau$. Thus there is a subsegment $\sigma \subseteq \tau$, with $\operatorname{length}(\sigma) \gg 0$ such that $r\sigma$ is a within a bounded neighbourhood of π, oriented in the same sense. Since pq displaces π a distance much more than $t + 2L_0$, we see that pqr displaces σ a distance much more than $t + L_0$. It follows that $l(pqr) \gg t + L_0 \geq L_0$. But $l(pqr) \leq L_0$ giving a contradiction.

In other words, we have shown that if $l(\alpha), l(\beta) \leq t$ then $l(\gamma) \preceq 2t + 4L_0$.

Part (2):

Let $\alpha, \beta, \gamma, \delta$ form a rhombus. Let $p, q, r, s \in \Gamma$ be such that $pqrs = 1$ and $l(p), l(q), l(r), l(s) \leq L_0$ and $\alpha, \beta, \gamma, \delta$ represented respectively by pq, qr, pr, qs.

We can assume that either $l(\alpha) \gg L_0$ or $l(\beta) \gg L_0$, otherwise there is nothing to prove. So we can suppose, without loss of generality, that $l(\alpha) = l(pq) \gg L_0$.

Suppose, for contradiction, that $\max\{l(pq), l(qs)\} \gg l(pq) + l(qr) - 2L_0$. In particular, we can assume, without loss of generality, that $l(pq) \gg 2L_0$. Thus, $D(p, q) \gg 2L_0$. Let $\pi = \pi(pq) = -\pi(rs)$. Set $P(p) = P_\pi(F(p))$, $P(q) = P_\pi(F(q))$, $P(r) = P_\pi(F(r))$ and $P(s) = P_\pi(F(s))$. By Lemma 8.5, we have $\operatorname{diam} P(p), \operatorname{diam} P(q), \operatorname{diam} P(r), \operatorname{diam} P(s) \preceq L_0$. Also $\sigma(p, q)$ and $\sigma(s, r)$ lie a bounded distance from π, are oriented in the same sense, and each has length approximately $D(p, q) \simeq D(s, r) \gg 2L_0$. We see that the segments of π between $P(p)$ and $P(q)$ and between $P(r)$ and $P(s)$ each have approximately equal length much more than $2L_0$ and oriented in the same sense. It follows that, along π, the most extreme among $P(p), P(q), P(r), P(s)$ are either $P(p)$ and $P(r)$ or else $P(q)$ and $P(s)$.

In the former case, we have $D(p,r) \simeq d(F(p), F(r)) \succeq d(F(p), F(q)) + d(F(q), F(r)) \simeq D(p,q) + D(q,r) \simeq l(pq) + l(qr) \gg 0$. Thus $l(pr) \succeq l(pq) + l(pr)$. In other words, $l(\gamma) \succeq l(\alpha) + l(\beta)$ as required.

Similarly, in the latter case, we have $D(q,s) \simeq d(P(q), P(s)) \succeq d(P(p), P(r)) + d(P(r), P(s))$, and so $l(qs) \succeq l(qr) + l(rs) = l(qr) + l(pq)$. This time, we get $l(\delta) \succeq l(\alpha) + l(\beta)$. Either way, $\max\{l(\gamma), l(\delta)\} \succeq l(\alpha) + l(\beta)$.

Thus, in all cases, we have $\max\{l(\gamma), l(\delta)\} \succeq l(\alpha) + l(\beta) - 2L_0$.

This proves Lemma 8.3.

As observed, Theorems 8.1 and 8.2 now follow, applying the logic of [Bo2].

9. Hierarchies.

For applications (in [Bo4]) we need to bound the lengths of all curves in a "hierarchy". (This is the form of the a-priori bounds theorem for hyperbolic 3-manifolds as given in [Mi].) This is a simple consequence of what we have already done. We state it here explicitly for reference.

Suppose $Q \subseteq X(\Sigma)$. We construct a larger subset $J(Q) \subseteq X(\Sigma)$ as follows. If we can find a subsurface $\Phi \in \mathcal{F}$ of Σ with $X(\partial_\Sigma \Phi) \subseteq Q$ and curves $\alpha, \beta, \gamma \in X(\Sigma)$ such that $\alpha, \beta \in Q \cap X(\Phi)$ (identifying $X(\Phi) \subseteq X(\Sigma)$) and such that γ lies on a tight geodesic from α to β in $\mathcal{G}(\Phi)$, then we include γ in $J(Q)$. In other words, $J_n(Q)$ is the set of all γ arising in this way. Note that $Q \subseteq J(Q)$. (If $\gamma \in Q$ then set $\Phi = \Sigma$ and $\alpha = \beta = \gamma$.) For any $n \in \mathbf{N}$, we define $J_n(Q)$ inductively by $J_0(Q) = Q$ and $J_{n+1}(Q) = J(J_n(Q))$.

(This description differs slightly from that in Section 8 of [Bo2] in that we are allowing Φ to be a 1HT or 4HS in the above, in which case $\mathcal{G}(\Sigma)$ is taken to be the Farey graph.)

For the purposes of stating the result, we can regard a "hierarchy" associated to a subset $Q \subseteq X(\Sigma)$ to be some canonically defined subset, $\mathcal{H}(Q)$, such that $\mathcal{H}(Q) \subseteq J_n(Q)$ where n depends only on the topological type of Σ.

Theorem 9.1 : *Suppose that M is a coarse hyperbolic manifold. Given any $k \geq 0$, there is some $k' \geq 0$, depending only k and the parameters of M such that if $Q \subseteq X(M,k)$, then $\mathcal{H}(Q) \subseteq X(M,k')$.*

(Recall that $X(M,r) = \{\gamma \in X(\Sigma) \mid l^S_M(\gamma) \leq r\}$.)

The proof is just an inductive application of Theorem 0.1. Note that if $\Phi \in \mathcal{F}$ with $X(\partial_\Sigma \Phi) \subseteq X(M,r)$ for some r, then we get a coarse hyperbolic manifold $M_\Phi = \tilde{M}/\pi_1(\Phi)$. The parameters of M_Φ depend only on those of M. In fact only the parameter associated to (M3) has changed, and that is contolled by the inductive hypotheses.

References.

[Bo1] B.H.Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces* : in "Group theory from a geometrical viewpoint" (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 64–167.

[Bo2] B.H.Bowditch, *Length bounds on curves associated to tight geodesics* : Geom. Funct. Anal. **17** (2007) 1001–1042.

[Bo3] B.H.Bowditch, *Train tracks and surface groups acting on hyperbolic spaces* : preprint, Warwick (2010).

[Bo4] B.H.Bowditch, *Coarse hyperbolic models for 3-manifolds* : preprint, Warwick (2010).

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

[GhH] E.Ghys, P.de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov* : Progress in Mathematics 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 (1988) 75–263.

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

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

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

[NS] H.Namazi, J.Souto, *Revisiting Thurston's uniform injectivity theorem* : in preparation.

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