Length bounds on curves arising from tight geodesics

Brian H. Bowditch
School of Mathematics, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
http://www.maths.soton.ac.uk/staff/Bowditch
[First draft: November 2003. Revised: March 2007]

0. Introduction.

The main aim of this paper is to describe certain bounds on lengths of curves in hyperbolic 3-manifolds. The curves arise from tight geodesics in the curve complex associated to a surface. The results themselves have considerable overlap with the bounds described in [Mi4], where they are termed "a-priori bounds". They play a central role in the proof of Thurston's Ending Lamination Conjecture [Mi4,BrCM1.BrCM2], and also have applications to the geometry of the curve complex as we describe in [Bow4]. Our methods are different, and we give proofs of these bounds which bypass much of the sophisticated machinery developed in [MaM2,Mi4] and elsewhere. Another proof of the Ending Lamination Conjecture, based on these bounds, is given in [Bow5,Bow6]. As in [Mi4], we confine our attention here to the indecomposable case. However, the essential points can be adapted to decomposable case, as described in [Bow6].

Let Σ be a compact surface, possibly with boundary, $\partial\Sigma$. Let $X = X(\Sigma)$ be the set of homotopy classes of non-trivial non-peripheral simple closed curves in Σ (which we usually refer to simply as "curves"). By the *curve graph*, $\mathcal{G} = \mathcal{G}(\Sigma)$, of Σ we mean the 1-skeleton of the curve complex defined by Harvey [Har]. Thus the vertex set of $\mathcal{G}$ is identified with X, and two curves are deemed to be adjacent in $\mathcal{G}$ if they can be realised disjointly in Σ. The curve complex has been much used in the study of the mapping class group, the geometry of Teichmüller space, and hyperbolic 3-manifolds.

We define the *complexity*, $\kappa(\Sigma)$, of Σ as $\kappa(\Sigma) = 3g + p - 3$ where g is the genus of Σ and p is the number of boundary components. If $\kappa(\Sigma) > 1$, then $\mathcal{G}$ is connected and we write d for the combinatorial path metric on $\mathcal{G}$. It was shown in [MaM1] that $\mathcal{G}$ is hyperbolic in the sense of Gromov [Gr,GhH]. Other proofs are given in [Bow2] and [Ham].

A complicating factor in applying the machinery of hyperbolic groups [G] to the curve complex is that (unlike the Cayley graph of a finitely generated group) the curve complex is not locally finite. Indeed, there may be infinitely many geodesics connecting two given vertices. One can, however, formulate weaker finiteness conditions. In [MaM2], the authors describe a certain class of "tight" geodesics. One of their key results is that the set of tight geodesics connecting any two vertices of $\mathcal{G}$ is finite. Their proof makes much use of their sophisticated theory of "hierarchies".

In this paper, we offer a direct proof of finiteness. Our statement is non-constructive, but can be refined in a number ways. For example in [Sh1,Sh2] purely combinatorial arguments are given, which provide explicit computable bounds. In a different direction, one of the consequences of later results in this paper is that one can give certain uniform (though still non-constructive) bounds [Bow4]. This in turn has consequences for the acylindricity of the action of the mapping class group on $\mathcal{G}$, as well as the uniform rationality of stable lengths. Other applications of this uniformity show that the curve complex has finite asymptotic dimension [BelF] and has Yu's property A [Ki]. The statements of [Bow4] are purely combinatorial, but their proofs make use of the theory of hyperbolic 3-manifolds, in particular, the a-priori bounds, which we now summarise.

Let M be a complete hyperbolic 3-manifold admitting a homotopy equivalence to Σ. We suppose that each boundary component of Σ corresponds to a parabolic cusp in M (i.e. it is "type preserving"). For some applications (e.g. those of [Bow4]), one could assume in addition that there are no accidental parabolics, i.e. each cusp of M corresponds to a boundary curve of Σ (i.e. it is "strictly type preserving"). In the latter case, each curve in Σ can be realised uniquely as a closed geodesic in M, whose length we denote by $L(M, \alpha)$. It will be shown that if $\alpha, \beta, \gamma \in X(\Sigma)$ and γ lies in a tight geodesic from α to β, then $L(M, \gamma)$ can be bounded above as a function of $\kappa(\Sigma)$, $L(M, \alpha)$ and $L(M, \beta)$. (For the applications we have in mind, $L(M, \alpha)$ and $L(M, \beta)$ can be assumed to be "small", so the relevant bounds depend only on $\kappa(\Sigma)$.) If we allow accidental parabolics, we deem the length of the corresponding simple closed curve to be 0. There are a number of variations on this result, which we describe in Section 1.

In Section 9, we shall give an account of the "exceptional" cases where Σ is either a one-holed torus or a four-holed sphere (so that $\kappa(\Sigma) = 0$). In this case one needs to modify the definition of the curve graph, and the methods of proof are quite different (cf. [Mi1,Bow1,Z]). This will be logically independent of the remainder of the paper, except for its application to "hierarchies" in Section 8 (see Corollary 8.3).

We should make a number of comments on these results. Firstly, the notion of "tight" geodesic we use here is slightly weaker than that of [Mi4]. Thus, in principle, the results are stronger, though it is unclear whether the additional information has practical significance.

As in [MaM2], (in the non-exceptional case) we make essential use of geometric limit arguments. As a consequence, the bounds obtained are not a-priori effectively computable. It would be interesting to find a means of bypassing these limit arguments. (The bounds obtained for the exceptional surfaces are computable.)

One of the main applications of the curve complex has been the proof of Minsky, Brock and Canary of Thurston's Ending Lamination Conjecture. The indecomposable case can be reduced to considering product manifolds of the type we have described. Such a 3-manifold, M, has two "end invariants" associated to it. These describe the asymptotic geometry of the respective ends. For such a product manifold, the Ending Lamination Conjecture says that M is determined up to isometry by these two invariants. This was proven in [Mi4,BrCM1]. Another proof was later given in [Bow5]. Both arguments involve a detailed analysis of the geometry of M. While many of the ideas are similar, the logic of the two approaches is different. A key fact is that the end invariants, via the a-priori bounds theorem, determine a canonical set of curves in Σ, such that the lengths of the corresponding geodesics in M are uniformly bounded. In [Mi4], this statement is embedded in the general analysis of the geometry of M, and makes use of results from [MaM2]. In [Bow5], however, it is taken as the starting point. As a result, some of the more technical issues of [BrCM1] can be avoided. We remark that the decomposable case of the Ending Lamination Conjecture is treated in the respective sequels [BrCM2] and [Bow6].

Our account also includes a brief discussion of hierarchies. We use a much simplified notion, which, as far as statements about a-priori bounds are concerned, include the case described by [MaM2,Mi4]. Essentially, a hierarchy can be thought of as a set of curves obtained by an inductive procedure involving tight geodesics. The a-priori bounds result extends to hierarchies via a fairly simple inductive argument. It is really these bounds that are needed in [Bow5].

Section 9 of this paper was written at the Max-Planck-Institut in Bonn, and various revisions to the paper were carried out at the Tokyo Institute of Technology. I am grateful to both institutions for their support and hospitality. I am indebted to Ser Peow Tan for suggesting a means of simplifying the argument of Section 9. I also thank the referee for many helpful comments, as well some interesting observations regarding the Uniform Injectivity Theorem.

1. Statement of results.

We give a statement of the main results. At the end of the section, we explain briefly how these relate to to those of Minsky.

We begin be describing the notion of a "tight geodesic" in $\mathcal{G}(\Sigma)$ when $\kappa(\Sigma) > 0$. Before giving a formal definition, we explain the idea as follows. Suppose, for example, $\alpha, \beta \in X(\Sigma)$ with $d(\alpha, \beta) = 2$. Then $\alpha \cup \beta$ determines a subsurface, Φ, of Σ which it fills. (Glue in all complementary discs and peripheral annuli and then take a small regular neighbourhood.) Any curve, γ, in Σ disjoint from Φ will satisfy $d(\alpha, \gamma) = d(\beta, \gamma) = 1$, i.e. α, γ, β is a geodesic in $\mathcal{G}$. In general, they may be infinitely many such γ, but we can reduce to finitely many by considering only those that are homotopic to a boundary component of Φ in Σ. This can be expressed in terms of $\mathcal{G}$ by saying that any fourth curve that is adjacent in $\mathcal{G}$ to both α and β is adjacent to, or equal to, γ. In principle, one could try to define a "tight geodesic", $\gamma_0, \gamma_1, \ldots, \gamma_p$ in $\mathcal{G}$, by making this hypothesis for each segment $\gamma_{i-1}, \gamma_i, \gamma_{i+1}$ as i ranges from 1 to $p-1$. The problem is that it is not clear that such a geodesic always exists between any two vertices that are distance greater than 2 apart. (Indeed, as far as I am aware, this question remains open.) For this reason, as in [MaM2], we reformulate everything in terms of "multigeodesics", as follows.

We say that two curves in $X = V(\mathcal{G})$ are *disjoint* if they can be realised disjointly, i.e. are adjacent in $\mathcal{G}$. Otherwise, we say that they *cross*. Any non-empty pairwise disjoint set of curves can be simultaneously realised so as to be disjoint. We refer to such a realisation, $\alpha \subseteq \Sigma$, defined up to homotopy, as a *multicurve*. We write $X(\alpha) \subseteq X$ for the set of components of α. We write MX for the set of multicurves in Σ.

We say that $\alpha, \beta \in MX$ are *(exactly) distance r apart* if $d(\gamma, \delta) = r$ for all $\gamma \in X(\alpha)$ and $\delta \in X(\beta)$, where d is the combinatorial metric on $\mathcal{G}(\Sigma)$. A *multigeodesic* consists of a sequence $(\gamma_i)_i$ of multicurves indexed by a set of consecutive integers such that for all i, j, γ_i and γ_j are exactly distance $|i-j|$ apart. We say that $(\gamma_i)_i$ is *tight* at index i if each curve crossing some curve of $X(\gamma_i)$ also crosses some curve of $X(\gamma_{i-1}) \cup X(\gamma_{i+1})$. (Here we differ slightly from [MaM2], where it was assumed that $X(\gamma_i)$ be maximal in this respect, i.e. to be the set of all relative boundary components of the subsurface of Σ filled by $\alpha \cup \beta$. We have no reason to assume that here.) We say that $(\gamma_i)_i$ is *tight* if it is tight at all indices other than the first and last. A geodesic, $(\alpha_i)_i$, in X is *tight* if there is a tight multigeodesic, $(\gamma_i)_i$, such that $\alpha_i \in X(\gamma_i)$ for all i. (Thus a tight geodesic need not be tight as a multigeodesic.)

It is shown in [MaM2] that any two vertices of $\mathcal{G}$ can be connected by a tight (multi)geodesic. They also show:

Theorem 1.1 : *If $\alpha, \beta \in X$, there are only finitely many tight geodesics from α to β.*

Of course, this is equivalent to the same statement for multigeodesics. In Section 3, we give another proof of this result. A constructive proof of this result is given in [Sh]. We also have the following variation on Theorem 1.1.

Theorem 1.2 : *Given $\alpha, \beta \in X(\Sigma)$ and $r \in \mathbf{N}$ there is some finite subset $A \subseteq X(\Sigma)$ such that if $(\gamma_i)_{i=1}^p$ is a tight geodesic in $\mathcal{G}$ with $d(\alpha, \gamma_0) \leq r$ and $d(\beta, \gamma_p) \leq r$, then $\gamma_i \in A$ for all i with $12r \leq i \leq p - 12r$.*

The arguments we give were inspired by ideas in [BesF], which the authors in turn say are inspired by the argument of Luo for showing that the curve complex has infinite diameter (see [MaM1]). I understand from the referee that Luo has attributed some of his ideas to the work of Kobayashi [Ko]. The main purpose of this paper will be to adapt these ideas to hyperbolic 3-manifolds.

Let M be a complete hyperbolic 3-manifold admitting a homotopy equivalence, $\chi : M \longrightarrow \Sigma$. We assume that, under this equivalence, each boundary component of M corresponds to a parabolic cusp. The remaining cusps of M are called *accidental*. Any accidental cusp is homotopic to an element of $X(\Sigma)$, and the set of such cusps is finite. (These statements follow, for example, from purely topological considerations such as the relative Scott Core Theorem [Mc].) We write $X_A(\Sigma) \subseteq X(\Sigma)$ for the set of curves correspoinding to accidental parabolics. Recall that $L(M, \alpha)$ denotes the length of the closed geodesic in M corresponding to a curve $\alpha \in X$. If $\alpha \in X_A(\Sigma)$, we set $L(M, \alpha) = 0$. Given $k \geq 0$, we write $X(M, k) = \{\alpha \in X \mid L(M, \alpha) \leq k\}$. It turns out [Mi3] that for all k sufficiently large (depending only on $\kappa(\Sigma)$) the set $X(M, k)$ is quasiconvex in $\mathcal{G}$, where the constant of quasiconvexity depends only on $\kappa(\Sigma)$ and k. (This is discussed further in Section 4.)

Theorem 1.3 : *Given $\kappa \in \mathbf{N}$ and $k \geq 0$, there is some $K = K(\kappa, k)$ with the following property. Suppose that $\kappa(\Sigma) = \kappa$, and that M is a hyperbolic 3-manifold with a homotopy equivalence to Σ as above. If $(\gamma_i)_{i=0}^p$ is a tight geodesic in $\mathcal{G}$ with $\gamma_0, \gamma_p \in X(M, k)$, then $\gamma_i \in X(M, K)$ for all $i = 1, \ldots, p - 1$.*

Note that the quasi-convexity of $X(M, k)$ tells us immediately that each γ_i lies a bounded distance from $X(M, k)$. This is true without any tightness assumptions. Tightness is needed to bound the length of the realisation of γ_i itself.

The argument presented here does not give an effective means of computing K from

κ and k. This also applies to the other results presented here.

Given $A, B \subseteq X$, write $d(A,B) = \min\{d(a,b) \mid a \in A, b \in B\}$. If $\alpha \in MX$, we write $d(\alpha, B) = d(X(\alpha), B)$.

We have the following variation on Theorem 1.3:

Theorem 1.4 : *Given $\kappa \in \mathbf{N}$ and $k \geq 0$, there is some $K' = K'(\kappa, k)$ and $r_0 = r_0(\kappa, k)$ such that if Σ and M are as in Theorem 1.3, and $(\gamma_i)_{i=0}^p$ is a tight geodesic in $\mathcal{G}$, then $\gamma_i \in X(M, K')$ for all i with $r+r_0 \leq i \leq p-r-r_0$, where $r = \max\{d(\gamma_0, X(M,k)), d(\gamma_p, X(M,k))\}$.* ∎

In other words, we are weakening the hypotheses on γ_0 and γ_p, but the conclusion is also weakened through the introduction of the constant r_0. We remark that the hyperbolicity of $\mathcal{G}$ and the quasiconvexity of $X(M,k)$ (provided k is large enough) are sufficient to bound the distance of each such γ_i from $X(M,k)$ in terms of κ and k, though, of course, this statement is much weaker.

We also note the following refinement of Theorem 1.3 relating to curves in subsurfaces.

Suppose that Φ is a closed subset of Σ, homeomorpic to a compact surface. We write $\partial_\Sigma \Phi$ for the boundary of Φ is Σ (in the general sense of topological spaces).

Definition : We say that Φ is *incompressible* if each component of $\partial_\Sigma \Phi$ is a non-trivial, non-peripheral curve.

Writing $\partial \Sigma$ and $\partial \Phi$ for the manifold boundaries, then if Φ is incompressible, we have $\partial_\Sigma \Phi = \partial \Sigma \setminus \partial \Phi$. In this case, we refer to $\partial_\Sigma \Phi$ as the *relative boundary* of Φ in Σ. (Note that our assumptions imply that any annular component of $\Sigma \setminus \Phi$ is bounded by two components of $\partial_\Sigma \Phi$.)

We write $X(\Sigma, \partial_\Sigma \Phi) \subseteq X$ for the set of curves arising in this way. (Note that two components of $\partial_\Sigma \Phi$ might get identified in $X(\Sigma, \partial_\Sigma \Phi)$.) We note that $X(\Sigma, \partial_\Sigma \Phi)$ forms a multicurve in Σ. We write $X(\Phi) \subseteq X$ for the set of curves that can be homotoped into Φ, and are not peripheral in Φ. Thus, $X(\Phi)$ is the vertex set of the curve graph $\mathcal{G}(\Phi)$, so we can talk about tight geodesics in $X(\Phi)$.

Theorem 1.5 : *Given $\kappa \in \mathbf{N}$ and $k \geq 0$, there is some $K'' = K''(\kappa, k)$ such that if $\Phi \subseteq \Sigma$ is an incompressible surface, M is a hyperbolic 3-manifold admitting a homotopy equivalence to Σ as above, and if $(\gamma_i)_{i=0}^p$ is a tight geodesic in $X(\Phi)$ with $X(\partial_\Sigma \Phi) \subseteq X(M,k)$ and $\gamma_0, \gamma_p \in X(M,k)$, then $\gamma_i \in X(M, K'')$ for all i.*

(In fact, the hypothesis on $X(\partial_\Sigma \Phi)$ can be omitted, using Theorem 1.3. If $p \geq 3$, then $\gamma_0 \cup \gamma_p$ fills Φ and so $\gamma_0, \partial_\Sigma \Phi, \gamma_p$ is a tight multigeodesic in Σ. If $p = 2$, then any curve arises in a tight geodesic in Φ also arises in a tight geodesic in Σ. In either case, the lengths of such curves in M are bounded using Theorem 1.3.)

In the exceptional cases, when $\kappa(\Sigma) = 0$, the modified curve graph is a Farey graph, and all geodesics are deemed "tight". In this case we prove a variation of Theorem 1.3, namely Proposition 9.3.

We should comment on the relation of these results with those of Minsky as presented in [Mi4]. Firstly, in [Mi4], everything is expressed in terms of hierarchies, thereby incorporating the bounds for tight geodesics for Σ, and for subsurfaces of Σ, into a single statement. The bounds are stated only with reference to the end invariants of M. The key point in the argument, however, is the quasisconvexity of $X(M,k)$. From this (in [Mi4]) it follows that any geodesic, π, in $\mathcal{G}(\Sigma)$ connecting the two end invariants (in an appropriate sense) lies a bounded distance from $X(M,k)$ (for suitable k). This in turn means that any point of π gets mapped a bounded distance under a certain quasi-projection map to $X(M,k)$. This is the key fact that is used in subsequent discussions. However, the same reasoning would apply to any geodesic segment, π, whose initial and final vertices lie a bounded distance from $X(M,k)$. This would then give statements corresponding to our Theorems 1.3 and 1.4. In other words, they are essentially proven in [Mi4], even if not explicitly stated in this way. In all of this, we are making the qualification that we are using here a slightly weaker notion of tight geodesic, as we noted earlier. We remark that Minsky's statement can, conversely, be recovered from our Corollary 8.3. This is explained in [Bow5].

The strategy of the proofs, in the non-exceptional case, will be to argue by contradiction. If no such bounds existed, then one could find a hyperbolic 3-manifold, M, and a tight geodesic, $(\gamma_i)_i$, with the curves, γ_i, extremely long. They will thus tend to "fill up" a subset of M, which on projecting to Σ gives rise to a subsurface. These subsurfaces can then be used to shortcut our geodesic $(\gamma_i)_i$, hence giving a contradiction. Much of the technical argument of the paper will be involved in recognising these subsurfaces. The basic idea behind this will be illustrated first in the 2-dimensional context in proving Theorem 1.1.

Throughout Sections 2 to 8, we shall assume that $\kappa(\Sigma) > 0$.

2. Subsurfaces.

In this section, we make some observations regarding subsurfaces of Σ. The main goal will be a proof of Lemma 2.1 — the "2/3 lemma". Note that subsurfaces are not in general assumed to be connected, nor indeed to be non-empty. Usually they are considered as defined up to homotopy.

Recall that a connected subsurface, $F \subseteq \Sigma$, is "incompressible" if each relative boundary component of F is essential and non-peripheral.

Definition : A subsurface, $F \subseteq \Sigma$ is *efficient* if each connected component of F is incompressible, and no two annular components of F are homotopic.

(Note that we allow an annular component of F to be homotopic to the boundary component of a non-annular component of F.)

Given two subsurfaces, F, G, defined up to homotopy, we write $F \subseteq G$ mean that F can be homotoped into G. We write $F \leq G$ to mean that F is a union of components of G, and say that F is a *full subsurface* of G. Any full subsurface of an efficient subsurface is efficient.

Definition : We say that two (efficient) subsurfaces, F, G, of Σ are *compatible* if there is an efficient subsurface, H, of Σ, such that $F \leq H$ and $G \leq H$.

In other words the components of F and G are either equal or disjoint.

In this definition we can always assume that $H = F \cup G$, in which case, H is well defined up to homotopy. We can thus talk about the *efficient union*, $F \cup G$, of two compatible efficient surfaces. Note that in this case, the intersection, $F \cap G$, is also well-defined, and $F \cap G \leq F \leq F \cup G$.

We say that two efficient surfaces, F, G, are *disjoint* if they are compatible and $F \cap G = \emptyset$ in the above sense. In other words, F and G are disjoint if and only if they have no common annular components (up to homotopy) and can be homotoped to be genuinely disjoint.

We can apply the above terminology to any finite collection of efficient subsurfaces.

To any subsurface, $F \subseteq \Sigma$, we can canonically associate an efficient surface $\Phi(F)$ as follows. First throw away any component of F that can be homotoped to a point or into $\partial\Sigma$. Now add to F any components of $\Sigma \setminus F$ that are discs or peripheral annuli in Σ. Finally add to F any component of $\Sigma \setminus F$ that is an annulus lying between two (homotopic) annuli components of F. This is a well-defined operation — if F is homotopic to F' then $\Phi(F)$ is homotopic to $\Phi(F')$.

Suppose F is any subset of Σ. We shall write $X(\Sigma, F) \subseteq X(\Sigma)$ for the set of curves that can be homotoped into F. (We only care here about nice locally connected subsets.)

Definition : The *component structure* of $X(\Sigma, F)$ consists of the the collection of subsets $X(\Sigma, G)$ as G ranges over the set of components of F.

Thus, formally it is a subset of the power set of $X(\Sigma, F)$. Note that it is possible that different components of F might give rise to the same (singleton) subset of $X(\Sigma, F)$. The component structure does not take account of such multiplicities.

Note that if Φ is an incompressible surface then $X(\Sigma, \Phi)$ is the disjoint union of $X(\Phi)$ and $X(\partial\Phi)$ as defined in Section 1.

If F is a subsurface, then $\Phi(F)$ is uniquely determined as the efficient surface with $X(\Sigma, \Phi(F)) = X(\Sigma, F)$, respecting the component structure of $X(\Sigma, F)$. (By "respecting the component structure" we mean that the component structures arising from $\Phi(F)$ and F are identical.)

Given $F, G \subseteq \Sigma$, we write $d(F, G) = d(X(\Sigma, F), X(\Sigma, G))$. Thus, if F, G are compatible, then $d(F, G) \leq 1$.

Definition : A sequence, $(F_i)_{i=1}^{p}$, of efficient surfaces is *compatible* if F_i and F_{i+1} are compatible for all $i = 1, \ldots, p-1$.

Definition : A compatible sequence $(F_i)_{i=1}^p$ is *taut* at index i if $F_i = (F_i \cap F_{i-1}) \cup (F_i \cap F_{i+1})$, i.e. if each component of F_i is also a component of either F_{i-1} or F_{i+1}.

Definition : A sequence $(F_i)_{i=1}^p$ is *taut* if it is taut for all i.

We interpret this to mean that $F_1 \subseteq F_2$ and $F_p \subseteq F_{p-1}$. In particular, $p \geq 2$.

(Later in the paper, an efficient sequence of surfaces will usually arise out of a geodesic in $\mathcal{G}(\Sigma)$. Tautness then arises out of the tightness assumption on such a geodesic.)

If $(F_i)_{i=1}^p$ is a compatible sequence of non-empty surfaces, then it is easily seen that $d(\alpha, F_p) \leq p - 1$ for all $\alpha \in X(F_1)$. Given tautness, we can do better:

Lemma 2.1 : *If $F_1, \ldots, F_p$ is a taut sequence of non-empty surfaces, then $d(F_1, F_p) \leq \left[\frac{2}{3}p\right] - 1$.*

Here $[.]$ denotes integer part. We first make a couple of observations about the function $f : \mathbf{N} \longrightarrow \mathbf{N}$ given by $f(p) = \left[\frac{2}{3}p\right] - 1$ for $p \geq 2$.

First note that $f(p) \leq p - 2$ (so that $d(F_1, F_p) \leq p - 2$).

For the proof, we also note that $f(p) \geq \left[\frac{p-2}{2}\right]$. Also, given $p, q \geq 2$, we have $\left[\frac{2}{3}p\right] + \left[\frac{2}{3}q\right] \leq \left[\frac{2}{3}(p+q)\right]$ and so $f(p) + f(q) + 1 \leq f(p+q)$.

Proof of Lemma 2.1 : Let us first deal with the case where $F_i \cap F_{i+1} \neq \emptyset$ for all $i = 1, \ldots, p-1$. If $p = 2$, then $F_1 = F_2$, so $d(F_1, F_2) = 0 = f(2)$ as required. If $p \geq 3$, set $r = \left[\frac{p-1}{2}\right]$, so that p is either $2r+1$ or $2r+2$. For $i = 1, \ldots, r$ let $G_i = F_{2i} \cup F_{2i+1}$. Thus each G_i is an efficient surface and for all $i = 1, \ldots, r-1$, G_i and G_{i+1} have a common component. Moreover, $F_1 \subseteq G_1$ and $F_p \subseteq G_r$. It now follows easily by induction that $d(F_1, F_p) \leq d(F_1, G_r) \leq r = \left[\frac{p-1}{2}\right] \leq f(p)$ as required.

We can now deal with the general case by induction on p. By the above, we can suppose that for some $q < p$, $F_q \cap F_{q-1} = \emptyset$. Since $(F_i)_{i=1}^p$ is taut, it follows that $F_q \leq F_{q-1}$, and so the subsequence $(F_i)_{i=1}^q$ is taut. Similarly, $F_{q+1}, \ldots, F_p$ is taut. (In particular, $2 \leq q \leq p-2$). By the inductive hypothesis, we have $d(F_1, F_q) \leq f(q)$ and $d(F_{q+1}, F_p) \leq f(p-q)$. Since $F_q \cap F_{q+1} = \emptyset$, it follows that $d(F_1, F_p) \leq f(q) + f(p-q) + 1 \leq f(p)$ as required. $\diamondsuit$

Corollary 2.2 : *If $(F_i)_{i=1}^p$ is a compatible sequence of efficient surfaces that is taut for $i = 2, \ldots, p-1$, then $d(F_1, F_p) \leq \left[\frac{2}{3}p\right] + 1$.*

Proof : Setting $F_0 = F_1$ and $F_{p+1} = F_p$, the sequence $F_0, \ldots, F_{p+1}$ is taut (for all indices). Thus, $d(F_1, F_p) = d(F_0, F_{p+1}) \leq \left[\frac{2}{3}(p+2)\right] - 1 \leq \left[\frac{2}{3}p\right] + 1$. $\diamondsuit$

3. Laminations.

In this section we discuss how laminations fill subsurfaces and hence give proofs of Theorems 1.1 and 1.2. Our account of the proof will be a bit more elaborate than necessary, though the machinery we described will be re-used for our applications to 3-manifolds.

We fix a hyperbolic structure on Σ such that $\partial\Sigma$ is totally geodesic (i.e. in the riemannian sense of having zero extrinsic curvature). The constructions we describe will, up to homotopy, be independent of this structure. Here we deal only with geometric laminations. (We will have no use for transverse measures.)

A *lamination* consists of a non-empty compact subset $\lambda \subseteq \Sigma$ that is a disjoint union of geodesic leaves. (See, for example, [CasB] for general background.) Each leaf is either a bi-infinite simple (local) geodesic, or else a non-peripheral simple closed geodesic. (Note that no leaf spirals onto a boundary component.) We say that a lamination is *minimal* if it contains no proper sublamination. A minimal lamination is either a closed geodesic or has uncountably many leaves (transversely a Cantor set).

Any lamination, λ, contains a finite number of disjoint minimal laminations whose union we denote by $\nu(\lambda)$. Moreover $\lambda \setminus \nu(\lambda)$ consists of a finite number of bi-infinite "isolated" leaves spiralling into $\nu(\lambda)$.

We write $N(\lambda, t)$ for the t-neighbourhood of λ in Σ. For all sufficiently small $t > 0$, the inclusion of $N(\lambda, u)$ into $N(\lambda, t)$ is a homotopy equivalence for all $u \in (0, t)$. This gives rise to a well-defined subsurface, $\Phi(N(\lambda, u)) = \Phi(N(\lambda, t))$, which we denote by $\Phi(\lambda)$. We refer to $\Phi(\lambda)$ as the subsurface *filled* by λ. Up to isotopy, it can be equivalently defined by taking the union of λ together with all complementary components that are discs or peripheral annuli, and then thickening up each closed curve component of λ to an annulus.

Note that the components of λ are in bijective correspondence to the components of $\Phi(\lambda)$. If μ is a sublamination of λ, then clearly $\Phi(\mu) \subseteq \Phi(\lambda)$ (in the sense described in Section 2). In particular, $\Phi(\nu(\lambda)) \subseteq \Phi(\lambda)$. Moreover, $\Phi(\nu(\mu)) \leq \Phi(\nu(\lambda))$ — each minimal sublamination of μ is also a minimal sublamination of λ. Note also that any multicurve, $\gamma \in MX$, can be viewed as (i.e. uniquely realised as) a lamination, and that $\Phi(\gamma)$ is a regular neighbourhood of γ.

Remark : In what follows, the existence of spiralling leaves in a lamination, λ, i.e. those of $\lambda \setminus \nu(\lambda)$, will be a complicating factor. The issues involved are mainly technical, and one could follow the overall logic by simply pretending that such things did not exist, i.e. that $\nu(\lambda) = \lambda$ for any lamination λ. In this case, the subsurfaces F and G defined below will be identical. In this way, one could eliminate certain technical arguments from a first reading: for example Lemmas 6.2, 6.3 and 6.4. In Section 7, one could forget about the arcs a^n and c^n, and ignore Lemmas 7.3 and 7.4. Various bits of Section 8 are similarly simplified. Of course, such an assumption is unjustified in general.

Before continuing, we recall the general definition of "Hausdorff convergence". Let K be a compact metric space. Let $\mathcal{C} = \mathcal{C}(K)$ be the set of closed subsets of K. Given $P, Q \in \mathcal{C}$ let $\mathrm{hd}(P, Q)$ be the minimum $r \geq 0$ such that $P \subseteq N(Q, r)$ and $Q \subseteq N(P, r)$. In other words, hd is the Hausdorff distance on $\mathcal{C}$. With this structure, (K, hd) is a compact metric space. In the case where $K = \Sigma$, the set of all laminations is closed in the Hausdorff topology. In particular, any sequence of multicurves has a subsequence converging on a lamination.

Suppose that $\underline{\alpha} = (\alpha^n)_{n=0}^{\infty}$ is a sequence of multicurves in Σ. After passing to a subsequence, we can suppose that α^n can be written as a disjoint union $\alpha^n = \check{\alpha}^n \sqcup \hat{\alpha}^n$ of possible empty multicurves, where the total length of $\check{\alpha}^n$ remains bounded, and where the minimal length of the components of $\hat{\alpha}^n$ tends to ∞. Passing to a further subsequence, we can suppose that $\check{\alpha}^n = l \in MX \cup \{\emptyset\}$ is constant, and that $\hat{\alpha}^n$ converges on a lamination, $\lambda \subseteq \Sigma$ (i.e. with respect to Hausdorff distance). Note that l and λ are disjoint, so that $l \sqcup \lambda$ is itself a lamination. Assume $\underline{\alpha}$ satisfies the above, and write $F(\underline{\alpha}) = \Phi(\nu(\lambda))$, $G(\underline{\alpha}) = \Phi(\lambda)$ and $H(\underline{\alpha}) = \Phi(l)$. Thus, $F(\underline{\alpha}) \subseteq G(\underline{\alpha})$ and $G(\underline{\alpha}) \cap H(\underline{\alpha}) = \emptyset$. Moreover, for all sufficiently large n, $\alpha^n \subseteq G(\underline{\alpha}) \cup H(\underline{\alpha})$, and if J is any component of $F(\underline{\alpha})$, then α^n crosses some element of $X(\Sigma, J)$. (In other words, some component of α^n either crosses a boundary component of J or lies in $X(J)$.) Thus, intuitively, H represents the bounded part of α^n, while the rest runs around G getting longer and longer, with most of it filling up the subsurface F of G.

Now suppose that $\underline{\beta} = (\beta^n)_{n=0}^{\infty}$ is another sequence of multicurves compatible with $\underline{\alpha}$, i.e. for all n, $\alpha^n \cup \beta^n$ is a multicurve. After passing to subsequences as above, we can suppose that $\beta^n \longrightarrow m \sqcup \mu$ for disjoint laminations, m, μ, so we have subsurfaces $F(\underline{\beta})$, $G(\underline{\beta})$ and $H(\underline{\beta})$. Now $\alpha^n \cup \beta^n \longrightarrow (l \cup m) \sqcup (\lambda \cup \mu)$. In particular, $\nu(\lambda \cup \mu)$ is a union of components of $\nu(\lambda)$ and $\nu(\mu)$. We conclude that $H(\underline{\alpha})$ and $H(\underline{\beta})$ are compatible with each other and compatible and disjoint from both $G(\underline{\alpha})$ and $G(\underline{\beta})$. Moreover $F(\underline{\alpha})$ and $F(\underline{\beta})$ are compatible. In fact, each component of $F(\underline{\alpha})$ is either a component of $F(\underline{\beta})$ or disjoint from $G(\underline{\beta})$.

Now suppose that for each $n \in \mathbf{N}$, we have a path of multicurves $(\alpha_i^n)_{i=1}^p$ for $p \in \mathbf{N}$ fixed, in other words, α_i^n and α_{i+1}^n are compatible for all i and n. Now, after passing to a subsequence in n, for each $i = 1, \ldots, p$, we can suppose that each of the sequences $\underline{\alpha}_i$ satisfies the above condition. We thus obtain efficient surfaces, $F_i = F(\underline{\alpha}_i)$, $G_i = G(\underline{\alpha}_i)$ and $H_i = H(\underline{\alpha}_i)$. We can suppose that $\check{\alpha}_i^n$ is independent of n and that the property of $\hat{\alpha}_i^n$ being non-empty is also independent of n.

Now $(F_i)_i$, $(G_i)_i$ and $(H_i)_i$ satisfy the following for all i, j with $|i - j| = 1$ and for all (sufficiently large) n:

(F1) $F_i \subseteq G_i$,
(F2) $G_i \cap H_i = \emptyset$,
(F3) $\hat{\alpha}_i^n \subseteq G_i$ and $H_i = \Phi(\check{\alpha}_i^n)$ (i.e. a regular neighbourhood of $\check{\alpha}_i^n$),
(F4) If J is a component of F_i then α_i^n crosses some element of $X(\Sigma, J)$,
(F5) If $\hat{\alpha}_i^n = \emptyset$ then $G_i = \emptyset$.
(F6) If $\hat{\alpha}_i^n \neq \emptyset$, then $F_i \neq \emptyset$,
(F7) $G_i \cap H_j = \emptyset$,
(F8) If J is a component of F_i, then either $J \cap G_j = \emptyset$, or J is a component of F_j. (In particular, F_i and F_j are compatible.)
(F9) If $(\alpha_i^n)_{i \in I}$ is tight at $i \in I$ for all n, then either $F_i = \emptyset$ or $(F_i)_{i \in I}$ is taut at i.

Properties (F1)–(F8) follow from the earlier observations. We need to verify (F9). The tightness hypothesis means that every curve crossing α_i^n crosses either α_{i-1}^n or α_{i+1}^n (or both). Let J be a component of F_i. We want to show that J is also a component of

F_{i-1} or F_{i+1} (or both). If not, then by (F8), $J \cap G_{i-1} = J \cap G_{i+1} = \emptyset$, so that by (F3), $J \cap \alpha^n_{i-1} = J \cap \alpha^n_{i+1} = \emptyset$. By (F4), α^n_i crosses some element, $\beta \in X(J)$. By tightness, β crosses either α^n_{i-1} or α^n_{i+1}, giving a contradiction.

We have already done more work than we need to prove Theorems 1.1 and 1.2.

Proof of Theorem 1.1 : Suppose $\alpha, \beta \in X$ and that there are infinitely many tight geodesics from α to β. We can find tight multigeodesics $(\alpha^n_i)^p_{i=0}$ with $\alpha^n_0 = \alpha$ and $\alpha^n_p = \beta$ such that the total length of the α^n_i go to ∞ for at least one i. We write I_∞ for the set of such i.

After passing to a subsequence we have surfaces $(F_i)_i$, $(G_i)_i$, $(H_i)_i$ as above, and where $I_\infty \neq \emptyset$. Note that $H_0 = \Phi(\alpha)$ and $H_p = \Phi(\beta)$. Let $s, s+1, \ldots, t$ be a maximal set of consecutive indices in I_∞. Thus $F_{s-1} = F_{t+1} = \emptyset$, and by (F9), $F_s, F_{s+1}, \ldots, F_t$ is a taut compatible sequence of surfaces. Any taut sequence of surfaces must have length at most 2, and so $t > s$. By Lemma 2.1, we have $d(F_s, F_t) \leq \left[\frac{2}{3}(t-s+1)\right] - 1 \leq t - s - 1$. Moreover, $H_{s-1} = \Phi(\alpha^n_{s-1})$ is disjoint from F_s and $H_{t+1} = \Phi(\alpha^n_{t+1})$ is disjoint from F_t. Thus $d(\alpha^n_{s-1}, \alpha^n_{t+1}) \leq d(F_s, F_t) + 2 \leq t - s + 1$. But $(\alpha^n_i)_i$ is geodesic, so in particular, $d(\alpha^n_{s-1}, \alpha^n_{t+1}) = (t+1) - (s-1) = t - s + 2$ giving a contradiction. $\diamondsuit$

Proof of Theorem 1.2 : Suppose the conclusion fails. Then we can find tight geodesics, $(\gamma^n_i)^p_{i=1}$ with $d(\alpha, \gamma^n_0) \leq r$ and $d(\beta, \gamma^n_0) \leq r$, and such that the total length of γ^n_i goes to ∞ for at least one i with $12r \leq i \leq p - 12r$. We can suppose that $d(\alpha, \gamma^n_0) = u$ and $d(\beta, \gamma^n_p) = v$ are fixed. We can then extend to a path of multicurves, $(\gamma^n_i)^{p+v}_{i=-u}$ with $\gamma^n_{-u} = \alpha$ and $\gamma^n_{p+v} = \beta$. Passing to a subsequence we can find subsurfaces $(F_i)_i$, $(G_i)_i$ and $(H_i)_i$ as before, with $H_{-u} = \Phi(\alpha)$, and $H_{p+v} = \Phi(\beta)$. Note that $d(\alpha, \beta) \geq p - 2r$.

Let I_∞ be the set of $i \in \{0, 1, \ldots, p\}$ such that the length of the γ^n_i go to ∞. Let $\{s, s+1, \ldots, t\}$ be a maximal set of consecutive indices in I_∞, meeting $\{12r, \ldots, p - 12r\}$. Using Theorem 1.1, we see that either $s = 0$ or $t = p$ (or both). In particular, $t - s \geq 12r$. By Lemma 2.1, we have $d(F_s, F_t) \leq [\frac{2}{3}(t-s+1)] - 1 < \frac{2}{3}(t-s)$, and so $(t-s) - d(F_s, F_t) > \frac{1}{3}(t-s) \geq 4r$. We can therefore shortcut as in the proof of Theorem 1.1 to obtain a path from α to β in $\mathcal{G}$ of length less than $(p + u + v) - 4r \leq (p + 2r) - 4r = p - 2r$. But $d(\alpha, \beta) \geq p - 2r$, giving a contradiction. $\diamondsuit$

The essential point of the above arguments is that when closed geodesics in Σ are very long, they fill up subsurfaces which can be used to shortcut paths in the curve complex. To prove the remaining results, we will want to carry out a similar argument in a 3-manifold.

4. 3-manifolds.

Before giving formal definitions, we set them in context by giving an outline of what we hope to achieve.

The overall strategy is to adapt the geometric arguments of Section 3 to the context of hyperbolic 3-manifolds. One problem in doing this that in order to get uniform bounds, we need to allow the geometry on our 3-manifold to change. Thus, if we were to apply the same idea directly, we would need to pass to a geometric limit of 3-manifolds. However, such a limiting manifold can be quite complicated (cf. [So]). Instead, we shall stop the process at some sufficiently late stage (large n), for which we are able to recognise the subsurfaces filled by our closed multicurves.

There is however, a particular situation in which passing to a limit is indeed feasible. In order to motivate later definitions, we consider this first. Suppose, for example, we were to weaken the statement of Theorem 1.3, by allowing the constant K to depend also on the length p of our geodesic in $\mathcal{G}$, and also on the injectivity radius of M. We argue again by contradiction, similarly as in Section 3. This time we have closed geodesics realised in a 3-manifold M^n, whose structure may depend on n. We can however, assume here that the injectivity radii of the M^n are uniformly bounded below by a positive constant. Up to the action of the mapping class group of Σ, we can pass to a geometric limit manifold, M, also homeomorphic to $\operatorname{int}\Sigma \times \mathbf{R}$. The closed geodesics in M^n, converge to a sequence of laminations, $(\lambda_i)_i$, each realised in M. It is not hard to describe subsurfaces of Σ filled by such laminations. (This is most conveniently done by lifting to the projectivised tangent bundle of M, where the realised laminations become embedded.) One can now go on to derive a similar contradiction via the 2/3 lemma.

For the general case, we need to remove the dependence on p, and on the injectivity radius. For the former, we use another argument, described in Section 8, which shows that, without any restriction on p, we can get length bound for some uniformly spaced subsequence of vertices of our geodesic in $\mathcal{G}$. After this, we can then apply the result for bounded p as above.

Removing the dependence on injectivity radius is more involved, and is the reason behind much of the discussion of this section. We must avoid geometric limits, and therefore be able to recognise our subsurfaces of Σ at a sufficiently late, but finite, stage M^n. This will use a version of Thurston's Uniform Injectivity Theorem (see [T]). Since the theorem applies only to the thick part of the 3-manifold, we need to be able to focus our attention on this. This will be justified by the "Tube Penetration Lemma" of Section 5.

We now move on to more formal definitions.

Let M be a complete hyperbolic 3-manifold admitting a type preserving homotopy equivalence $\chi : M \longrightarrow \Sigma$. To simplify the discussion, let us assume, for the moment, that there are no accidental cusps, so that the cusps of M correspond exactly to the boundary components of Σ. Any non-peripheral closed curve in Σ can be realised uniquely as a closed geodesic in M which we usually denote also by α. We write $L(M,\alpha)$ for its length. We can also realise a multicurve γ as a disjoint union of closed geodesics in M, and write $L(M,\gamma)$ for the maximum of $L(M,\alpha)$ as α varies over the components of γ. Given $k \geq 0$, we write $X(M,k) = \{\alpha \in X(\Sigma) \mid L(M,\alpha) \leq k\}$.

If we fix certain "Margulis constants" described below, then we obtain a set, $\mathcal{P}$, of Margulis cusps, and a set $\mathcal{T}$ of Margulis tubes, all of which are mutually disjoint. The boundary, ∂P, of each $P \in \mathcal{P}$ is euclidean cylinder, foliated by euclidean circles of fixed length. We write $\Psi(M) = M \setminus \bigcup_{P \in \mathcal{P}} \operatorname{int} P$ for the *non-cuspidal* part of M. We can assume that χ maps $\Psi(M)$ onto Σ, and $\partial\Psi(M)$ onto $\partial\Sigma$. Indeed, by tameness [Bon], $\Psi(M)$ is homeomorphic to $\Sigma \times \mathbf{R}$. A result of Otal [O] tells us that (the core curves of) the Margulis tubes are unknotted and unlinked in $M \cong \Sigma \times \mathbf{R}$ (in particular, the core curves are all simple in Σ). If $T \in \mathcal{T}$, then ∂T is a euclidean torus. It has a preferred homotopy class of longitude (homotopic to infinity in $M \setminus \mathrm{int}\, T$) and hence a canonical foliation by euclidean circles. We write $\Theta(M) = \Psi(M) \setminus \bigcup_{T \in \mathcal{T}} \mathrm{int}\, T$ for the *thick part* of M.

We should comment on the choice of Margulis constants. Firstly, we can choose $\eta_0 > 0$ small and suppose that the core curves of the Margulis tubes are precisely the elements of $X(M, \eta_0)$. Given $\delta \in X(M, \eta_0)$, we get a one-parameter family of tubes about δ of differing radii. A convenient way to normalise is to let $T(\delta, \eta)$ be the tube such that each euclidean longitude of $\partial T(\delta, \eta)$ (as described above) has length η. As discussed in [Bow4], using ideas of Otal, we can choose the Margulis tube about δ to have the form $T(\delta, \eta_1)$ for some fixed $\eta_1 > 0$. (Here, η_1 may depend on $\kappa(\Sigma)$.) We can also assume that the longitudes of ∂P have length η_1 for all $P \in \mathcal{P}$.

For future reference, we note that if $\eta, r > 0$ satisfy $L(M, \delta) \leq \eta < e^r \eta \leq \eta_1$, then $N(T(\delta, \eta), r) \subseteq T(\delta, \eta e^r)$. (Indeed $N(T(\delta, \eta), r) = T(\delta, \eta')$ for some $\eta' \in [\eta, e^r \eta]$.) In particular, it follows easily that given $l \geq 0$, there is some $\eta(l) \geq 0$ such that if $\eta \leq \eta(l)$ and $\gamma \in X(M, l) \setminus \{\delta\}$ then $\gamma \cap T(\delta, \eta) = \emptyset$. Moreover, each closed geodesic in M corresponding to a curve in $X(\Sigma)$ lies in $\Psi(M)$. (This is a consequence of the existence of pleated surfaces, as discussed below.) Thus, if we have some fixed bound, l, we can assume we have chosen η so that if $\gamma \in X(M, l)$, then either γ lies in the thick part of M or is the core curve of a Margulis tube.

We will need to use some notion of a pleated surface (see, for example, [CanEG].) Here we are only interested in the case where the "pleating locus" is a multicurve, which will avoid many of the technicalities. It will also be enough for us to define pleated surfaces as uniformly lipschitz maps. We shall (for the moment) use the following formulation. (We shall modify this a little in Section 7.)

Suppose $\gamma \in MX$ is realised as a multigeodesic, $\gamma \subseteq M$. There is a finite area complete hyperbolic surface S, a geodesic multicurve, $\alpha \subseteq S$, and a uniformly lipschitz homotopy equivalence, $\phi : S \longrightarrow M$ such that $\phi|\alpha$ is a local isometry onto γ. Note that $\omega = \omega_\phi = \chi \circ \phi : S \longrightarrow \Sigma$ is a homotopy equivalence. Thus α is the unique geodesic realisation of the multicurve $\omega^{-1}\gamma$. Clearly the length, $L(S, \alpha)$, of α in S equals $L(M, \gamma)$. If $P \in \mathcal{P}$, then $\phi^{-1} P$ is a neighbourhood of a cusp in S. Indeed we can assume that it is bounded by a horocycle, and that ϕ maps this horocycle to a longitude of ∂P. By "uniformly lipschitz", we mean μ-lipschitz, where $\mu > 0$ depends only on $\kappa(\Sigma)$. In fact, here we can take $\mu = 1$, though the extra generality will be useful in Section 7. We refer to S, ϕ as a "pleating surface" for γ. We write ρ_S for the metric on the surface S. Note that we can choose the Margulis constant defining the (boundary) cusps of M to be a fixed number, which will remain constant throughout this paper.

A useful fact about pleated surfaces is the "Uniform Injectivity Theorem". There are several versions of this, the original being due to Thurston [T]. Here we only need a fairly weak form. It is a special case of that described in [Mi2].

Let E be the projectivised unit tangent space (or 1-dimensional Grassmannian) of M which carries a natural riemannian metric, ρ_E. Let $\pi : E \longrightarrow M$ be the projection map. We can lift any geodesic multicurve of γ in M to an embedded multicurve in E, which we also denote by γ. Composing $\phi|\alpha$ with this lift, we get a locally isometric homeomorphism $\psi : \alpha \longrightarrow \gamma$. This map is also uniformly lipschitz with respect to the metric ρ_S and ρ_E. Moreover, the uniform injectivity theorem tells us that:

Theorem 4.1 : *There is a homeomorphism, $g : [0, \infty) \longrightarrow [0, \infty)$, depending only on $\kappa(\Sigma)$ and the choice of Margulis constant such that if $\gamma \subseteq \Theta(M)$ is a multicurve lying in the thick part of M, and $\phi : S, \alpha \longrightarrow M, \gamma$ is a pleating surface for γ, then the lift, $\psi : \alpha \longrightarrow \gamma$ is g-injective, i.e. for all $x, y \in \alpha$, $\rho_S(x, y) \leq g(\rho_E(\psi x, \psi y))$.*

The usual proofs of this rely on geometric limit arguments, and do not enable us to explicitly estimate $g(t)$ as a function of $\kappa(\Sigma)$ and t (though the referee has indicated that it could be made explicit, using a somewhat different argument).

Pleating surfaces give us a means of defining a "quasiprojection" in the curve complex. Given $k \geq 0$, and a hyperbolic surface, $S \cong \operatorname{int} \Sigma$, we write $X(S, k) = \{\alpha \in X(\Sigma) \mid L(S, \alpha) \leq k\}$. A simple argument gives a bound on its diameter, $\operatorname{diam} X(S, k)$, in terms of $\kappa(\Sigma)$ and k. (This is sometime called the "Bers Lemma".) Moreover, there is some $k_1 > 0$, depending only on κ such that $X(S, k_1) \neq \emptyset$.

Suppose $\phi : S, \alpha \longrightarrow M, \gamma$ is a pleating surface for the geodesic multicurve $\gamma \subseteq M$. Note that $\omega_\phi X(S, k) \subseteq X(M, \mu k)$, where μ is the uniform lipschitz constant. In particular, setting $k_0 = \mu k_1$, we get a non-empty bounded subset, $\omega_\phi X(S, k_1)$ of $X(M, k_0)$. We need to observe that this is coarsely well defined. More precisely:

Lemma 4.2 : *If $\phi : S, \alpha \longrightarrow M, \gamma$ and $\phi' : S', \alpha' \longrightarrow M, \gamma$ are pleating surfaces for the same curve γ, then $\operatorname{diam}(\omega_\phi X(S, k_1) \cup \omega_{\phi'} X(S', k_1))$ is bounded by a constant depending only on κ and k_1.*

This can be proven via the Uniform Injectivity Theorem. Arguments along these lines can be found in Section 3 of [Min3] or Section 6.2 of [Min4]. However, as the referee noted, one can also give a direct proof, employing an argument found in Thurston [T] (which features in his proof of the Uniform Injectivity Theorem). Moreover, this gives explicit constants. The idea is given below.

We use, implicitly, the lemma of Thurston that says that the η-thick part of a pleating surface gets mapped into the η'-thick part of M, where η' depends only on η and the topological type. (See the discussion after Lemma 5.2 for more details.)

Proof : It is enough to find some $\epsilon \in \mathcal{G}(\Sigma)$ whose length in both S and S' is bounded (via the homotopies ϕ and ϕ'). We can assume that γ does not penetrate too deeply into any Margulis tube in M, otherwise we could take ϵ to be the core of this tube. We can therefore assume that α and α' lie in the η_0-thick parts of S and S' respectively, for some η_0 depending only on $\kappa(\Sigma)$. By a simple volume argument, there is a bound on the number of $(\eta_0/2)$-balls we can pack into the product $S \times S'$. From this, a simple argument, as in [T], gives us a constant L such that if γ has length greater than L, we can find subarcs β of α and β' of α', both mapping to the same subarc in γ, and such that the initial and final points of α are η_0-close in S', and the initial and final points of α' are η_0-close in S. We can assume the lengths of these arcs to be greater than η_0 but less than L. We can now take short arcs δ in S and δ' is S' so as to give us non-trivial closed curves $\alpha \cup \delta$ and $\alpha' \cup \delta'$. Now $\phi(\alpha \cup \delta)$ and $\phi'(\alpha' \cup \delta')$ are homotopic in M, and so $\omega_\phi(\alpha \cup \delta)$ and $\omega(\alpha' \cup \delta')$ are homotopic in Σ. This curve has bounded length in both surfaces. It might not be simple, but some subarc of it closes up to a simple curve, again of bounded length in both surfaces. This will serve for our curve ϵ, provided γ has length greater than L. Otherwise, we can simply take $\epsilon = \gamma$. This proves Lemma 5.2. $\diamondsuit$

Using the above, we get a curve, $\mathrm{proj}(\gamma) \in \omega_\phi X(S, k_1) \subseteq X(M, k_0)$, well-defined up to bounded distance. Moreover, if $\gamma' \in X$ is adjacent to γ, then $d(\mathrm{proj}(\gamma), \mathrm{proj}(\gamma'))$ is bounded (by considering a pleated surface of the multicurve $\gamma \cup \gamma'$.) We also note that if $\gamma \in X(M, k)$, then $d(\gamma, \mathrm{proj}(\gamma))$ is bounded in terms of k.

It was shown in [MaM1] that the curve graph is hyperbolic in the sense of Gromov (see also [Bow1]). If $k \geq k_0$, then standard hyperbolic arguments show that $X(M, k)$ is uniformly quasiconvex. ("Quasiconvex" means that any geodesic in $\mathcal{G}(\Sigma)$ with endpoints in $X(M, k)$ lies in a fixed h-neighbourhood of $X(M, k)$. "Uniform" means that this h depends only on $\kappa(\Sigma)$ and k.) Moreover, $\mathrm{proj} : X \longrightarrow X(M, k_0) \subseteq X(M, k)$ is a quasiprojection to $X(M, k)$. In other words, $\mathrm{proj}(\gamma)$ is a bounded distance from any nearest point of $X(M, k)$ to γ (see [Mi3]). A consequence we shall use later is the following. Suppose $(\gamma_i)_{i=1}^p$ is any geodesic in X and set $r = \max\{d(\gamma_0, X(M, k)), d(\gamma_p, X(M, k))\}$. Then for all i with $r \leq i \leq p - r$, γ_i is a bounded distance from $X(M, k)$. In particular, $d(\gamma_i, \mathrm{proj}(\gamma_i))$ is bounded. Here, all bounds depend only on $\kappa(\Sigma)$ and k.

The above discussion goes through in essentially the same way if we allow for accidental parabolics. We just need to modify some definitions and reinterpret some terminology. In this case, the set $\mathcal{P}$ of cusps is a disjoint union $\mathcal{P} = \mathcal{P}_\partial \sqcup \mathcal{P}_A$, where $\mathcal{P}_\partial$ is in bijective correspondence to the boundary curves of Σ. We refer to the elements of $\mathcal{P}_\partial$ and $\mathcal{P}_A$ as *boundary cusps* and *accidental cusps* respectively. Again $\mathcal{P}$ is finite. In this situation, the elements of $\mathcal{P}_A$ behave like Margulis tubes. We define $\Psi(M) = M \setminus \bigcup_{P \in \mathcal{P}_\partial} \mathrm{int}(P)$ and $\Theta(M) = M \setminus \bigcup_{C \in \mathcal{P} \cup \mathcal{T}} \mathrm{int}(C)$. Again, $\Psi(M) = \Sigma \times \mathbf{R}$, and we refer to $\Theta(M)$ as the "thick part" of M.

We write $X_A(\Sigma) \subseteq X(\Sigma)$ for the set of curves corresponding to accidental parabolics. If $\alpha \in X_A(\Sigma)$, we set $l_M(\alpha) = 0$. In this case, α can be realised as an arbitrarily short curve. For the purposes of our arguments here, we can fix some realisation of α in M which is sufficiently short in relation to the various constants introduced later. This realisation plays the same role as the core curve of a Margulis tube. In this way we still have a pleating surface realising any multicurve, even if some of its elements lie in $X_A(\Sigma)$, and Theorem 4.1 remains valid.

Given $\alpha \in X_A(\Sigma)$ and some $\eta > 0$, we write $P = P(\alpha, \eta)$ for the corresponding cusp with Margulis constant η. Note that, whereas the constant determining the boundary cusps is fixed throughout the paper, the one determining the accidental cusps will depend on other factors. At any given stage in the argument, we will use the same Margulis constant, $\eta > 0$ for the Margulis tubes and the accidental cusps.

5. Tube penetration.

The "Tube Penetration Lemma" says roughly that closed geodesics in M lying on tight geodesics in the curve graph cannot enter too deeply into Margulis tubes unless they happen to be Margulis core curves. Similarly they cannot enter too deeply into any accidental cusp. Thus, choosing the Margulis constant appropriately, we can assume that they lie in the thick part, $\Theta(M)$, of M. We describe a number of variations here. The arguments apply similar principles to those of [Mi4], though the overall logic is different. We will first describe the case of Margulis tubes, though essentially the same argument will work for accidental cusps.

Let k_0 be the constant described in Section 4, and let $k \geq k_0$. One version of the tube penetration lemma says the following:

Lemma 5.1 : *There is some $\eta > 0$, depending only on k and $\kappa(\Sigma)$ with the following property. Suppose $\gamma_0, \ldots, \gamma_p$ is a tight geodesic in $X(\Sigma)$ with $\gamma_0, \gamma_p \in X(M,k)$. If $T = T(\delta, \eta)$ is any Margulis tube, then for all i, $\gamma_i \cap T \subseteq \delta$.*

In other words, either $\gamma_i \cap T = \emptyset$, or γ_i is the core curve, δ, of T. Here, the radius of the Margulis tube, T, is determined by η, as described in Section 4.

We have stated the result for tight geodesics in X, but we shall verify the statement for a tight multigeodesic $(\gamma_i)_i$. We only need tightness at index i.

In fact, the statement can be modified:

Lemma 5.2 : *Given $\eta > 0$, $h \geq 0$ and $k \in \mathbf{N}$, there is some $\eta' \in (0, \eta]$ with the following property. Suppose that $\gamma_0, \ldots, \gamma_p$ is a multigeodesic which is tight at index i and with $d(\gamma_0, X(M,k)) \leq h$ and $d(\gamma_p, X(M,k)) \leq h$. Suppose that $\gamma_0 \cap T(\delta, \eta) \subseteq \delta$ and $\gamma_p \cap T(\delta, \eta) \subseteq \delta$. Then $\gamma_i \cap T(\delta, \eta') \subseteq \delta$.*

To deduce Lemma 5.1, note that, by definition, any tight geodesic lies inside a tight multigeodesic. Moreover, as discussed in Section 4, if $\gamma_0, \gamma_p \in X(M,k)$, then we can find some Margulis constant, $\eta(k)$, so that $(\gamma_0 \cup \gamma_p) \cap T(\delta, \eta(k)) \subseteq \delta$ for all Margulis tubes, $T(\delta, \eta(k))$. We can thus apply Lemma 5.2 starting with $\eta = \eta(k)$ and with $h = 0$.

Before starting the proof of Lemma 5.2, we make some preliminary observations. Note that by quasiconvexity of $X(M,k)$, each γ_i is a bounded distance from $X(M,k)$ (depending on $h, k, \kappa(\Sigma)$, and hence a bounded distance from its quasiprojection to $X(M,k)$.)

It is well known that if a point in a pleated surface lies deep inside a Margulis tube, then this point lies inside an essential short curve on the surface, homotopic to the core curve of the tube. We can quantify this by the following argument. To simplify notation, we will assume that all pleated surfaces are 1-lipschitz. Once can easily incorporate uniform lipschitz constants into the discussion.

First, let S be a complete finite-area hyperbolic surface homeomorphic to $\operatorname{int} \Sigma$. Given $c > 0$, there is some some $l \geq 0$, depending only on $\kappa(\Sigma)$ and c with the following property. For all $x \in S$ the image of $\pi_1(N(x,l))$ in $\pi_1(\Sigma)$ is non-trivial, and if it is cyclic, then x lies on an essential simple closed curve, ϵ, of length at most c in S. We can choose $c < l$ sufficiently small (independently of $\kappa(\Sigma)$) such that any closed simple geodesic in S that meets ϵ must in fact cross ϵ. (These observations are just basic hyperbolic geometry.)

In what follows, we identify the set of homotopy classes of curves in M and in Σ via the homotopy equivalence $\chi : M \longrightarrow \Sigma$. In particular, we can talk of two such curves as "crossing" if their homotopy classes cross in Σ.

Suppose that $T = T(\delta, \eta)$ is some Margulis tube in M and that $\gamma \subseteq M$, $\gamma \in X(\Sigma) \setminus \{\delta\}$ is a closed geodesic meeting $T(\delta, \eta_0 e^{-l})$ in some point y. Note that $N(y, l) \subseteq T$. Let $\phi : S, \alpha \longrightarrow M, \gamma$ be a (1-lipschitz) pleating surface for γ. Let $x \in \alpha$ with $\phi(x) = y$. Thus, $N(x, l) \subseteq \phi^{-1}T$. Since ϕ is a homotopy equivalence, the image of $\pi_1(N(x, l))$ is cyclic, and we can apply the paragraph before last to obtain a loop ϵ based at x of length at most c. This must lie in the free homotopy class of $[\phi]^{-1}\delta$ (where $[\phi]$ denotes the homotopy class of ϕ). Now α meets, and hence crosses, ϵ, and so γ crosses the curve δ. We can assume that $c \leq k_0$, and so $\epsilon \equiv [\phi]^{-1}\delta \in X(S, k_0)$. Thus, δ would serve as a quasiprojection of the curve γ. In particular, $d(\gamma, \delta)$ is bounded above in terms of $d(\gamma, X(M, k))$. In particular, if $(\gamma_i)_i$ is a multigeodesic as in the hypotheses of Lemma 5.2, and $\gamma \in X(\gamma_i)$ for some i, with $\gamma \cap T(\delta, \eta e^{-l}) \neq \emptyset$ then $d(\gamma, \delta) \leq R$ for some fixed constant, R (depending only on $\kappa(\Sigma)$, k and h).

Now let $\eta' = \eta e^{-(2R+1)c-l}$, so that $N(T(\delta, \eta'), (2R + 1)c + l) \subseteq T(\delta, \eta)$.

Proof of Lemma 5.2 : The basic idea is to apply the observations of the previous paragraphs. If one of the curves of $\gamma_0, \ldots, \gamma_p$ penetrates deeply into a Margulis tube, then we get trapped: all the neighbouring curves in the sequence also penetrate deeply. The only way to escape is via a curve equal to or adjacent to δ in $\mathcal{G}$. This will give a contradiction to $(\gamma_i)_i$ being geodesic in $\mathcal{G}$. Here is a formal argument.

For notational convenience, we shift indices so that we have a multigeodesic, $\gamma_{-q}, \ldots, \gamma_0, \ldots, \gamma_r$, where $(\gamma_{-q} \cup \gamma_r) \cap T(\delta, \eta) \subseteq \delta$ and this multigeodesic is tight at index 0. Suppose for contradiction that γ_0 meets $T(\delta, \eta')$ at some point, $y_0 \notin \delta$. By the above discussion, γ_0 crosses δ. By tightness, either γ_{-1} or γ_1 also crosses δ. Let $\gamma_{-s}, \ldots, \gamma_0, \ldots, \gamma_t$ be a maximal sequence of consecutive multicurves, all of which cross δ. Thus, $s + t \geq 1$. We claim that either $s = q$ 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, $\gamma_r \cap T(\delta, \eta) \subseteq \delta$, we see that, in fact, $\gamma_r \cap T(\delta, \eta) = \emptyset$.

Now let $\phi : S, \alpha_0 \sqcup \alpha_1 \longrightarrow M, \gamma_0 \sqcup \gamma_1$ be a pleating surface for $\gamma_0 \sqcup \gamma_1$, and let $x \in \alpha_0$ map to y_0. There is a loop, ϵ, through x in S, of length at most c, with $\phi\epsilon$ in the free homotopy class of δ. By assumption, γ_1 crosses δ and hence α_1 also crosses ϵ. Thus $\gamma_1 \cap \phi\epsilon \neq \emptyset$, and so, in particular (since ϕ is 1-lipschitz) $d(y_0, \gamma_1) \leq c$. Thus, γ_1 meets $N(T(\delta, \eta'), c)$ at some point, y_1. Note that, since $\eta' e^c < \eta e^{-l}$, we have $y_1 \in T(\delta, \eta e^{-l})$.

We can continue this process inductively to show that γ_i meets $N(T(\delta, \eta'), ic)$ in some y_i, provided $i \leq 2R + 1$, (so that $ic + l \leq (2R + 1)c + l$ and so $y_i \subseteq T(\delta, \eta e^{-l})$). We know that, $\gamma_r \cap T(\delta, \eta) = \emptyset$ and so it follows that $r \geq 2R + 1$.

By our earlier discussion of quasiprojections and our choice of R, we see that $d(\gamma_0, \delta) \leq R$ and $d(\gamma_{2R+1}, \delta) \leq R$. Thus, $d(\gamma_0, \gamma_{2R+1}) \leq 2R$, contradicting our assumption that $(\gamma_i)_i$ is geodesic. $\diamondsuit$

Lemma 5.1 now follows.

We note the following variation on Lemma 5.1:

Lemma 5.3 : *There is some $\eta > 0$ and $r_1 \in \mathbf{N}$, depending only on $\kappa(\Sigma)$ and k, with the following property. Suppose $\gamma_0, \ldots, \gamma_p$ is a multigeodesic in M, and let $r =$*

$\max\{d(\gamma_0, X(M,k)), d(\gamma_p, X(M,k))\}$. *Suppose that* $r + r_1 \leq i \leq p - r - r_1$ *and that* $(\gamma_i)_i$ *is tight at index* i. *If* $T = T(\delta, \eta)$ *is a Margulis tube, then* $\gamma_i \cap T \subseteq \delta$.

Proof : Since $X(M,k)$ is quasiconvex, a standard hyperbolicity argument shows that the sub-multigeodesic, $\gamma_r, \gamma_{r+1}, \ldots, \gamma_{p-r}$ lies a bounded distance from $X(M,k)$. This gives a constant $R \geq 0$ as in the proof of Lemma 5.2, and we set $r_1 = 2R + 1$. The proof now proceeds exactly as for Lemma 5.2. This time, we have not assumed that γ_r or γ_{p-r} meet T at most in δ. However this assumption was only used in ensuring that $i - r$ or $p - r - i$ was at least $2R + 1$, so as to give a contradiction. Here, this is ensured instead by our choice of r_1. $\diamondsuit$

These lemmas have obvious variants for accidental cusps. The proofs remain essentially unchanged. We interpret "Margulis tube" to mean "Margulis tube or accidental cusp". In the latter case, $T = T(\delta, \eta)$ becomes $P(\delta, \eta)$, where $\delta \in X_A(\Sigma)$, and the statement that $\gamma \cap T \subseteq \delta$ should be interpreted to that $\gamma \cap T$ is either empty or consists of a single curve in the homotopy class of T.

6. Hausdorff convergence.

In this section, we discuss Hausdorff convergence of sets, in particular, in relation to curves and laminations. One of the aims of the discussion is to be able to associate a subsurface of Σ to a lamination or a sufficiently close approximating curve. We can perform similar constructions in the surface or in the projectivised tangent bundle to the 3-manifold. A convenient way of formally describing the subsurface will be to consider the homotopy classes of closed curves lying in a small neighbourhood of the lamination, or approximating curve.

To make this more precise, recall from Section 3 the notion of Hausdorff distance on the set of compact subsets of a metric space.

Suppose (Q_1, ρ_1) and (Q_2, ρ_2) are metric spaces, and $g : [0, \infty) \longrightarrow [0, \infty)$ is a homeomorphism.

Definition : A function $\theta : Q_1 \longrightarrow Q_2$ is *g-continuous* if for all $x, y \in Q_1$, $\rho_2(\theta(x), \theta(y)) \leq g(\rho_1(x,y))$.
A bijective function, θ is *g-bicontinuous* if both θ and θ^{-1} are g-continuous (so that θ is a homeomorphism).

Let (K_1, ρ_1) and (K_2, ρ_2) be compact metric spaces and write $(K_1 \times K_2, \rho)$ for the product space with the sup-metric, i.e. $\rho((x_1, x_2), (y_1, y_2)) = \max\{\rho_1(x_1, y_1), \rho_2(x_2, y_2)\}$ for all $x_1, y_1 \in K_1$, $x_2, y_2 \in K_2$. Let $\pi_i : K_1 \times K_2 \longrightarrow K_i$ be the projection map.

We say that a closed subset, $G \subseteq K_1 \times K_2$ is a *g-graph* if for all $(x_1, x_2), (y_1, y_2) \in G$, we have: $\rho_2(x_2, y_2) \leq g(\rho_1(x_1, y_1))$ and $\rho_1(x_1, y_1) \leq g(\rho_2(x_2, y_2))$. Thus a g-graph is the graph of a g-bicontinuous function from $\pi_1(G)$ to $\pi_2(G)$. Note that the set of g-graphs is closed in the Hausdorff topology on $\mathcal{C}(K_1 \times K_2)$.

Lemma 6.1 : *Suppose $Q_1^n \subseteq K_1$ and $Q_2^n \subseteq K_2$ are sequences of closed subsets and for each n there is a g-bicontinuous homeomorphism θ^n from Q_1^n to Q_2^n. Then we can find subsets $Q_1 \subseteq K_1$ and $Q_2 \subseteq K_2$ such that $Q_1^{n_i} \to Q_1$ and $Q_2^{n_i} \to Q_2$ in the Hausdorff topologies on $\mathcal{C}(K_1)$ and $\mathcal{C}(K_2)$ respectively, and such that there is a g-bicontinuous homeomorphism, θ, from Q_1 to Q_2.* $\diamondsuit$

Proof : Let $G^n = \{(x, \theta^n(x)) \mid x \in K_1\}$ be the graph of the g-bicontinuous homeomorphism from $\pi_1(G^n) = Q_1^n$ to $\pi_2(G^n) = Q_2^n$. Thus G^n is a g-graph. We pass to a subsequence G^{n_i} which converges to some $G \subseteq K_1 \times K_2$ in the Hausdorff topology on $K_1 \times K_2$. Now G is a g-graph. Moreover, it is easily seen that $\pi_j(G^{n_i}) \to \pi_j(G)$ in the Hausdorff topology on K_j. Thus setting $Q_j = \pi_j(G)$, G is the graph of a g-bicontinuous homeomorphism from Q_1 to Q_2, and $Q_1^{n_i} \to Q_1$ and $Q_2^{n_i} \to Q_2$ as required. $\diamondsuit$

In fact, the following slight refinement of Lemma 6.1 follows by essentially the same argument.

Lemma 6.2 : *Suppose, with the hypotheses of Lemma 6.1, we are given, in addition closed subsets, $P^n \subseteq Q_1^n$. Then we can choose the subsequence n_i so that there is a closed subset, $P \subseteq Q_1$, with $P^n \to P$ and $\theta^n(P^n) \to \theta(P)$ in the Hausdorff topologies on $\mathcal{C}(K_1)$ and $\mathcal{C}(K_2)$ respectively.* $\diamondsuit$

We want to apply these results to laminations on a complete finite-area hyperbolic surface, S. We are only dealing with compact laminations. Any such lamination will lie inside a fixed compact subset of S (cutting off the cusps along horocycles of length 1). Moreover, the set of such laminations is closed, and hence compact, in the Hausdorff topology.

Given a lamination, $\lambda \subseteq S$, recall that $\nu(\lambda)$ is the sublamination consisting of the union of all minimal sublaminations.

We need the following observation:

Lemma 6.3 : *Suppose $\lambda, \mu \subseteq S$ are laminations and $\theta : \lambda \longrightarrow \mu$ is a homeomorphism. Then $\theta(\nu(\lambda)) = \nu(\mu)$.*

Proof : The property of being a sublamination of a given lamination is an intrinsic property of the lamination. Thus, the homeomorphic image of any minimal sublamination is a minimal sublamination. Moreover, $\nu(\lambda)$ and $\nu(\mu)$ have the same number of components, and the result follows. $\diamondsuit$

Any sequence $(\alpha^n)_n$ of multicurves has a subsequence $(\alpha^{n_i})_i$ converging to a lamination λ. We can also choose subsets $a^{n_i} \subseteq \alpha^{n_i}$ which converge on $\nu(\lambda)$. There is a sense in which such a choice of subset is natural:

Lemma 6.4 : *Suppose $(\alpha^n)_n$, $(\beta^n)_n$ are sequences of geodesic multicurves in (S, ρ). Suppose that for each n, there is a g-bicontinuous homeomorphism, θ^n, from α^n to β^n (with respect to the metric ρ). Thus we can find subsequences so that α^{n_i} and β^{n_i}*

converge respectively on laminations λ and μ. Moreover, we can find (possibly empty) closed subsets, $a^{n_i} \subseteq \alpha^{n_i}$ such that $a^{n_i} \to \nu(\lambda)$ and $\theta^{n_i} a^{n_i} \to \nu(\mu)$.

Proof : After passing to a subsequence, we can suppose that α^n converges to a lamination λ, and so we can choose $a^n \subseteq \alpha^n$ converging on $\nu(\lambda)$. Lemma 6.2 gives us a subsequence $(n_i)_i$, with $\alpha^{n_i} \to \lambda$, $\beta^{n_i} \to \mu$ and a g-bicontinuous homeomorphism, $\theta : \lambda \to \mu$ and that $\theta^n a^n \to \theta(\nu(\lambda))$. By Lemma 6.3, $\theta(\nu(\lambda)) = \nu(\mu)$. $\diamondsuit$

We remark that we can choose the sets a^{n_i} independently of the sequence β^n.

Let P be a riemannian manifold with path-metric $\rho = \rho_P$. Given $x \in P$, write $\mathrm{inj}_P(x)$ for the injectivity radius of P at x. Given $\epsilon > 0$, write $\mathrm{thick}_\epsilon(P) = \{x \in P \mid \mathrm{inj}_P(x) \geq \epsilon\}$. We write $Y(P)$ for the set of free homotopy classes of closed curves in P, which we refer to as *loops*. For $C \subseteq P$, we shall write $Y(P,C)$ for the image of $Y(C)$ in $Y(P)$. Given a closed subset, $Q \subseteq P$ and $t > 0$, a *t-chain* in Q is a finite sequence, $\underline{x}$, of points, $x_0, x_1, \ldots, x_p = x_0$ in Q with $\rho(x_i, x_{i+1}) \leq t$ for all i. If $Q \subseteq \mathrm{thick}_t(P)$, then this determines a closed curve, $\sigma(\underline{x})$ in P, by connecting each x_i to x_{i+1} by the unique shortest geodesic segment between them. We can think of this loop as an element of $Y(P)$. We write $Y_t(P,Q) \subseteq Y(P)$ for the set of loops in P of the form $\sigma(\underline{x})$ where $\underline{x}$ is a t-chain in Q. We can verify that $Y_t(P,Q) = Y(P, N(Q, t/2))$. In fact $Y_t(P,Q)$ carries an additional structure consisting of a collection of subsets of $Y_t(P,Q)$, namely the sets $Y(P,C)$ as C varies over the set of components of $N(Q,t/2)$. We refer to this as the *component structure* on $Y_t(P,Q)$.

Suppose $\epsilon > 0$ and $I \subseteq (0, \epsilon]$ is a subinterval.

Definition : A subset $Q \subseteq P$ is *I-stable* if for all $t, u \in I$, $Y_t(P,Q) = Y_u(P,Q)$ and the respective component structures are equal.

We write $Y_I(P,Q)$ for this subset. Clearly it would be enough that for $t < u \in I$, the inclusion of $N(Q,t/2)$ into $N(Q,u/2)$ be a homotopy equivalence.

We make the following observation:

Lemma 6.5 : *Suppose that $(Q^n)_{n=0}^\infty$ is a sequence of closed subsets of P converging to $Q \subseteq P$ in the Hausdorff topology. Suppose Q is I-stable, and that $J \subseteq \mathrm{int}\, I$ is a closed subinterval. Then Q^n is J-stable for all sufficiently large n. Moreover, $Y_J(P,Q^n) = Y_I(P,Q)$.*

Proof : Let $J = [t,u]$ and let $t', u' \in I$ with $t' < t$ and $u < u'$. For all sufficiently large n, we have $N(Q,t'/2) \subseteq N(Q^n,t/2) \subseteq N(Q^n,u/2) \subseteq N(Q,u'/2)$, and it follows easily that $Y_t(P,Q^n) = Y_u(P,Q^n)$, respecting component structures. $\diamondsuit$

Again, we will want to apply this to laminations. We note:

Lemma 6.6 : *Suppose that S is a complete finite area hyperbolic surface, and $\lambda \subseteq S$ is a lamination. There is some $\tau(\lambda) > 0$ such that for all t, u with $0 < t \leq u \leq \tau(\lambda)$, the inclusion of $N(\lambda, t/2)$ into $N(\lambda, u/2)$ is a homotopy equivalence.*

Proof : Take $\tau(\lambda)$ less that the injectivity radius outside the cusps of S, and less than the length of any arc in $S \setminus \lambda$ that connects two non-asymptotic components. (Note that we only need to consider finitely many complementary components of λ, each of which has only finitely many boundary leaves, and hence only finitely many pairs on non-asymptotic leaves.) $\diamondsuit$

In particular, we see that λ is $(0, \tau(\lambda)]$-stable.

Note that if $t \in (0, \tau(\lambda)]$, then $Y_t(S, \lambda) = Y(S, N(\lambda, t/2))$, and from the discussion of laminations is Section 3, we see that this equals $Y(S, \Phi(\lambda))$, where $\Phi(\lambda)$ is the efficient subsurface of S filled by λ. We see (restricting to simple curves) that $X(S) \cap Y_t(S, \lambda) = X(S, \Phi(\lambda))$.

The following set-up will feature in Section 7.

Suppose P, P_1, P_2 are riemannian manifolds, and that $\zeta_i : P_i \longrightarrow P$ is a map inducing an isomorphism of fundamental groups, hence a bijection between $Y(P_i)$ and $Y(P)$. In this way, we may identify the sets $Y(P_1)$ and $Y(P_2)$. Suppose that $h : [0, \infty) \longrightarrow [0, \infty)$ is a homeomorphism, and that ζ_i is h-continuous on $\zeta_i^{-1}(\mathrm{thick}_\epsilon(P))$. Suppose there exists ϵ' such that $\mathrm{thick}_{\epsilon'}(P_i) \subseteq \zeta_i^{-1}(\mathrm{thick}_\epsilon(P))$. We can take $\epsilon' \leq h^{-1}(\epsilon)$. Suppose also that $Q_i \subseteq \mathrm{thick}_{\epsilon'}(P_i)$, and that $\theta : Q_1 \longrightarrow Q_2$ is a g-bicontinuous homeomorphism with $\zeta_2 \circ \theta = \zeta_1$ on Q_1, where $g : [0, \infty) \longrightarrow [0, \infty)$ is a homeomorphism.

Now any t-chain, $\underline{x} = (x_i)_i$, in Q_1 gives rise to a $g(t)$-chain, $\theta(\underline{x}) = (\theta(x_i))_i$ in Q_2. If $g(t) < \epsilon'$, then we have loops $\sigma(\underline{x})$ and $\sigma(\theta(\underline{x}))$ in P_1 and P_2 respectively. These project to loops $\zeta_1(\theta(\underline{x}))$ and $\zeta_2(\sigma(\theta(\underline{x})))$ in P. Since $\zeta_2 \circ \theta = \zeta_1$ on Q_1, these loops agree on the $h(t)$-chain $\zeta_1(\underline{x}) = \zeta_2(\theta(\underline{x}))$. Since $h(t) \leq h(\epsilon') \leq \epsilon$, and this chain lies in $\mathrm{thick}_\epsilon(P)$, we see that $\zeta_1(\sigma(\underline{x}))$ and $\zeta_2(\sigma(\theta(\underline{x})))$ must be homotopic in P. Thus, $\underline{x}$ and $\theta(\underline{x})$ give rise to the same element of $Y(P_1) \equiv Y(P_2)$ under the identification described above.

Note that the same argument applies to any t-chain, $\underline{y}$ in Q_2. This gives rise to a $g(t)$-chain, $\theta^{-1}(\underline{y})$ in Q_1.

Now suppose that Q_1 is J-stable in P_1, where $J = [r, s]$ is a closed interval. Suppose that $g^2(r) < s < \epsilon'$.

Lemma 6.7 : *If $g(r) \leq u \leq g^{-1}(s)$, then $Y_u(P_2, Q_2) = Y_J(P_1, Q_1)$, respecting the component structures.*

Proof : Each element of $Y_J(P_1, Q_1)$ is represented by a r-chain, $\underline{x}$, in Q_1. This gives us a $g(r)$-chain $\theta(\underline{x})$ in Q_2, representing the same element under the identification of $Y(P_1)$ with $Y(P_2)$. Similarly, any u-chain, $\underline{y}$, in Q_2 gives us a $g(u)$-chain $\theta^{-1}(\underline{y})$ representing the same element. These identifications are easily seen to respect the component structures. $\diamondsuit$

We shall apply these results in the case where $P = M$ is a hyperbolic 3-manifold; $P_1 = S$ is a hyperbolic surface; $\alpha = Q_1 \subseteq S$ is a geodesic multicurve; $P_2 = E$ is the projectivised unit tangent bundle of M; $\zeta_2 : E \longrightarrow M$ is the projection map; $Q_2 \subseteq E$ is the lift of a geodesic multicurve, γ, in M; $\zeta_1 = \phi : S, \alpha \longrightarrow M, \gamma$ is a pleating surface for γ; and $\theta = \psi$ is the lift of $\phi|\alpha$. Since ζ_1 is by hypothesis uniformly lipschitz, and ζ_2 is 1-lipschitz, we can take h to have the form $[t \mapsto \mu t]$ for some fixed $\mu > 0$. The existence of

$\epsilon, \epsilon' > 0$ follow from the discussion of pleating surfaces, in Section 4. The existence of the map g is given by the uniform injectivity theorem.

7. Construction of subsurfaces.

In this section, we use the constructions of Section 6 to show how sequences of multicurves in a 3-manifold give rise to subsurfaces in Σ.

We fix Margulis constants as described in Section 4. Suppose that M is a complete hyperbolic 3-manifold with a homotopy equivalence $\chi : M \longrightarrow \Sigma$ as before. Let $\mathcal{T}$ be the set of Margulis tubes and let $\Theta(M) = \Psi(M) \setminus \bigcup_{C \in \mathcal{T} \cup \mathcal{P}_A} \operatorname{int} C$, (so that $\operatorname{thick}_\eta(M) \subseteq \Theta(M) \subseteq \operatorname{thick}_{\eta'}(M)$ for certain constants, $\eta, \eta' > 0$).

We say that a geodesic multicurve, γ, in M is *non-penetrating* if, for all $T \in \mathcal{T} \cup \mathcal{P}_A$, $\gamma \cap T$ is either empty or the core curve of T. In this situation, it will be convenient to modify our definition of a pleating surface. To this end, we fix once and for all, a finite area hyperbolic structure, (S, ρ_S), on a surface homeomorphic to $\operatorname{int} \Sigma$.

Modified definition : Let $\gamma \subseteq M$ be a non-penetrating multicurve. A *pleating surface* for γ consists of a geodesic multicurve, $\alpha \subseteq S$ and a homotopy equivalence, $\phi : S \longrightarrow M$, such that ϕ maps α onto γ by a local homeomorphism, with the following properties. The map $\phi|(S \cap \phi^{-1}(\Theta(M)))$ is uniformly lipschitz, the map $\phi|(\alpha \cap \phi^{-1}(\Theta(M)))$ is locally uniformly bilipschitz.

With some modification, we can also assume that if $T \in \mathcal{T} \cup \mathcal{P}_A$, then $\phi^{-1}T$ is either empty or is an annulus in S, though we won't explicitly need this. The various constants depend only on $\kappa(\Sigma)$, the Margulis constants, and our particular choice of (S, ρ_S).

Lemma 7.1 : *Let $\gamma \subseteq M$ be a non-penetrating geodesic multicurve in M. Then there is a pleating surface, $\phi : S, \alpha \longrightarrow M, \gamma$ for γ, in this modified sense.*

Proof : We start with a pleating surface, $S', \alpha \longrightarrow M, \gamma$, in the previous sense. The injectivity radius on $S' \cap \phi^{-1}\Theta(M)$ is uniformly bounded below. Thus, the non-trivial components of $S' \cap \phi^{-1}\Theta(M)$ remain in a bounded region of the moduli space (after straightening the boundaries). We can thus find a homeomorphism of S to S' for which the metrics are uniformly bilipschitz related on $S' \cap \phi^{-1}\Theta(M)$. Composing with this homeomorphism gives us a pleated surface in the modified sense. $\diamondsuit$

We also note that the Uniform Injectivity Theorem (Theorem 4.1) remains true in this context:

Lemma 7.2 : *Suppose $\gamma \subseteq M$ is a non-penetrating curve and $\phi : S, \alpha \longrightarrow M, \gamma$ is a pleating surface in the above sense. Let $\psi|\alpha \longrightarrow \gamma \subseteq E$ be the lifted homeomorphism to $\gamma \subseteq E$, where E is the projectivised unit tangent bundle to M. We put the metric ρ_S on α. Then $\psi|(\alpha \cap \phi^{-1}(\Theta(M)))$ is g-bicontinuous, where $g : [0, \infty) \longrightarrow [0, \infty)$ depends* *only on $\kappa(\Sigma)$ and the Margulis constants (and the constants featuring in the definition of pleating surface).* $\diamondsuit$

We want to apply these constructions to sequences of multicurves in 3-manifolds.

Suppose that $(M^n)_n$ is an infinite sequence of hyperbolic 3-manifolds, each with its own homotopy equivalence, $\chi^n : M^n \longrightarrow \Sigma$.

Suppose that, for each n, we have a non-penetrating geodesic multicurve, γ^n, in M^n. After passing to a subsequence, we can suppose that $\gamma^n = \check{\gamma}^n \sqcup \hat{\gamma}^n$, where the total length of the $\check{\gamma}^n$ remains bounded, but where the minimal length of the components of $\hat{\gamma}^n$ tend to ∞. We can suppose that $\hat{\gamma}^n \subseteq \Theta(M^n)$ for all n.

Now let $\phi^n : S, \alpha^n \longrightarrow M^n, \gamma^n$ be pleating surfaces for γ^n (in the above sense). Let $\omega^n = \chi^n \circ \phi^n$ be the induced homotopy equivalence from S to Σ. We can decompose α^n as $\alpha^n = \check{\alpha}^n \sqcup \hat{\alpha}^n$, where $\check{\gamma}^n = \phi^n \check{\alpha}^n$ and $\hat{\gamma}^n = \phi^n \hat{\alpha}^n$. Thus the length of $\check{\alpha}^n$ remains bounded, and the minimal lengths of the components of $\hat{\alpha}^n$ tend to ∞. As in Section 3, after passing to a further subsequence, we can suppose that $\check{\alpha}^n = l$ is a fixed (possibly empty) multicurve, and that $\hat{\alpha}^n$ tends to a lamination, $\lambda \subseteq S$. We thus obtain subsurfaces $F(\underline{\alpha}) = \Phi(\nu(\lambda))$, $G(\underline{\alpha}) = \Phi(\lambda)$ and $H(\underline{\alpha}) = \Phi(l)$ in S. Let $F^n = \omega^n F(\underline{\alpha})$, $G^n = \omega^n G(\underline{\alpha})$, and $H^n = \omega^n H(\underline{\alpha})$. These are all efficient subsurfaces of Σ with $F^n \subseteq G^n$ and $G^n \cap H^n = \emptyset$.

For all sufficiently large n, we would like to have a means of recognising the surfaces F^n, G^n, H^n directly in terms of the geodesics, γ^n, in M^n and the maps χ^n, i.e. without explicit reference to pleating surfaces. Intuitively, the idea is that the lifts of γ^n to the projectivised unit tangent bundles E^n "fill up" subsets which project to these subsurfaces under the maps χ^n. To make this precise, we use the machinery of Section 6.

Firstly, we make the observation that $\omega^n l = \omega^n \check{\alpha}^n = \chi^n \check{\gamma}^n$, and so $H^n = H(\chi^n \check{\gamma}^n)$ is just a regular neighbourhood of the multicurve $\chi^n \check{\gamma}^n$.

Next, we want to recognise the surfaces G^n.

Given $t > 0$, we have the set, $Y_t(E^n, \hat{\gamma}^n) \subseteq Y(E^n)$. Denoting the composition of the projection of E^n to M with the homotopy equivalence from M^n to Σ also by χ^n, we see that χ^n induces a map between $Y(E^n)$ and $Y(\Sigma)$. This map is two-to-one (since the fibres of E^n are projective planes), so we should quotient out $Y(E^n)$ by an involution that identifies the non-trivial curve in the projective plane to a point. In future references to $Y(E^n)$ we assume that we have carried out this identification. We thus get a subset $\chi^n Y_t(E^n, \hat{\gamma}^n) \subseteq Y(\Sigma)$. We set $X_t(E^n, \gamma^n) = X(\Sigma) \cap \chi^n Y_t(E^n, \hat{\gamma}^n)$. Note that this comes equipped with a component structure (arising from the components of $N(\hat{\gamma}^n, t/2)$ in E^n).

Lemma 7.3 : *There is some $\tau > 0$ such that for all $t \in (0, \tau]$, for all sufficiently large n (depending of t) and for all $u \in [t, \tau]$, we have $X(\Sigma, G^n) = X_u(E^n, \hat{\gamma}^n)$, respecting component structures.*

Recall that $X(\Sigma, G^n)$ is the set of curves in $X(\Sigma)$ that are homotopic into G^n. Together with its component structure, this determines G^n as an efficient subsurface (up to homotopy).

Proof : Let ψ^n be the homeomorphism from α^n to $\gamma^n \subseteq E^n$ that projects to $\phi^n | \alpha^n$. By Lemma 7.2, $\psi^n | \hat{\alpha}^n$ is g-bicontinuous with respect to the metric ρ_S on $\hat{\alpha}^n$ and $\rho_n = \rho_{E^n}$

on $\hat{\gamma}^n$.

Let $\tau(\lambda)$ be the constant given by Lemma 6.6, so that the lamination, λ, is $(0, \tau(\lambda)]$-stable. We can suppose that $\tau(\lambda)$ is less the Margulis constant, ϵ', featuring in the hypotheses of Lemma 6.7, as discussed at the end of Section 6. Fix any $\tau \in (0, g^{-1}(\tau(\lambda)))$, and suppose that $t \in (0, \tau]$. Let J be the interval $[g^{-1}(t), g(\tau)]$. Now $J \subseteq (0, \tau(\lambda))$ and so, by Lemma 6.5, $\hat{\alpha}^n$ is J-stable for all sufficiently large n. Indeed, $Y_J(S, \hat{\alpha}^n) = Y_I(S, \lambda) = Y(S, \Phi(\lambda)) = Y(S, G(\underline{\alpha}))$, where $I = (0, \tau(\lambda)]$.

We are now in the set-up described at the end of Section 6. We set $P = M^n$; $P_1 = S$; $\zeta_1 = \phi^n$; $Q_1 = \hat{\alpha}^n$; $P_2 = E^n$; $Q_2 = \hat{\gamma}^n \subseteq E^n$; ζ_2 is the projection of E^n to M^n; and $\theta = \psi^n|\hat{\alpha}^n$. Thus, θ is g-bicontinuous, ζ_1 is uniformly lipschitz and ζ_2 is 1-lipschitz. Note that $Y(S)$ is identified with $Y(E^n)$ via the map ϕ^n and projection of E^n to M^n.

Suppose $u \in [t, \tau] = [g(g^{-1}(t)), g^{-1}(g(\tau))]$. Since $\hat{\alpha}^n$ is $[g^{-1}(t), g(\tau)]$-stable, by Lemma 6.7 we see that $Y_u(E^n, \hat{\gamma}^n) = \phi^n Y_J(S, \hat{\alpha}^n) = \phi^n Y(S, G(\underline{\alpha}))$. Thus, $\chi^n Y_u(E^n, \hat{\gamma}^n) = \chi^n \phi^n Y(S, G(\underline{\alpha})) = \omega^n Y(S, G(\underline{\alpha})) = Y(\Sigma, \omega^n G(\underline{\alpha})) = Y(\Sigma, G^n)$. Thus $X_u(E^n, \hat{\gamma}^n) = X(\Sigma) \cap \chi^n Y_u(E^n, \hat{\gamma}^n) = X(\Sigma) \cap Y(\Sigma, G^n) = X(\Sigma, G^n)$. ♢

Next, we want to recognise the surfaces F^n in terms of γ^n and M^n.

Recall that $\hat{\alpha}^n \to \lambda$ in S, and so we can find (possibly empty) subsets $a^n \subseteq \hat{\alpha}^n$, with $a^n \to \nu(\lambda)$. Let $c^n = \psi^n a^n \subseteq \gamma^n \subseteq E^n$. We now apply the argument of Lemma 7.3, with a^n replacing α^n and c^n replacing γ^n, to show:

Lemma 7.4 : *There is some $\tau' > 0$ such that for all $t \in (0, \tau']$, for all sufficiently large n (depending on t) and all $u \in [t, \tau']$, $X(\Sigma, F^n) = X_u(E^n, c^n)$, respecting the component structure.*

Here of course, $X_u(E^n, c^n)$ denotes $X(\Sigma) \cap \chi^n Y_u(E^n, c^n) \subseteq X(\Sigma)$. This set, together with its component structure is sufficient to determine F^n.

In this, however, we have cheated in that c^n might depend on the choice of pleating surfaces ϕ^n. This can be fixed by the following observation.

Lemma 7.5 : *After passing to a subsequence of M^n, γ^n, the sets c^n featuring in Lemma 7.4 can be chosen independently of the sequence of pleating surfaces.*

Proof : Suppose we have another sequence, $\bar{\phi}^n : S, \bar{\alpha}^n \longrightarrow M^n, \hat{\gamma}^n$ for $\hat{\gamma}^n$. Let $\bar{\psi}^n : \bar{\alpha}^n \longrightarrow \hat{\gamma}^n \subseteq E^n$ the corresponding lift of of $\bar{\phi}^n|\bar{\alpha}^n$. Let $\bar{a}^n = (\bar{\psi}^n)^{-1} c^n \subseteq \bar{\alpha}^n$. Let $\theta^n = (\bar{\psi}^n)^{-1} \circ \psi^n : \hat{\alpha}^n \longrightarrow \bar{\alpha}^n$. This is a g^2-bicontinuous homeomorphism. Applying Lemma 6.4, we can pass to a subsequence so that $\bar{\alpha}^n$ converges to some lamination $\mu \subseteq S$ and so that $\theta^n a^n \to \nu(\mu)$. But $\theta^n a^n = \bar{a}^n$ and $\bar{\psi}^n a^n = c^n$. Thus the same sets $c^n \subseteq \gamma^n$ will serve for both sequences of pleating surfaces, $(\psi^n)_n$ and $(\bar{\psi}^n)_n$.

Of course, this is somewhat weaker than the claim we made since we are only considering two non-penetrating sequence of pleated surfaces. To complete that argument, we would need to observe that the sequence a^n (and so c^n) can be chosen before knowing $\bar{\psi}^n$ (cf. the observation made after Lemma 6.4), and that the same subsequence will serve for all sequences $(\bar{\psi}^n)_n$. We shall not however to go into the details of this, since we shall only ever deal with at most two sequences of pleated surfaces at a time. ♢

Next, consider the situation where we have two sequences, $(\gamma_0^n)_n$ and $(\gamma_1^n)_n$ in M^n, so that for all n, γ_0^n and γ_1^n are compatible, i.e. $\chi^n(\gamma_0^n)$ and $\chi^n(\gamma_1^n)$ do not cross, so that $\chi^n(\gamma_0^n \cup \gamma_1^n)$ is a multicurve in Σ.

We have pleating surfaces, $\phi^n : S, \alpha^n \cup \beta^n \to M, \gamma_0^n \cup \gamma_1^n$, with $\gamma_0^n = \phi^n \alpha^n$ and $\gamma_1^n = \phi^n \beta^n$. Passing to a subsequence, we have $\alpha^n \to l \sqcup \lambda$, $\beta^n \to m \sqcup \mu$ and $\alpha^n \cup \beta^n \to (l \cup m) \sqcup (\lambda \sqcup \mu)$ as in Section 2. We thus get efficient surfaces, $F(\underline{\alpha})$, $G(\underline{\alpha})$, $H(\underline{\alpha})$, $F(\underline{\beta})$, $G(\underline{\beta})$, $H(\underline{\beta})$ in S. We set $F_0^n = \omega^n F(\underline{\alpha})$, $G_0^n = \omega^n G(\underline{\alpha})$, $H_0^n = \omega^n H(\underline{\alpha})$, $F_1^n = \omega^n F(\underline{\beta})$, $G_1^n = \omega^n G(\underline{\beta})$, $H_1^n = \omega^n H(\underline{\beta})$. These are all efficient surfaces in Σ. Moreover, as in Section 2, we see that they satisfy all of the properties (F1)–(F8). with $\{i,j\} = \{0,1\}$ and with F_i^n replacing F_i etc., and with $\chi^n(\gamma_i^n)$ replacing α^n.

Moreover, we have sequences $c_0^n \subseteq \gamma_0^n$ and $c_1^n \subseteq \gamma_1^n$ of $\gamma_0^n, \gamma_1^n \subseteq E^n$, such that for $i \in \{0,1\}$ we have $H_i^n = \Phi(\chi^n \check{\gamma}_i^n)$ and there is some $\tau > 0$ such that for all $t \in (0,\tau]$, for all sufficiently large n, for all $u \in [t,\tau]$, $X(\Sigma, G_i^n) = X_u(E^n, \hat{\gamma}_i^n)$ and $X(\Sigma, F_i^n) = X_u(E^n, c_i^n)$.

We can now apply the same argument to a sequence of paths of multicurves.

Suppose for each n, we have a sequence $(\gamma_i^n)_{i=0}^p$ of multicurves in M^n, so that $\chi^n \gamma_i^n$ and $\chi^n \gamma_{i+1}^n$ are compatible for all n and for all $i = 0, \ldots, p-1$. For each $i \in \{0, \ldots, p-1\}$ we choose a sequence $\phi_i^n : S, \alpha_i^n \cup \beta_i^n \longrightarrow M, \gamma_i^n \cup \gamma_{i+1}^n$ of pleating surfaces for $\gamma_i^n \cup \gamma_{i+1}^n$. Using induction over i, we pass to successive subsequences, so as to give us surfaces, F_i^n, G_i^n, H_i^n in Σ as above. In fact, we can find some $u > 0$ so that $X(\Sigma, F_i^n) = X_u(E^n, c_i^n)$ and $X(\Sigma, c_i^n) = X_u(E^n, \hat{\gamma}_i^n)$ for all sufficiently large n.

Note that, by Lemma 7.5, the same sequence $c_i^n \subseteq \hat{\gamma}_i^n \subseteq E^n$ serves for both pleated surfaces $\phi_i^n : S, \hat{\alpha}_i^n \longrightarrow M^n, \gamma_i^n$ and $\phi_{i-1}^n : S, \hat{\beta}_{i-1}^n \longrightarrow M^n, \gamma_i^n$. This means that the surfaces $F_i^n = \omega_i^n F(\underline{\alpha}_i) = \omega_{i-1}^n F(\underline{\beta}_{i-1})$ are well defined.

The earlier discussion shows that for consecutive indices, i, j, the surfaces satisfy all the conditions (F1)–(F8) of Section 3. Moreover, property (F9) was deduced directly from (F1)–(F8) and so also holds.

We have shown:

Lemma 7.6 : *Suppose we have a sequence, M^n, of complete hyperbolic 3-manifolds, each admitting a strictly type preserving homotopy equivalence $\chi^n : M^n \longrightarrow \Sigma$. Suppose, for each n, we have a sequence $(\gamma_i^n)_{i=0}^n$ of geodesic multicurves in M^n with γ_i^n and γ_{i+1}^n compatible for all i and n. Suppose that γ_i^n is non-penetrating for all i and n with respect to fixed Margulis constants. Then, for an infinite subsequence of n, we can construct efficient subsurfaces F_i^n, G_i^n, H_i^n of Σ satisfying properties (F1)–(F9) of Section 3, where F_i, G_i, H_i are replaced by F_i^n, G_i^n, H_i^n and where $\alpha_i^n = \check{\alpha}_i^n \sqcup \hat{\alpha}_i^n$ are replaced by $\gamma_i^n = \check{\gamma}_i^n \sqcup \hat{\gamma}_i^n$.* $\diamondsuit$

8. Proofs of the main results.

In this section, we prove the main results stated in Section 1, namely Theorems 1.3, 1.4 and 1.5.

The key lemma will be the following. Suppose M is a hyperbolic 3-manifold with a homotopy equivalence to Σ. We can talk about a "(tight) multigeodesic" in M, by which we mean a sequence of geodesic multicurves in M realising a (tight) multigeodesic in the curve graph of Σ. We write $L(M, \gamma_i)$ for the maximal length of the components of γ_i in M.

We fix a set, $\mathcal{T}$, of Margulis tubes for M. In the case where there are accidental cusps, we interpret this to include the accidental cusps as well.

Lemma 8.1 : *Given $k \geq 0$ and $p \in \mathbf{N}$, there is some $l \geq 0$ such that if $(\gamma_i)_{i=0}^p$ is a tight multigeodesic in M with γ_0 and γ_p non-penetrating with respect to $\mathcal{T}$, and $p \geq 12r + 19$, where $r = \max\{d(\gamma_0, X(M,k)), d(\gamma_p, X(M,k))\}$, then there is some $i \in \{0, \ldots, p\}$ such that $L(M, \gamma_i) \leq l$. Here, l depends on $k, p, \kappa(\Sigma)$ and the constants defining $\mathcal{T}$.*

Proof : Suppose the conclusion fails.

In that case, we can find a sequence, $(M^n)_n$ of 3-manifolds with homotopy equivalences, $\chi^n : M^n \longrightarrow \Sigma$, and tight multigeodesics, $(\gamma_i^n)_{i=0}^p$ in M^n, with p fixed, so that γ_0^n and γ_p^n are non-penetrating, $p \geq 12r_n + 19$, where $r_n = \max\{d(\gamma_0^n, X(M,k)), d(\gamma_p^n, X(M,k))\}$, and with $L(M^n, \gamma_i^n) \to \infty$ as $n \to \infty$ for each i. Here "non-penetrating" refers to sets, $\mathcal{T}^n$, of Margulis tubes and accidental cusps in M^n with fixed constants (independent of n). Passing to a subsequence, we may as well assume (to simplify notation) that $s = d(\gamma_0^n, X(M,k)) + 1$ and $t = d(\gamma_p^n, X(M,k)) + 1$ are fixed. Note that $p \geq 12r_n + 19 \geq 6(s+t-2) + 19 = 6(s+t) + 7$.

Some component, γ_{-1}^n, of γ_0^n satisfies $d(\gamma_{-1}^n, X(M^n, k)) = s - 1$. (Recall that if γ is a multicurve and $Q \subseteq X(\Sigma)$, $d(\gamma, Q)$ is defined as $d(X(\Sigma, \gamma), Q)$.) By assumption, γ_0^n, and hence γ_{-1}^n is non-penetrating with respect to $\mathcal{T}^n$.

Let $\gamma_{-1}^n, \ldots, \gamma_{-s}^n$ be a tight multigeodesic connecting γ_{-1}^n to some curve $\gamma_{-s}^n \in X(M^n, k)$. Now γ_{-s}^n is non-penetrating with respect to some other set of Margulis tubes depending on k. Thus, by Lemma 5.2, we can find yet another set of tubes, $\mathcal{T}_0^n$, with respect to which the multicurves γ_{-i}^n are non-penetrating for all $1 \leq i \leq s$. Since s is bounded in terms of p, the constants of $\mathcal{T}_0^n$ depend only on those of $\mathcal{T}^n$ and on k, p and $\kappa(\Sigma)$. We can assume that each tube of $\mathcal{T}_0^n$ is contained in one of $\mathcal{T}^n$. (Though, of course, not every tube of $\mathcal{T}^n$ need contain one $\mathcal{T}_0^n$.)

We get a similar multigeodesic, $\gamma_{p+1}^n, \ldots, \gamma_{p+t}^n$, with $\gamma_{p+1}^n \subseteq \gamma_p^n$, with $L(M^n, \gamma_{p+t}^n) \leq k$ and with all curves non-penetrating with respect to $\mathcal{T}_0^n$. Lemma 5.2 also tells us that $\gamma_0^n, \ldots, \gamma_p^n$ is non-penetrating with respect to $\mathcal{T}_0^n$.

Thus, for each n, we have a path $(\gamma_i^n)_{i=-s}^{p+t}$ of multicurves in M^n. It is non-penetrating with respect to $\mathcal{T}_0^n$ for all i, and tight for all $i \in \{1, \ldots, p-1\}$. In the notation of Section 7, $\hat{\gamma}_i^n \neq \emptyset$ for all $i \in \{0, \ldots, p\}$. We are now in the situation described in Lemma 7.6. We thus obtain a sequence of efficient subsurfaces, F_i^n, G_i^n, H_i^n of Σ, satisfying (F1)–(F9). In particular, (by (F6)), $F_i^n \neq \emptyset$ for all $i \in \{0, \ldots, p\}$, and (by (F9)) the sequence $(F_i^n)_i$ is taut for all $i \in \{1, \ldots, p-1\}$.

Applying Corollary 2.2, we see that $d(F_0^n, F_p^n) \leq \left[\frac{2}{3}(p+1)\right] + 1 < \frac{2}{3}p + 2$. In other words, there are curves $\alpha \subseteq F_0^n$ and $\beta \subseteq F_p^n$ with $d(\alpha, \beta) < \frac{2}{3}p + 2$.

Now $F_0^n \subseteq G_0^n$ and the surfaces $G_{-i}^n \sqcup H_{-i}^n$ and $G_{-i-1}^n \cup H_{-i-1}^n$ are compatible for all $i \in \{0, \ldots, s-1\}$. Moreover $\gamma_{-s}^n \in X(M^n, k)$, so $\hat{\gamma}_{-s}^n = \emptyset$, and so (by (F5)) $G_{-s}^n = \emptyset$. Moreover, by (F6), $H_{-s}^n = \Phi(\gamma_{-s}^n)$. It follows that $d(\alpha, \gamma_{-s}^n) \leq s$. Since $d(\gamma_0^n, \gamma_{-s}^n) < s$, we get that $d(\alpha, \gamma_0^n) < 2s$.

Similarly, we see that $d(\beta, \gamma_p^n) < 2t$. Thus $d(\gamma_0^n, \gamma_p^n) < \frac{2}{3}p + 2 + 2(s+t)$.

But $(\gamma_i)_{i=0}^p$ is a multigeodesic, and so γ_0^n and γ_p^n are exactly distance p apart. Thus, $p < \frac{2}{3}p + 2 + 2(s+t)$, giving $p < 6(s+t) + 6$.

This contradicts the earlier assertion that $p \geq 6(s+t) + 7$. $\diamondsuit$

We will also need the following "interpolation" lemma.

Lemma 8.2 : *Given $l \geq 0$ and $q, \kappa(\Sigma) \in \mathbf{N}$, there is some $K \geq 0$ such that if M is a hyperbolic 3-manifold with a homotopy equivalence to Σ, and if $(\gamma_i)_{i=0}^p$ is a tight multigeodesic with $p \leq q$, and $L(M, \gamma_0) \leq l$ and $L(M, \gamma_p) \leq l$, then for all $i \in \{0, \ldots, p\}$, we have $L(M, \gamma_i) \leq K$.*

Proof : Suppose the conclusion fails.

We find a sequence of manifolds, M^n, each with a tight multigeodesic $(\gamma_i^n)_{i=0}^p$ for some fixed $p \leq q$, with $L(M^n, \gamma_0^n) \leq l$ and $L(M^n, \gamma_p^n) \leq l$, and with $L(M^n, \gamma_i^n) \to \infty$ for at least one i. Now Lemma 5.1 tells us that each γ_i^n is non-penetrating with respect to a set of Margulis tubes with fixed constants (depending only on l and $\kappa(\Sigma)$). We now apply Lemma 7.6 to give us, for some n, a sequence of surfaces, F_i^n, G_i^n, H_i^n satisfying (F1)–(F9). By (F6), at least one of the F_i^n is non-empty. Moreover, $G_0^n = G_p^n = \emptyset$ and $H_0^n = \Phi(\gamma_0^n)$ and $H_p^n = \Phi(\gamma_p^n)$. We now derive a contradiction exactly as in the proof of Theorem 1.1 given in Section 3. $\diamondsuit$

Proof of Theorem 1.3 : From the definition of a tight geodesic, we can prove the statement for a tight multigeodesic, $(\gamma_i)_{i=0}^p$. By assumption, $L(M, \gamma_0) \leq k$ and $L(M, \gamma_p) \leq k$. By Lemma 5.1, the sequence $(\gamma_i)_i$ is non-penetrating with respect to fixed Margulis constants, depending on k and $\kappa(\Sigma)$. We can assume that $k \geq k_0$, so that by the quasiconvexity of $X(M, k)$ as discussed in Section 4, there is some R depending only on k and $\kappa(\Sigma)$ such that $d(\gamma_i, X(M, k)) \leq R$ for all i. Let $D = 12R + 19$. By Lemma 8.1 there is some $l \geq k$ such that any sequence of consecutive $(\gamma_i)_i$ of length D must contain an element γ_i with $L(M^n, \gamma_i) \leq l$. In other words, we can find a subsequence of indices, $0 = i(0) < i(1) < \cdots < i(q) = p$, with $i(j+1) \leq i(j) + D$ for all $j = 0, \ldots, q-1$, and with $L(M, \gamma_{i(j)}) \leq l$ for all j.

Now interpolating using Lemma 8.2, we find some K, depending on l, D and $\kappa(\Sigma)$, so that $L(M, \gamma_i) \leq K$ for all $i \in \{0, \ldots, p\}$. Thus K depends ultimately only on k and $\kappa(\Sigma)$. $\diamondsuit$

Proof of Theorem 1.4 : This is just a variation of the proof of Theorem 1.3.

First, note that by the hyperbolicity of the curve graph and the quasiconvexity of $X(M, k)$, the multicurves γ_i are at most some bounded distance from $X(M, k)$ for all i between r and $p - r$, where this bound, R, depends only on k and $\kappa(\Sigma)$. Lemma 5.3 now gives us some $r_1 \geq 0$ so that for all i between $r + r_1$ and $p - r - r_1$ the multicurves γ_i are non-penetrating. Again r_1 and the Margulis constants depend only on k and $\kappa(\Sigma)$. Let $D = 12R + 19$. We now apply Lemma 7.6 as before to give us a subsequence of indices, $r + r_1 \leq i(0) < i(1) < \cdots < i(q) \leq p - r - r_1$ with $i(0) \leq r + r_1 + D$, $p - r - r_1 \leq i(q) + D$, and $i(j+1) \leq i(j) + D$ for all i, such that $L(M^n, \gamma_{i(j)})$ is bounded for all j. We now use Lemma 8.2 to interpolate between the indices $i(j)$. Setting $r_0 = r_1 + D$, we see that $L(M^n, \gamma_i)$ is bounded for all i between $r + r_0$ and $p - r - r_0$ as required. ♢

One can also give relative versions of these results for fixed subsurfaces of Σ. An example of a relative version of Theorem 1.3 has been formulated as Theorem 1.5. The proof is essentially the same except that we now deal with the curve graph, $\mathcal{G}(\Phi)$, associated to Φ, rather than that associated to Σ. This is hyperbolic with constants depending on $\kappa(\Phi)$, and hence bounded in terms of $\kappa(\Sigma)$.

We should make some remarks regarding the interpretation of "pleating surfaces" in this context.

Suppose that $\Phi \subseteq \Sigma$ is a fixed connected efficient surface. Write δ for the multicurve associated to the relative boundary, $\partial_\Sigma \Phi$, of Φ in Σ, i.e. $\partial_\Sigma \Phi$ with any pair of homotopic curves identified. Suppose that γ is a multicurve in $XM(\Phi)$. Thus $\gamma \cup \delta$ is a multicurve which we can realise in M. Let $\phi : S, \alpha \sqcup \beta \longrightarrow M, \gamma \sqcup \delta$ be a pleating surface for $\gamma \sqcup \delta$. In the hypotheses of Theorem 1.5, $L(M, \delta)$ and hence $L(S, \beta)$ is bounded. Since $\omega_\phi(\beta) = \chi \circ \phi(\beta) = \delta$ is fixed, we can assume that β is a fixed multicurve in S. Indeed we can define a fixed map, $\xi : \Phi \longrightarrow S$, possibly identifying relative boundary components but otherwise injective, so that $\omega_\phi \circ \xi | \partial_\Sigma \Phi$ is homotopic to the inclusion of $\partial_\Sigma \Phi$ into Σ. (Recall that $\omega_\phi = \chi \circ \phi : S \longrightarrow \Sigma$.) We can thus talk about a pleating surface relative to Φ as a composition of $f = \phi \circ \xi : \Phi, \beta \longrightarrow M, \delta$, with the property that $\omega_f | \partial_\Sigma \Phi$ is homotopic to the identity, where $\omega_f = \chi \circ f$, and such that the curves $f(\partial_\Sigma \Phi)$ have bounded length in M. The earlier discussion of pleating surfaces now applies, where the mapping class group of Σ is replaced by that of Φ (fixing setwise each component of $\partial_\Sigma \Phi$). In particular, in the "modified definition" in Section 7, we can take a fixed hyperbolic structure on Φ where each component of $\partial_\Sigma \Phi$ is totally geodesic, and is mapped to the corresponding geodesic in M. The lipschitz constants involved depend on $\kappa(\Sigma)$ and on the constant, k, which bounds the length of these boundary geodesics in M, as given in the hypotheses of the theorem.

Proof of Theorem 1.5 : This now follows exactly as in the proof of Theorem 1.3. Pleating surfaces are interpreted as maps from Φ into M as above. In the combinatorial arguments, the curve graph, $\mathcal{G}(\Sigma)$, is replaced by the curve graph, $\mathcal{G}(\Phi)$. Efficient surfaces are defined in Φ. ♢

As remarked in Section 1, together with Proposition 9.3, this gives the bound on lengths of curves in a hierarchy. The notion of a "hierarchy" was defined in [MaM2]. In that paper, a hierarchy has a lot more structure, but here we will just consider the underlying set of curves, i.e. a subset of $\mathcal{G}(\Sigma)$. This is defined by an inductive procedure. Here, we consider a simplified version. As far as length bounds go, this includes the case of [MaM2] — the underlying set of a hierarchy there will always be a set of the type described below.

Suppose $Q \subseteq X(\Sigma)$. We construct a larger subset, $J(Q) \subseteq X(\Sigma)$, as follows. Suppose we can find a connected efficient subsurface, $\Phi \subseteq \Sigma$, and two curves $\alpha, \beta \in X(\Phi) \subseteq X(\Sigma)$, such that α, β and all components of $\partial_\Sigma \Phi$ lie in Q. If $\gamma \in X(\Phi)$ lies on a tight geodesic in $\mathcal{G}(\Phi)$ from α to β then we include γ in $J(Q)$. Thus $J(Q)$ consists of Q together with all curves, γ, arising in this way. If α and β are any two elements of $X(\Sigma)$, we can always take $\Phi = \Sigma$. Thus $J(\{\alpha,\beta\})$ is the union of all tight geodesics from α to β in Σ (or vertices thereof). For each $n \in \mathbf{N}$, we define inductively the sets $J_n(Q)$ by $J_{n+1}(Q) = J(J_n(Q))$.

Now it follows from Theorem 1.5 that given $k \geq 0$, there is some $K \geq 0$ such that $J(X(M,k)) \subseteq X(M,K)$. Thus there is some K_n depending on $\kappa(\Sigma)$, k and n such that $J_n(X(M,k)) \subseteq X(M,K_n)$.

To apply this to hierarchies, we need a slight modification, that will use a result we postpone until Section 9. We define $J'(Q) \subseteq J(Q)$ as for $J(Q)$, except that we allow $\kappa(\Phi) = 0$, and in this case take replace $\mathcal{G}(\Phi)$ by the modified curve graph $\mathcal{H}(\Phi)$ defined in Section 9. We define $J'_n(Q)$ by iterating this construction. Bringing Proposition 9.3 into play, the previous paragraph also applies to $J'_n(Q)$.

If $\alpha, \beta \in X(\Sigma)$, the vertices of a "hierarchy" between α and β form a canonical subset $H(\alpha,\beta) \subseteq X(\Sigma)$. All we care about here is that $H(\alpha,\beta) \subseteq J'_n(\{\alpha,\beta\})$ for some n depending only on $\kappa(\Sigma)$. (This applies to both [Mi4] and [Bow5].) We can thus deduce:

Corollary 8.3 : *Given $k \geq 0$, there is some $L \geq 0$, depending only on $\kappa(\Sigma)$ and k such that if $\alpha, \beta \in X(M,k)$, then $H(\alpha,\beta) \subseteq X(M,L)$.* $\diamondsuit$

This is essentially Lemma 7.9 of [Mi4], which is one of the key steps towards the Ending Lamination Conjecture. This also gives the "a-priori bounds" statement required for [Bow5].

9. Length bounds for exceptional surfaces.

In this section, we give an account of length bounds for the exceptional surfaces, namely the one-holed torus and four-holed sphere. In these cases the curve graphs as defined earlier are just countable sets of points, and it is natural to consider the modified curve graph (denoted $\mathcal{H}$) defined below. Our proof is based on certain trace identities that can be found, for example, in [Go].

Let 1HT and 4HS denote respectively the one-holed torus and four-holed sphere. If Φ is one of these surfaces, we write $X(\Phi)$ for the set of non-peripheral simple curves as before. The modified graph, $\mathcal{H} = \mathcal{H}(\Phi)$ has vertex set $V(\mathcal{H}) = X(\Phi)$, and $\alpha, \beta \in X(\Phi)$ are deemed adjacent if they have minimal possible intersection, i.e. can be realised so that $|\alpha \cap \beta| = 1$ for 1HT and $|\alpha \cap \beta| = 2$ for 4HS. In both cases, $\mathcal{H}$ is isomorphic to the Farey graph.

Suppose $\rho : \pi_1(\Phi) \longrightarrow SL(2,\mathbf{C})$ is any homomorphism. Given $\alpha \in X(\Phi)$, write $t(\alpha) = \mathrm{tr}\,\rho(\underline{\alpha})$, where $\underline{\alpha} \in \pi_1(\Phi)$ is any representative of the free homotopy class of α. The *complex length*, $\lambda(\alpha)$ is defined by the formula $2\cosh\lambda(\alpha) = 2\,\mathrm{tr}\,\alpha$ and is well defined in $\mathbf{C}/2\pi i\mathbf{Z}$. The *(real) length*, $L(\alpha)$, of α is the real part of $\lambda(\alpha)$ and is well defined in $\mathbf{R}$. Indeed $L(\alpha)$ and $|t(\alpha)|$ depend only on the projection of ρ to $PSL(2,\mathbf{C})$. Note also that $L(\alpha)$ is bounded above in terms of $|t(\alpha)|$ and conversely.

If Φ is a 1HT, write ζ for its boundary curve, and set $L(\partial\Phi) = L(\zeta)$ and $|t|(\partial\Phi) = |t(\zeta)|$. If Φ is a 4HS, write $\zeta_1, \zeta_2, \zeta_3, \zeta_4$ for the boundary curves, and set $L(\partial\Phi) = \max\{L(\zeta_i) \mid i = 1,2,3,4\}$ and $|t(\partial\Phi)| = \max\{|t(\zeta_i)| \mid i = 1,2,3,4\}$.

We first show:

Lemma 9.1 : *For all l there exists L such that if $\rho : \pi_1(\Phi) \longrightarrow SL(2,\mathbf{C})$ is any representation, and $\alpha, \beta \in X(\Phi)$ with $L(\alpha), L(\beta), L(\partial\Phi) \le l$, there is some arc, π, from α to β in $\mathcal{H}(\Phi)$ such that $L(\gamma) \le L$ for all $\gamma \in \pi \cap V(\mathcal{H})$.*

Note that by an earlier observation, this can be translated into an equivalent statement about absolute values of traces, which is what we actually prove.

We shall focus on the case of the 4HS. The case of 1HT can be dealt with by a similar, but somewhat simpler argument, as described in [Bow1]. (That paper assumed the peripheral curves to be parabolic, though that does not affect this particular argument.) A detailed discussion of discrete 1HT groups is given in [Mi1] (again under the assumption that peripheral subgroups are parabolic), where a similar result is proven by geometric arguments. A more general discussion, including a proof of Lemma 9.1 for the 1HT, is given in [Z].

For the 4HS, our argument is based on the following trace identity. Let $P, Q, R, S \in SL(2,\mathbf{C})$, with $PQRS = I$. Let $p = \operatorname{tr} P$, $q = \operatorname{tr} Q$, $r = \operatorname{tr} R$, $s = \operatorname{tr} S$, $a = \operatorname{tr} PQ = \operatorname{tr} RS$, $b = \operatorname{tr} QR = \operatorname{tr} SP$, $c = \operatorname{tr} PR$, $d = \operatorname{tr} QS$, then

$$ab + c + d = pr + qs.$$

We can interpret this in terms of the Farey graph, $\mathcal{H} = \mathcal{H}(\Phi)$. We say that four curves, $\alpha, \beta, \gamma, \delta$, in $V(\mathcal{H})$ form a *rhombus* if α, β are adjacent and γ, δ are the opposite vertices of the two triangles meeting along the edge $\alpha\beta$. In this case, for suitable choice of representatives $\underline{\zeta}_i$ or ζ_i, we can represent $\alpha, \beta, \gamma, \delta$ by $\underline{\alpha} = \underline{\zeta}_1\underline{\zeta}_2$, $\underline{\beta} = \underline{\zeta}_2\underline{\zeta}_3$, $\underline{\gamma} = \underline{\zeta}_1\underline{\zeta}_3$ and $\underline{\delta} = \underline{\zeta}_2\underline{\zeta}_4$. Suppose $\rho : \pi_1(\Phi) \longrightarrow SL(2,\mathbf{C})$ is any representation. Write $z_i = t(\zeta_i)$, $a = t(\alpha)$, $b = t(\beta)$ $c = t(\gamma)$, $d = t(\delta)$. We see that $ab + c + d = z_1 z_3 + z_2 z_4$. In particular, $|ab + c + d| \le 2k^2$, where $k = \max\{|z_i| \mid i = 1, 2, 3, 4\}$. Note that k is bounded in terms of $L(\partial\Phi)$.

We now put transverse orientations in the edges of the Farey graph. (This can be interpreted in terms of orienting edges of the dual tree as in [Bow1].) Throughout we use the convention that Greek letters $\alpha, \beta, \gamma, \delta$ etc. denote vertices of $\mathcal{H}$ and the corresponding Latin letters denote traces, i.e. $a = t(\alpha)$, $b = t(\beta)$ etc.

Suppose that $\alpha, \beta, \gamma, \delta$, form a rhombus in $\mathcal{H}$. We say the edge $\alpha\beta$ is transversely oriented from γ to δ if $|c| \ge |d|$. (If $|c| = |d|$ one can assign the orientation arbitrarily.) If $|ab| \ge 6k^2$, this implies that $|c| \ge \frac{1}{3}|ab|$. To see this, note that $|ab + c + d| \le 2k^2 \le \frac{1}{3}|ab|$, so $|c + d| \ge |ab| - \frac{1}{3}|ab| \ge \frac{2}{3}|ab|$, Since $|c| \ge |d|$, we get $|c| \ge \frac{1}{3}|ab|$ as claimed.

We also note that if $|c|, |d| \le 3k^2$, then $\min\{|a|, |b|\} \le 3k$, for if not, we would derive the contradiction $3k^2 = 9k^2 - 6k^2 \le |ab| - |c| - |d| \le |ab + c + d| \le 2k^2$.

Proof of Lemma 9.1 : Suppose k is any constant greater than $\max\{|t(\partial\Phi)|, 1\}$.

Let $V_0 = \{\epsilon \in V(\mathcal{H}) \mid |t(\epsilon)| \le 3k\}$. Let $\mathcal{H}_0 \subseteq \mathcal{H}$ be the full subgraph on vertex set V_0. The second observation above tells us that if $\alpha, \beta, \gamma, \delta$ form a rhombus with $\alpha, \beta \notin V_0$, then either $\gamma \notin V_0$ or $\delta \notin V_0$. In other words any two distinct components of V_0 are distance at least 3 apart.

We claim that $\mathcal{H}_0$ is, in fact, connected. For if not, we can find a triangle abc in $\mathcal{H} \setminus \mathcal{H}_0$ such that the edges ab and ac are oriented away from c and b respectively. (This most easily seen in terms of the dual tree, T. Note that every edge of $\mathcal{H}$ separates, so crosses an edge of the dual tree. Let T_0 be the union of edges such that the crossing edge of $\mathcal{H}$ meets $\mathcal{H}_0$. If $\mathcal{H}_0$ is not connected, then neither is T_0. Consider an arc, τ, connecting two components of T_0. At the initial and final edges of τ, the flow is out of τ, and so there is some vertex of τ at which the flow on the two incident edges is outward. This gives us our triangle.) But, by the first observation, we have $|ab| \leq 3|c|$ and $|ac| \leq 3|b|$ so that $|a| \leq 3 \leq 3k$, giving the contradiction that $a \in V_0$.

To apply this to Lemma 9.1, we just take L to be bigger than k and l. ♢

We have shown that lengths are bounded on some arc from α to β. We would like to say that they are bounded on all geodesics. The following argument, suggested by Ser Peow Tan, greatly simplifies my original.

We use the following trace identity. Given $P, Q, R, S \in SL(2, \mathbf{C})$ with $PQRS = I$, and traces a, b, c, p, q, r, s as defined earlier, then

$$a^2 + b^2 + c^2 + abc - (pq + rs)a - (ps + qr)b - (pr + qs)c + p^2 + q^2 + r^2 + s^2 - pqrs = 4.$$

We only need note that it has the form $c^2 + Ac + B = 0$, where A, B are polynomials in the other traces. Thus, c is bounded above in terms of a, b, p, q, r, s. We obtain:

Lemma 9.2 : *If $\alpha, \beta, \gamma \in V(\mathcal{H})$ are the vertices of a triangle, then $L(\gamma)$ is bounded above by a fixed function of* $\max\{L(\alpha), L(\beta), L(\partial\Phi)\}$. ♢

We can now prove the main result of this section:

Proposition 9.3 : *Given $l \geq 0$, there is some $L \geq 0$, such that if $\rho : \pi_1(\Phi) \longrightarrow PSL(2, \mathbf{C})$ is any representation, and if $\alpha, \beta \in V(\mathcal{H})$ with* $\max\{L(\alpha), L(\beta), L(\partial\Phi)\} \leq l$, *then $L(\gamma) \leq L$ for any γ lying in any geodesic from α to β.*

Proof : Since $\pi_1(\Phi)$ is free, ρ lifts to a representation to $SL(2, \mathbf{C})$, so by Lemma 9.1, there is a path π from α to β, with $L(\gamma)$ uniformly bounded in terms of l along π. Now if $\delta \in V(\mathcal{H}) \setminus V(\pi)$ lies in any other geodesic from α to β, then from the combinatorics of the Farey graph, we see that γ is adjacent to two vertices in $V(\pi)$. Thus $L(\delta)$ is bounded by Lemma 9.2 as required. ♢

The proofs we have given make reference only to 4HS. The case of 1HT is essentially the same. This time, we use the trace identity $\operatorname{tr} PQ + \operatorname{tr} PQ^{-1} = \operatorname{tr} P \operatorname{tr} Q$. Applied to a rhombus $\alpha, \beta, \gamma, \delta$, this gives $ab = c + d$, and the argument proceeds as before. See [Bow1,Mi1] for further discussion of this case.

References.

[BelF] G.Bell, K.Fujiwara, *The asymptotic dimension of a curve graph is finite* : preprint Greensboro/Sendai (2005).

[BesF] M.Bestvina, K.Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : Geom. Topol. **6** (2002) 69–89.

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

[Bow1] B.H.Bowditch, *Markoff triples and quasifuchsian groups* : Proc. London Math. Soc. **77** (1998) 697–736.

[Bow2] B.H.Bowditch, *Intersection numbers and the hyperbolicity of the curve complex* : J. reine angew. Math. **598** (2006) 105–129.

[Bow3] B.H.Bowditch, *Systems of bands in hyperbolic 3-manifolds* : preprint, Southampton (2003).

[Bow4] B.H.Bowditch, *Tight geodesics in the curve complex* : preprint, Southampton (2003).

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

[Bow6] B.H.Bowditch, *End invariants of hyperbolic 3-manifolds* : preprint, Southampton (2005).

[BrCM1] J.F.Brock, R.D.Canary, Y.N.Minsky, *Classification of Kleinian surface groups II: The ending lamination conjecture* : preprint (2004).

[BrCM2] J.F.Brock, R.D.Canary, Y.N.Minsky, *The classification of finitely generated kleinian groups* : in preparation.

[CanEG] R.D.Canary, D.B.A.Epstein, P.Green, *Notes on notes of Thurston* : in "Analytic and geometric aspects of hyperbolic space" London Math. Soc. Lecture Notes Series No. 111, (ed. D.B.A.Epstein) Cambridge University Press (1987) 3–92.

[CasB] A.J.Casson, S.A.Bleiler, *Automorphism of surfaces after Nielsen and Thurston* : London Math. Soc. Student Texts No. 9, Cambridge University Press (1988).

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

[Go] W.M.Goldman, *Ergodic theory on moduli spaces* : Ann. of Math. **146** (1997) 475–507.

[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.

[Ham] U.Hamenstädt, *Geometry of the complex of curves and of Teichmüller space* : Preprint, Bonn (2005).

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

[Ki] Y.Kida, *The mapping class group from the viewpoint of measure equivalence theory* : to appear in Mem. Amer. Math. Soc.

[Ko] T.Kobayashi, *Heights of simple loops and pseudo-Anosov homeomorphisms* : in “Braids (Santa Cruz, CA, 1986)”, Contemp. Math. No. 78, Amer. Math. Soc. (1988) 327–338.

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

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

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

[Mi1] Y.N.Minsky, *The classification of punctured-torus groups* : Ann. of Math. **149** (1999) 559–626.

[Mi2] Y.N.Minsky, *Kleinian groups and the complex of curves* : Geom. Topol. **4** (2000) 117–148.

[Mi3] Y.N.Minsky, *Bounded geometry for Kleinian groups* : Invent. Math. **146** (2001) 143–192.

[Mi4] Y.N.Minsky, *The classification of Kleinian surface groups I : Models and bounds* : preprint, Stony Brook (2002).

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

[Sh1] K.J.Shackleton, *Tightness and computing distances in the curve complex* : preprint, Southampton (2004).

[Sh2] K.J.Shackleton, *A computational acylindricity theorem for the mapping class group* : preprint, (2005).

[So] T.Soma, *Geometric limits of quasifuchsian groups* : preprint, Tokyo Denki.

[T] W.P.Thurston, *Hyperbolic structures on 3-manifolds I: Deformations of acylindrical 3-manifolds* : Ann. of Math. **124** (1986) 203–246.

[Z] Y.Zhang, *Hyperbolic cone surfaces, generalized Markoff maps, Schottky groups and McShane’s identity* : Thesis, National University of Singapore (2004).