Tight geodesics in the curve complex

Brian H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
http://www.maths.soton.ac.uk/staff/Bowditch
[First Draft: October 2003. Revised: November 2003]

Let Σ be a compact orientable surface. We define the *complexity* of Σ as $\kappa(\Sigma) = 3g + p - 4$, where g is the genus of Σ and p is the number of boundary components. Let $\Gamma(\Sigma)$ be the mapping class group of Σ, that is the group of orientation preserving self-homeomorphisms of Σ defined up to homotopy.

The *curve complex* [Harv] of Σ is the simplicial complex whose vertex set, $X = X(\Sigma)$, consists of homotopy classes of essential non-peripheral simple closed curves of Σ, and where a (finite) set of such curves is deemed to bound a simplex if its elements can be realised disjointly in Σ. In the "non-exceptional" cases, that is when $\kappa(\Sigma) > 0$, this complex is connected and has dimension $\kappa(\Sigma)$. Note that $\Gamma(\Sigma)$ acts simplicially on the curve complex. with finite quotient (after taking barycentric subdivision). This action has been exploited by various authors to study $\Gamma(\Sigma)$, see for example [Hare,I,BeF] and references therein.

All that is directly relevant to the present paper is the 1-skeleton, $\mathcal{G}(\Sigma)$, of the curve complex, which we refer to as the *curve graph*. We write d for the combinatorial metric on $\mathcal{G}(\Sigma)$ so that each edge has unit length. In [MaM1], it was shown that, provided $\kappa(\Sigma) > 0$, the graph $\mathcal{G}(\Sigma)$ is hyperbolic in the sense of Gromov [Gr,GhH] (see also [Bo2] for another proof).

Applications of this this result are complicated by the fact that $\mathcal{G}(\Sigma)$ is far from being locally finite. To some extent this is remedied in [MaM2], where certain hierarchical finiteness properties of the curve complex are investigated. In particular, they define a canonical set of "tight geodesics" connecting any two vertices, a, b of $\mathcal{G}$. They show that this set is always non-empty and finite. (The existence of tight geodesics is direct combinatorial argument. For another proof of finiteness, see [Bo4].)

In this paper, we shall refine the above finiteness statement, giving bounds on the number of points that can lie in a tight geodesics in a bounded neighbourhood of a given point (Theorems 1.1 and 1.2). We consider applications of this to the geometry of the action of $\Gamma(\Sigma)$ on $\mathcal{G}(\Sigma)$, in particular, acylindricity (Theorem 1.3) and uniform rationality of stable lengths (Theorem 1.4 and Corollary 1.5). We remark that acylindricity implies, in turn, the "weak proper discontinuity" condition used in [BeF].

1. Statement of results.

Let $\mathcal{G} = \mathcal{G}(\Sigma)$ be the curve complex associated to Σ, where $\kappa(\Sigma) > 0$. This is k-hyperbolic for some $k = k(\Sigma)$. We refer to the elements of $X = X(\Sigma)$ simply as *curves*. We say that two curves are *disjoint* if they are adjacent in $\mathcal{G}(\Sigma)$, otherwise, we say that they *cross*. A *multicurve* is a subset of X consisting of pairwise disjoint elements. (Such elements can be rendered simultaneously disjoint, and hence bound a simplex in the curve complex.)

A *multigeodesic* is a sequence, $(A_i)_{i=0}^n$ of multicurves such that for all $i \neq j$ and all $a \in A_i$ and $b \in A_j$, $d(a,b) = |i-j|$. We say that $(A_i)_i$ is *tight* at index $i \neq 0, n$ if, for all $a \in A_i$, each curve that crosses a also crosses some element of $A_{i-1} \cup A_{i+1}$.

We remark that this definition is weaker than that given in [MaM2]. We are effectively saying that each element of A_i is the boundary curve of the (connected) subsurface of Σ filled by $A_{i-1} \cup A_{i+1}$. In [MaM2] it is required that A_i consists of all such boundary curves (non-peripheral in Σ), though we do not need this additional information here.

We say that a *multigeodesic* $(A_i)_{i=0}^n$ is *tight* if it is tight for all indices in $\{1, \ldots, n-1\}$. A geodesic consisting of vertices $(a_i)_{i=0}^n$ is *tight* if there is some multigeodesic $(A_i)_{i=0}^n$ such that $a_i \in A_i$ for all i.

Given $a, b \in V(\mathcal{G}) = X$, we write $\mathcal{L}(a,b)$ for the set of all geodesics in $\mathcal{G}$ connecting a to b. Let $\mathcal{L}_T(a,b) \subseteq \mathcal{L}(a,b)$ be the subset of tight geodesics. We have observed that $\mathcal{L}_T(a,b)$ is non-empty and finite [MaM2]. We write $G(a,b) = \bigcup \mathcal{L}_T(a,b) \subseteq \mathcal{G}$.

Let $N(c;r)$ denote the r-ball about c in $\mathcal{G}$. There is some constant k_0, depending only on k and hence on $\kappa(\Sigma)$ such that if c lies in some geodesic from a to b, then each geodesic from a to b passes through $N(c;k_0)$.

We show:

Theorem 1.1 : *There is some K_0 depending only on $\kappa(\Sigma)$ such that if $a, b \in V(\mathcal{G})$ and $c \in G(a,b)$, then $G(a,b) \cap N(c;k_0)$ has at most K_0 elements.*

In view of the choice of k_0, this is equivalent to asserting that the slices of $G(a,b)$ have bounded cardinality (in terms of $\kappa(\Sigma)$). Here a *slice* of $G(a,b)$ can be defined as a subset of the form $\{x \in G(a,b) \mid d(a,x) = p\} = \{x \in G(a,b) \mid d(b,x) = n-p\}$ for some $p \in \{0, \ldots, n\}$.

One might still imagine that the set $G(a,b)$ could change dramatically if we move one of the endpoints a or b. The following variation of Theorem 1.1 gives some control on this, and seems to be of more practical use. Given $A, B \subseteq V(\mathcal{G})$, let $\mathcal{L}_T(A,B) = \bigcup \{\mathcal{L}_T(a,b) \mid a \in A, b \in B\}$ and $G(A,B) = \bigcup \mathcal{L}_T(A,B) = \bigcup \{G(a,b) \mid a \in A, b \in B\}$ etc. We write $G(a,b;r) = G(N(a;r), N(b;r))$. By hyperbolicity, we can again assume that k_0 is such that if $c \in \bigcup \mathcal{L}(a,b)$, then any geodesic connecting $N(a;r)$ to $N(b;r)$ must intersect $N(c,k_0)$.

Theorem 1.2 : *There are constants k_1 and K_1, depending only on $\kappa(\Sigma)$ such that if $a, b \in V(\mathcal{G})$, $r \in \mathbb{N}$ and $c \in G(a,b)$ with $d(c, \{a,b\}) \geq r + k_1$, then $G(a,b;r) \cap N(c;k_0)$ has at most K_1 elements.*

Note that Theorem 1.1 doesn't quite follow from Theorem 1.2 because of the constant k_1. However, it does follow if we include the additional information that there are boundedly many tight geodesics connecting two given vertices at distance at most $k_0 + k_1$ apart.

The proofs of Theorems 1.1 and 1.2 make use of 3-dimensional hyperbolic geometry.

We need to connect two curves $a, b \in X$ by some kind of geometric object, and a natural choice would seem to be a hyperbolic 3-manifold M, homotopy equivalent to Σ, in which the homotopy classes of curves corresponding to a and b are realised by "short" closed geodesics. We use two facts about the geometry of M. Firstly there is a system of "bands" in M as constructed in [Bo3]. (This is related to the block decomposition described in [Mi].) Secondly, we use the "a-priori" bounds on the length of closed geodesics in M associated to curves in $\bigcup \mathcal{L}_T(a,b)$. These are proven in [Mi] (at least using the stronger definition of tight geodesic given there). A more direct proof of this particular statement is given in [Bo4].

We remark that, although the various constants featuring in [Bo3] are (in principle) computable in terms of $\kappa(\Sigma)$, those in [Mi] or [Bo4] make use of geometrical limit arguments. Thus, we have no explicit means of computing the constants K_0 and K_1 featuring in Theorems 1.1 and 1.2. The same therefore applies to the results that follow, in particular Theorems 1.3 and 1.4 below.

The above results have consequences for the geometry of the action of $\Gamma(\Sigma)$ on $\mathcal{G}(\Sigma)$ which we go on to describe.

The following definition makes sense for any isometric action of a group Γ on a k-hyperbolic graph $\mathcal{G}$ (or indeed, any path-metric space).

Definition : We say that the action of Γ is *acylindrical* if for all $r \geq 0$, there exist $R, N \geq 0$ such that for all $a, b \in V(\mathcal{G})$ with $d(a,b) \geq R$ there are at most N distinct elements, g, of Γ such that $d(a, ga) \leq r$ and $d(b, gb) \leq r$.

This definition will be elaborated on in Section 2. We note the following consequences.

Clearly, setting $r = 0$, we obtain constants, R_0 and N_0 such that if $d(a,b) \geq R_0$, then $|\operatorname{stab} a \cap \operatorname{stab} b| \leq N_0$, where $\operatorname{stab} x$ denotes the Γ-stabiliser of x. Indeed, if $\mathcal{G}$ were a simplicial tree, this would be sufficient (cf. the acylindricity condition of Sela [S]). The same applies to a uniformly locally finite graph (i.e. every vertex has bounded degree). A situation that encompasses both is that of a uniformly fine hyperbolic graph [Bo1]. Curve graphs, however, fail to have this property.

A more significant consequence is that one can partition the elements of Σ into elliptic and loxodromic elements. Moreover, there is the positive lower bound on the stable lengths of loxodromic elements (Lemma 2.2). Recall that the *stable length*, $||g||$, of $g \in \mathcal{G}$ is defined as $\lim_{n\to\infty} \frac{1}{n} d(x, g^n x)$ for some (any) $x \in V(G)$.

We shall show:

Theorem 1.3 : *If* $\kappa(\Sigma) > 0$*, then the action of* $\Gamma(\Sigma)$ *on* $\mathcal{G}(\Sigma)$ *is acylindrical.*

In this case, the division into elliptic elements and loxodromic elements is a consequence of Thurston's classification of mapping classes. An element of $\Gamma(\Sigma)$ might be finite order or reducible. In the latter case, it preserves some simplex of the curve complex. Both cases are clearly elliptic. All other elements of $\Gamma(\Sigma)$ are pseudoanosov. In [MaM1] it is shown that pseudoanosov elements are loxodromic. The positive lower bound on stable lengths however seems to be new.

Note that if $a, b \in X(\Sigma)$ satisfy $d(a, b) \geq 3$, then $a \cup b$ fills Σ, and it is easily seen that $|\operatorname{stab} a \cap \operatorname{stab} b|$ is uniformly bounded in terms of $\kappa(\Sigma)$.

We also remark that a weaker version of acylindricity, termed "weak proper discontuity" is proven in [BeF]. In fact an elaboration of their argument shows that if $r \geq 0$, and $a, b \in X(\Sigma)$ with $d(a, b) \geq 2r + 3$, then $\{g \in \Gamma \mid d(a, ga) \leq r, d(b, gb) \leq r\}$ is finite. It is unclear, however, how one could use these methods to bound the cardinality of such sets.

It turns out that one can get more information about stable lengths than is implied directly by acylindricity. We show:

Theorem 1.4 : *Given $\kappa = \kappa(\Sigma) > 0$ there is some $m = m(\Sigma)$ such that if $g \in \mathcal{G}(\Sigma)$ is loxodromic (i.e. pseudoanosov) then g^m preserves some bi-infinite geodesic in $\mathcal{G}$.*

Thus g^m translates this geodesic some distance $p \in \mathbb{N}$, and so $||g|| = ||g^m||/m = p/m$. It follows that stable lengths are uniformly rational:

Corollary 1.5 : *Given $\kappa = \kappa(\Sigma) > 0$, there is some $n = n(\Sigma) > 0$ such that if $g \in \Gamma$, then $n||g|| \in \mathbb{N}$.*

The analogous statement for hyperbolic groups is given in [Gr], and an elegant proof can be found in [D].

2. Acylindrical actions.

Let $\mathcal{G}$ be a k-hyperbolic graph with vertex set $V(\mathcal{G})$ and combinatorial metric d. Let $\partial\mathcal{G}$ be its Gromov boundary, defined as equivalence classes of quasigeodesic rays, where two rays are deemed equivalent if the Hausdorff distance between them is finite.

Suppose that $\pi \subseteq \mathcal{G}$ is an (oriented) geodesic segment with vertices $x_0 x_1 \dots x_n$. Given $r \geq 0$, we refer to the geodesics $x_r x_{r+1} \dots x_{n-r}$ as the r-central segment of π. If $\pi' = y_0 y_1 \dots y_n$ is another segment (of the same length) we say that π and π' are *r-close* if $d(x_i, y_i) \leq r$ for all i. More generally, if $l \in \mathbb{Z}$, we say that π' is *r-closely translated a distance l* from π, if $d(y_i, x_{i+l}) \leq r$ for all $i \in \{0, \dots, n-l\}$.

We state the following observation as a lemma, since it is central to what follows.

Lemma 2.1 : *There is some constant $k_0 \geq 0$ such that if $\pi \in \mathcal{L}(a, b)$ and $\pi' \in \mathcal{L}(a', b')$, then the r-central segment of π is k_0-close to a segment of π', where $r = \max\{d(a, a'), d(b, b')\}$.* ◊

Here k_0 is a fixed (universal) multiple of k. Recall that $\mathcal{L}(a, b)$ is the set of all geodesics from a to b.

Now suppose the group Γ acts on $\mathcal{G}$. If $g \in \Gamma$, its stable length $||g||$ is defined as $\lim_{n\to\infty} \frac{1}{n} d(x, g^n x)$ for $x \in V(\mathcal{G})$. One verifies that this is well-defined, independently of x, that it is a conjugacy invariant, and that $||g^n|| = n||g||$ for all $n \in \mathbb{N}$.

Definition : An element $g \in \Gamma$ is *elliptic* if some (hence every) $\langle g \rangle$-orbit in $\mathcal{G}$ has finite diameter.
An element $g \in \Gamma$ is *loxodromic* if $||g|| > 0$.

Clearly these are two categories are mutually exclusive (though not exhaustive in general).

If g is elliptic, one can show that there is some $\langle g \rangle$-orbit whose diameter is bounded purely in terms of k. If g is loxodromic, then for any $x \in V(\mathcal{G})$, $(g^n x)_{n \in \mathbb{Z}}$ is quasigeodesic sequence, i.e. each point lies in a $\langle g \rangle$-invariant bi-infinite quasigeodesic. One deduces that g has precisely two fixed points in $\partial \mathcal{G}$.

Given $a, b \in V(\mathcal{G})$, and $r \geq 0$, let $\Delta(a, b; r) = \{g \in \Gamma \mid d(a, ga) \leq r, d(b, gb) \leq r\}$. Recall that the action of Γ on $\mathcal{G}$ is acylindrical if

$$(\forall r \geq 0)(\exists R, N \geq 0)(\forall a, b \in V(\mathcal{G}))(d(a, b) \geq R \Rightarrow |\Delta(a, b; r)| \leq N).$$

In view of Lemma 2.1, one verifies easily that if $g \in \Delta(a, b; r)$ and π_0 is the r-central segment of π, then π_0 is $(3k_0)$-closely translated some distance $l = l(g)$. (Up to an additive constant, $l(g)$ could be defined as $\frac{1}{2}(d(a, gb) - d(b, ga))$.)

From this observation, one can see that it is sufficient to verify the criterion of acylindricity for some fixed $r = r_0$ depending only on k.

One consequence of acylindricity is the following:

Lemma 2.2 : *Suppose Γ acts acylindrically on $\mathcal{G}$. Then each element of Γ is either elliptic or loxodromic. Moreover, there is some $\epsilon > 0$ depending only on k and the parameters of acylindricity such that if $g \in \Gamma$ is loxodromic, then $||g|| \geq \epsilon$.*

This can be proven by an argument similar to that which shows that any infinite-order element of a hyperbolic group is loxodromic (see [GhH]). Moreover, in the case of the mapping class group acting on the curve complex, the conclusions can be deduced from other considerations (see Corollary 3.5). For these reasons, we only give a sketch of the argument below.

Proof : Suppose $g \in \Gamma$. Let $x \in V(\mathcal{G})$, and let $D = d(x, gx)$. Let $h > 0$ be a constant depending only on k as described below, and set $r = 10h$. Let R and N be the constants given by the acylindricity hypothesis for this r, and set $L \geq R + 2(N+1)D + 2h$ plus a suitable constant depending on k.

If g is not elliptic, there is some power, g^p, of g with $d(x, g^p x) \geq R$. Set $y = g^p x$. Let $\pi \in \mathcal{L}(x, y)$ and let $\pi_0 \subseteq \pi$ be its $(N+1)D$-central segment. This has length greater than $R + 2h + 2D$.

Suppose $i \in \{1, \ldots, N+1\}$. Now $d(x, g^i x) = d(y, g^i y) \leq (N+1)D$, and so π_0 is $(3k_0)$-closely translated some distance $l_i \in \mathbb{Z}$. Moreover, since $d(g^i x, g^{i+1} x) = d(g^i y, g^{i+1} y) = D$, we see that $|l_i - l_{i+1}|$ is at most D plus a constant depending on k.

Now applying the acylindricity assumption to the endpoints of π_0, we see that for some $i \in \{1, \ldots, N+1\}$, g^i closely translates π_0 a distance at least h and at most $h + D$ up to additive constants. But π_0 has length at least $2(h+D)$. Thus, provided h is chosen large enough in relation to k, one can show that g is loxodromic, with $||g^i||$ bounded below by some positive constant, say $h/2$. (This uses the fact that paths that are quasigeodesic over all sufficiently long subpaths are globally quasigeodesic. From this one deduces that the translates of π_0 under $\langle g \rangle$ lie uniformly close to a bi-infinite $\langle g \rangle$-invariant quasigeodesic.) We deduce that $||g|| \geq h/2i \geq h/2(N+1)$, so we may set $\epsilon = h/2(N+1)$, which ultimately depends only on k and the acylindricity parameters as claimed. $\diamondsuit$

3. Tight geodesics.

Let $\mathcal{G}$ be a k-hyperbolic graph. Given $a, b \in V(\mathcal{G})$ let $\mathcal{L}(a,b)$ be the set of all geodesics from a to b. We suppose that we are given a non-empty subset, $\mathcal{L}_T(a,b) \subseteq \mathcal{L}(a,b)$ of geodesics, which we refer to as *tight*. We suppose that every subpath of a tight geodesic is tight. We write $G(a,b) = \bigcup \mathcal{L}_T(a,b)$. We define $\mathcal{L}_T(a,b;r)$ and $G(a,b;r)$ as before.

We can define a tight geodesic ray or bi-infinite geodesic as one for which every finite subpath is tight. We extend the notations $\mathcal{L}(a,b), \mathcal{L}_T(a,b), G(a,b)$ to allow $a, b \in V(\mathcal{G}) \cup \partial \mathcal{G}$.

Let us suppose that the collection $\mathcal{L}_T = (\mathcal{L}_T(a,b))_{a,b \in V(\mathcal{G})}$ satisfies the conclusions of Theorems 1.1 and 1.2, namely it has the following finiteness properties:

(F1): $(\exists P_0 \in \mathbb{N})(\forall a, b \in V(\mathcal{G}))(\forall c \in G(a,b))(|G(a,b) \cap N(c,k_0)| \leq P_0)$.

(F2): $(\exists P_1, k_1 \in \mathbb{N})(\forall r \geq 0)(\forall a, b \in V(\mathcal{G}))$ for all $c \in G(a,b)$ with $d(c, \{a,b\}) \geq r + k_1$, we have $|G(a,b;r) \cap N(c,k_0)| \leq P_1$.

Recall that k_0 is the constant of Lemma 2.1. Note also that in view of Lemma 2.1, it is enough to verify (F2) for $r = k_0$.

Lemma 3.1 : *For all $a, b \in V(\mathcal{G}) \cup \partial \mathcal{G}$, $\mathcal{L}_T(a,b) \neq \emptyset$.*

Proof : We prove this in the case where $a, b \in \partial \mathcal{G}$. The case where $a \in V(\mathcal{G})$ and $b \in \partial \mathcal{G}$ is similar.

Let $(x_n)_{n \in \mathbb{Z}}$ be a sequence of vertices with $x_{-n} \to a$ and $x_n \to b$ as $n \to \infty$. Given any $R \geq 0$, for all sufficiently large m and n, any geodesic from x_{-m} to x_m is k_0-close to any geodesic from x_{-n} to x_n in the ball $N(x_0, R)$. The result follows by taking tight geodesics, applying property (F2) and taking a diagonal subsequence of segments that converge on a bi-infinite tight geodesic. $\diamondsuit$

In particular, it follows that $\mathcal{L}(a,b) \neq \emptyset$ for all $a, b \in V(\mathcal{G}) \cup \partial \mathcal{G}$ — a fact not true for an arbitrary hyperbolic graph.

The following is also easily verified using Lemma 2.1:

Lemma 3.2 : *The finiteness properties (F1) and (F2) hold for $a, b \in V(\mathcal{G}) \cup \partial \mathcal{G}$.* $\diamondsuit$

Here we deem $d(c,a) = \infty$ if $a \in \partial \mathcal{G}$. The constants, P_0 and P_1 arising may differ from the originals. Henceforth, we fix a constant P that works in all cases.

Now suppose that Γ acts on $\mathcal{G}$ and suppose that $\mathcal{L}_T$ is Γ-equivariant, i.e. $g\mathcal{L}_T(a,b) = \mathcal{L}_T(ga, gb)$ for all $g \in \Gamma$ and $a, b \in V(\mathcal{G})$ (hence for all $a, b \in V(\mathcal{G}) \cup \partial \mathcal{G}$).

Let us suppose that the action of Γ satisfies:

(F3) There are constants, $R_0, N_0 \in \mathbb{N}$ such that if $a, b \in V(\mathcal{G})$ with $d(a,b) \geq R_0$, then $|\operatorname{stab} a \cap \operatorname{stab} b| \leq N_0$.

Lemma 3.3 : *Suppose that Γ and $\mathcal{G}$ satisfy (F2) and (F3). Then the action of Γ is acylindrical.*

Proof : Let k_1 and P be the constants of (F2) and let k_2 be a constant depending only on k as described below. Let $r \geq 0$, and let $R = N_0 + 4r + 2k_1 + 2k_2$.

Suppose $a, b \in V(\mathcal{G})$ with $d(a,b) \geq R$, and set $\Delta = \{g \in \Gamma \mid d(a, ga) \leq r, d(b, gb) \leq r\}$. We want to bound $|\Delta|$.

Let $\pi \in \mathcal{L}_T(a,b)$ and let $x, y \in \pi$ satisfy $d(x,a) = d(y,b) = 2r + k_1 + k_2$. Thus, $d(x,y) \geq N_0$.

Suppose $g \in \Delta$. Applying Lemma 2.1, we see that $gx \in N(p; k_0)$ and $gy \in N(q; k_0)$ for some $p, q \in \pi$. Moreover, $d(x,p) \leq r + k_2$ and $d(y,q) \leq r + k_2$, where k_2 just depends on k. In particular, $d(a,p) \geq r + k_1$ and $d(b,q) \geq r + k_1$. Now since $gx \in g\pi \in g\mathcal{L}_T(a,b) = \mathcal{L}_T(ga, gb)$, applying (F2) we see there are at most P possibilities for gx for a given p. Since $d(x,p) \leq r + k_1$ there are at most $(2r + 2k_1 + 1)P$ possibilites in total for the point gx. The same applies to gy. But $d(gx, gy) = d(x,y) \geq N_0$, so by (F3), at most N_0 elements of Δ can send x and y a given pair of points. We conclude that $|\Delta| \leq (2r + 2k_1 + 1)^2 N_0 P^2$ as required. ♢

In particular, this partitions the elements of Γ into elliptics and loxodromics, as described in Section 2.

Suppose $g \in \Gamma$ is loxodromic. It has two fixed points $a, b \in \partial\mathcal{G}$. Any two geodesics in $\mathcal{L}(a,b)$ are k_0-close. Moreover, $G(a,b)$ is a locally finite $\langle g \rangle$-invariant graph, and we see that $G(a,b)/\langle g \rangle$ is finite. Indeed if $c \in G(a,b)$ then $|G(a,b) \cap N(c; k_0)| \leq P$ and any geodesic in $\mathcal{L}_T(a,b)$ meets this set. We claim:

Lemma 3.4 : *If g is loxodromic, then there is some $m \leq P^2$ and some $\pi \in \mathcal{L}(a,b)$ with $g^m \pi = \pi$.*

(Note that we are not claiming that π is a tight geodesic.)

Proof : The argument follows that in [D] for hyperbolic groups. Let E be the set of directed edges in $G(a,b)$. Now $E/\langle g \rangle$ is finite, so we can label the orbits of E using a bijection to $\{1, \ldots, p\}$ for some $p \in \mathbb{N}$.

Let $\mathcal{L}_G(a,b) = \{\pi \in \mathcal{L}(a,b) \mid \pi \subseteq G(a,b)\}$. Clearly $\mathcal{L}_T(a,b) \subseteq \mathcal{L}_G(a,b)$ so this is non-empty. We say that $\pi \in \mathcal{L}_G(a,b)$ is *lexicographically least* for all vertices $x, y \in \pi$, the sequence of labels of directed edges in the segment $\pi_0 \subseteq \pi$ between x and y is lexicographically least among all geodesic segments from x to y in $G(a,b)$. Here we assume x separates from a from y in π, and that we are dealing with geodesic segments oriented from x to y. Let $\mathcal{L}_L(a,b) \subseteq \mathcal{L}_G(a,b)$ be the subset of such lexicographically least geodesics. This is $\langle g \rangle$-invariant. We claim:

(1) $\mathcal{L}_L(a,b) \neq \emptyset$.

To see this, start with any geodesic in $\mathcal{L}_G(a,b)$ and let $(x_n)_{n\in\mathbb{Z}}$ be its sequence of vertices. For each $n \in \mathbb{N}$, let π_n be a lexicographically least geodesic in $G(a,b)$ from x_{-n} to x_n. (This is necessarily geodesic in $\mathcal{G}$.) By the local finiteness of $G(a,b)$, a diagonal sequence argument gives us a subsequence $(\pi_{n_i})_i$ that converges over arbitrarily large subsets to a bi-infinite geodesic π. This must be lexicographically least, for if we could improve on some finite subsegment, π', then we could improve on any π_{n_i} containing π'.

(2) $|\mathcal{L}_L(a,b)| \le N^2$.

If not, we can find N^2+1 elements of $\mathcal{L}_L(a,b)$ which all differ in some sufficiently large compact subset of $G(a,b)$. In particular, we can find $x, y \in G(a,b)$ so that each of these N^2+1 geodesics has a subsegment connecting $N(x;k_0)$ to $N(y;k_0)$, and these subsegments are all distinct. Thus, at least two must share the same endpoints. Taking x and y far enough apart, we can also assume them to be oriented in the same direction, from x to y, say. But each of these is the unique lexicographically least geodesic segment in $G(a,b)$ connecting x to y, giving a contradiction.

Now take any $\pi \in \mathcal{L}_L(a,b)$. Its elements are permuted by $\langle g\rangle$, and so we have some $m \le P^2$ with $g^m\pi = \pi$.

Corollary 3.5 : *Suppose Γ acts on $\mathcal{G}$ with equivariant tight geodesics satisfying (F2). There is a constant, $q \in \mathbb{N}$, depending only on k and the parameters of (F2) such that if $g \in \Gamma$ is loxodromic, then $q||g|| \in \mathbb{N}$.*

Proof : Set $q = P^2!$. By Lemma 3.5, g^q translates some bi-infinite geodesic some distance $p \in \mathbb{N}$. Thus $q||g|| = ||g^q|| = p \in \mathbb{N}$. ♢

We remark that if $\mathcal{G}$ were uniformly fine, we could simply set $\mathcal{L}_T(a,b) = \mathcal{L}(a,b)$, i.e. every geodesic is deemed to be tight, and obtain the above results. This applies, in particular, to trees and uniformly locally finite graph (such as hyperbolic groups) but not to the curve graph.

4. Quasifuchsian groups.

In this section, we recall some facts and definitions from [Bo3] that are used in our proof of the main results.

Let $\mathcal{F}$ be the set of homotopy classes of non-empty connected non-annular proper subsurfaces of Σ such that each relative boundary component in Σ is essential and non-peripheral. Given $\Phi \in \mathcal{F}$ we write $\partial\Phi$ for its relative boundary in Σ. We allow for the possibilty that some component of $\Sigma \setminus \Phi$ may be a (non-peripheral) annulus, so two components of $\partial\Phi$ might correspond to the same curve of X. We write $X(\Sigma,\Phi) \subseteq X$ for the set of curves lying in (or more precisely that can be homotoped into) the surface Φ.

Now let M be a complete hyperbolic 3-manifold admitting a homotopy equivalence, $\chi : M \longrightarrow \Sigma$, with the property that parabolics in M correspond exactly to boundary components of Σ (i.e. "strictly type preserving").

We fix some appropriate Margulis constants as in [Bo3] (an additional requirement on these constants will be given in Section 5) and let $\mathcal{P}$ and $\mathcal{T}$ be the sets of Margulis cusps and Margulis tubes respectively. The boundaries of each cusp and tube are foliated by euclidean circles or *longitudes*, all of which we can assume to have length equal to a fixed constant. We write $M_P = M \setminus \bigcup_{P \in \mathcal{P}} \operatorname{int} P$ and $M_C = M \setminus \bigcup_{T \in \mathcal{T}} \operatorname{int} T$ respectively for the *non-cuspidal part* and the *thick part* of M. The injectivity radius of the thick part is bounded away from 0. Given $T \in \mathcal{T}$, $\chi(T)$ is homotopic to a curve of X which we denote by $\phi(T)$.

Let Y be the convex core of M. In this paper, we can restrict attention to the quasifuchsian case, where $Y \cap M_C$ is compact. Thus, $\mathcal{T}$ is finite, and $Y \cap M_P$ is compact. We write $\Theta = Y \cap M_P$ and $\Theta_C = Y \cap M_C$. There is a homeomorphism from $\Sigma \times [-1, 1]$ to Θ. We write $\partial_H \Theta = \partial Y \cap M_P \equiv \Sigma \times \{-1, 1\}$ and $\partial_V \Theta = Y \cap \partial M_P \equiv \partial\Sigma \times [-1, 1]$.

Let $\Phi \in \mathcal{F}$. Suppose that B is the image of a homeomorphism, $\theta : \Phi \times [-1, 1] \longrightarrow \Theta \setminus \partial_H \Theta$ such that the composition of $\Phi \hookrightarrow \Phi \times \{0\}$ with $\chi \circ \theta$ is homotopic to the inclusion of Φ into Σ. We write $\partial_H B = \theta(\Phi \times \{-1, 1\})$ and $\partial_V B = \theta(\partial\Phi \times [-1, 1])$ for the *horizontal* and *vertical* boundaries. We suppose that $\partial_V B \subseteq \bigcup_{T \in \mathcal{T}} \partial T \cup \bigcup_{P \in \mathcal{P}} \partial P$, and that each component of $\partial_V B$ is bounded two longitudes of the corresponding tube or cusp. As in [Bo3], we say that B is a *band* (or *primative strip*) if it also has the property that if $T \in \mathcal{T}$, then $T \cap B = \emptyset$ or $T \subseteq B$ or $T \cap B$ is a solid torus bounded in T by the annulus $T \cap \partial_H B$. (We can also assume that B is *unknotted* in Θ, but that need not concern us here.) We refer to Φ as the *base surface* of B, and denote it by $\phi(B)$.

By a *collared band* we mean a pair of bands, $(B, \hat{B})$ such that $\hat{B}$ is the union of three bands, $\hat{B} = B_- \cup B \cup B_+$ with $B_- \cap B_+ = \emptyset$ and $B \cap (B_- \cup B_+) = \partial_H B$ (so that B, B_-, B_+ and $\hat{B}$ all have the same base surface.) The bands B_- and B_+ are the *collars* of B.

Given a path π in Θ, we write $l(\pi)$ for the rectifiable length of $\pi \cap \Theta_C$, which refer to as the *exterior length* of π.

Given a band, B, the height, $H(B)$, is the infimum of $l(\pi)$ as π varies over paths in $B \setminus \partial_V B$ connecting the two components of $\partial_H B$. The *depth* of $Q \subseteq B$ denoted $D(Q, B)$ is the infimum of $l(\pi)$ as π varies over all paths from Q to $\partial_H B$ in $B \setminus \partial_V B$. We say that a *band* is *h-collared* if there is a collared band $(B, \hat{B})$ with $D(B, \hat{B}) \geq h$. This is equivalent to saying that $H(B_-) \geq h$ and $H(B_+) \geq h$.

Let's fix a contant H_0 to be determined later. Let $\mathcal{A}$ be the set of (outermost) bands in Θ constructed in [Bo3]. Among the properties of $\mathcal{A}$ described there, we note:

(1) The elements of $\mathcal{A}$ are mutually disjoint.

(2) Each element of $\mathcal{A}$ is H_0-collared.

(3) There is some W_0 such that for each $B \in \mathcal{A}$, each component of $\partial_H B \cap \Theta_C$ has diameter at most W_0 in Θ_C.

(4) There is a constant L_0 such that the total length of $(\partial T \cap \Theta) \setminus \bigcup \mathcal{A}$ is at most L_0.

Note that $(\partial T \cap \Theta) \setminus \bigcup \mathcal{A}$ is a union of annuli. Each annulus is bounded by two euclidean geodesiscs and its length is defined as the intrinsic euclidean distance between the two boundary components. The *total length* is thus the sum over all such annuli in ∂T. In practice, this also bounds the number of such annuli.

In the properties stated, the constant, W_0 depends only on $\kappa(\Sigma)$. We are free to choose H_0 as we please, but then L_0 depends on this as well as $\kappa(\Sigma)$. Everything will also depend on our choice of Margulis constants.

5. Closed geodesics.

We use the same notations as in Section 4. Our aim here will be to discuss how closed geodesics in Θ relate to bands in our band system.

Each curve $a \in X = X(\Sigma)$ can be realised uniquely as a closed geodesic, α, in Θ, i.e. so that $\chi(\alpha)$ is homotopic to a. We write $L(M, a)$ for its length in M. Given $t \geq 0$, write $X(M, t) = \{a \in X \mid L(M, a) \leq t\}$. Any multicurve in X can be realised as a union of closed geodesics in M.

Convention : Throughout the remainder of this paper, we use lower case Latin letters, a, b, c etc, for elements of X (or multicurves) and the corresponding Greek letters, α, β, γ etc, for the corresponding closed geodesics in M (or unions thereof).

We need also the notion of a pleated surface surface in M, see for example, [CEG]. The technicalities of this subject need not concern us here, since we can deal simply with 1-lipschitz maps, or uniformly lipschitz maps.

For us, a *pleated surface* consists of a complete finite-area hyperbolic metric, σ, on $\operatorname{int}\Sigma = \Sigma \setminus \partial\Sigma$, together with a 1-lipschitz map, $f : \operatorname{int}\Sigma \longrightarrow M$ such that $\chi \circ f$ is homotopic to the inclusion of $\operatorname{int}\Sigma$ into Σ. Indeed, it is sufficient that f be λ-lipschitz for some fixed constant λ.

Suppose that f is a pleated surface. For each $P \in \mathcal{P}$, $f^{-1}P$ is a neighbourhood of the corresponding cusp. If $T \in \mathcal{T}$, then $f^{-1}T$ is homotopic into $\phi(T)$. We can assume that f is in general position with respect to the boundaries of all such P and T. Thus, each component of $f^{-1}M_C$ is a compact subsurface. We say that such a component is *non-trivial* if it is not homotopic into any curve in Σ. After adjoining any complementary discs, such a surface is either Σ itself, or (homotopic to) an element of $\mathcal{F}$. There is combinatorial bound on the number of such surfaces. There is also a lower bound on the injectivity radius of such a surface coming from the bound in M_C. Thus, the diameter of each such surface is bounded above. From this, we conclude:

Lemma 5.1 : *If f is a pleated surface, and $x, y \in f(\Sigma) \cap M_P$, with $l(\pi)$ bounded above by some constant C_0, and such that π meets at most $\kappa(\Sigma)$ elements of $\mathcal{T}$.* $\diamondsuit$

(Recall that $\kappa(\Sigma) + 1$ bounds the number of elements in any multicurve in Σ.)

The following is a key fact concerning pleated surfaces (see for example [CEG]:

Lemma 5.2 : *If a is a multicurve in Σ, there is a pleated surface, $f : \operatorname{int}\Sigma \longrightarrow Y$ and a realisation of a in Σ such that $f|a$ maps a locally isometrically onto α.*

In other words, we can extend any geodesic multicurve in M to a pleated surface. (Recall that Y is the convex core of M.)

We note the following corollary:

Corollary 5.3 : *Suppose $a, b \in X$ and $d(a,b) \in X$ and $d(a,b) \leq n$. Suppose that $x \in \alpha \cap \Theta$ and $y \in \beta \cap \Theta$, with $l(\pi) \leq C_0 n$, and such that π meets at most $n(\kappa(\Sigma)+1)$ elements of $\mathcal{T}$.*

Proof : Let $a = a_0, a_1, \ldots, a_n = b$ be a path in X connecting a to b, and $x_i \in \alpha_i$ with $x_0 = x$ and $x_n = y$. Now $\{a_i, a_{i+1}\}$ is a multicurve, for each $i = 0, 1, \ldots, n-1$, so we can connect x_i to x_{i+1} by a path π_i using Lemma 5.1 set $\pi = \bigcup \pi_i$. $\diamondsuit$

Now fix some $t_0 > 0$, to be determined in later sections. Let $X_0 = X(M, t_0)$. We can choose our Margulis constants so that if $a \in X_0$, $P \in \mathcal{P}$ and $T \in \mathcal{T}$, then $\alpha \cap P = \emptyset$ and either $\alpha \cap T = \emptyset$ or α is the core geodesic of T.

Lemma 5.4 : *Suppose B is a band in Θ and let $\Phi = \phi(B)$ and suppose $a \in X_0 \cap X(\Sigma, \Phi)$. Suppose that there is some $\alpha \cap B$ with $D(x, B) > C_0$. If $b \in X_0$ is adjacent to a in $\mathcal{G}$, then b does not cross any curve of $\partial\Phi$.*

Proof : Note that $\{a, b\}$ is a multicurve, so there is a pleated surface $f : \Sigma \longrightarrow Y$ with $f(a) = \alpha$ and $f(b) = \beta$. Now there is some $p \in a \subseteq \Sigma$ with $f(p) = x$. Let $F \subseteq \Sigma$ be the component of $f^{-1}(B)$ containing p. The boundary of B in M is $\partial_H B \cup \partial_V B$, so $\partial F \subseteq f^{-1}(\partial_H B) \cup f^{-1}(\partial_V B)$. Now if $q \in f^{-1}(\partial_V B)$ we see, as in Lemma 5.1, that we could connect p to $f(q)$ by a path π in $B \backslash \partial_V B$ with $l(\pi) \leq C_0$. This gives the contradiction that $D(x, B) \leq C_0$. We therefore deduce that $\partial F \subseteq f^{-1}(\partial_H B)$. Thus each curve of ∂F is either homotopically trivial or peripheral in Σ or is homotopic to a curve in $\partial\Phi$. Moreover, each curve in F lies in $X(\Sigma, \Phi)$. It now follows easily (cf. [Bo4]) that F is equal to Φ up to homotopy, possibly with some discs removed. In particular, each boundary curve of Φ is also a boundary curve of F. Thus, if b were to cross such a curve it would also have to cross ∂F. Thus, β has to meet ∂T for some $T \in \mathcal{T}$. But, by the choice of Margulis constants and the assumption that $b \in X_0$, either $\beta \cap T = \emptyset$ or β is a core curve of T, either way giving a contradiction. $\diamondsuit$

We note that following addendum to Lemma 5.4:

Lemma 5.5 : *With the hypotheses if Lemma 5.4, if we assume that $b \in X(\Sigma, \Phi)$, then $\beta \subseteq B$ and $D(\beta, B) \leq D(x, B) - C_0$.*

Proof : In this case we have $\beta \subseteq f(F) \subseteq B$ as claimed. Moreover, if $z \in \beta$, as in Lemma 5.1, we can connect x to z by a path π in $B \setminus \partial_V B$ with $l(\pi) \leq C_0$. It follows that $D(x, B) \leq D(z, B) + l(\pi) \leq D(z, B) + C_0$. $\diamondsuit$

Lemma 5.6 : *Suppose $c_0 c_1 c_2 c_3$ is a tight multigeodesic in $\mathcal{G}$, where each element of each of the multicurves c_i lies in X_0. Suppose that $B \subseteq \Theta$ is a band and that $x \in B \cap \gamma_1$. Then $D(x, B) \leq 2C_0$.*

Proof : First note that $\bigcup(c_0 \cup c_2)$ and $\bigcup(c_1 \cup c_3)$ are connected in Σ and that $\bigcup(c_0 \cup c_3)$ fills Σ. Let $\Phi = \phi(B)$. Suppose $D(x, B) > 2C_0$, and let $x \in \alpha \subseteq \gamma_1$, and let $a \subseteq c_1$ be the corresponding curve. Thus a is adjacent to each of the curves in $c_0 \cup c_2$. Now applying Lemma 5.4, none of these curves can cross $\partial\Phi$. Since $\bigcup(c_0 \cup c_2)$ is connected, it either lies in, or is disjoint from Φ. But in the latter case, $c_0c_1c_2c_3$ cannot be tight at c_1, since any curve that crosses a and is contained in Φ will not cross either c_0 or c_2. Thus, c_0 and c_2 are contained in Φ. In particular, $c_2 \subseteq X(\Sigma, \Phi)$. Now applying Lemma 5.5, $\gamma_2 \subseteq B$ and if $y \in B$ then $D(y, B) > 2C_0 - C_0 > C_0$. We can thus apply the same argument again to show that c_1 and c_3 are contained in Φ. In other words $\bigcup(c_1 \cup c_3) \subseteq \Phi$, contradicting the fact that this fills Σ. $\diamondsuit$

6. Bounding numbers of curves.

Let $\mathcal{A}$ be a system of bands satisfying the conditions laid out in Section 4. In this section, we show:

Lemma 6.1 : *Given $C \geq 0$, there is some N such that if $x \in \Theta \setminus \bigcup \mathcal{A}$, then there are at most N curves $a \in X_0$ such that $\alpha \cap \mathcal{A} = \emptyset$ and x can be connected to α in Θ by a path π with $l(\pi) \leq C$.*

Recall that $\alpha \subseteq \Theta$ is the closed geodesic realising a, and that X_0 is the set of curves, a, such the α has length at most some fixed constant t_0 to be determined later. As in Section 5, we can assume that we have fixed the Margulis constants such that α is either a core curve of some element of $\mathcal{T}$ or else lies in Θ_C. Here, C will depend on $\kappa(\Sigma)$, t_0, the Margulis constants, and the constants W_0, H_0 and L_0 featuring in the properties of $\mathcal{A}$. All of these constants will ultimately depend only on $\kappa(\Sigma)$.

We shall denote the induced path metric on Θ_C by ρ.

Before proving Lemma 6.1, we first note that we can assume that $\pi \cap \partial_V B = \emptyset$ for all $B \in \mathcal{A}$, since we can replace it by a path π' with the same endpoints having this property, and with $l(\pi') \leq C'$ where C' depends only on C. To achieve this, note that there is a bound on the number of elements of $\mathcal{A}$ that π can enter. Moreover, the boundary of a band, $B \in \mathcal{A}$ in Θ is contained in $\partial_H B \cup \partial_V B$. Now $\partial_H B$ meets at most $2(\kappa(\Sigma) + 1)$ elements of $\mathcal{T}$ and each component of $\partial_H B \setminus \bigcup \mathcal{T}$ is assumed to have diameter at most W_0 in Θ_C. Thus any two points of $\partial_V B \cup \partial_H B$ can be connected by a path not meeting $\partial_V B$ anywhere else and of bounded exterior length. We can now easily modify π so as to avoid $\partial_V B$ completely.

We now construct a graph Δ as follows. Fix some $\epsilon > 0$ so that 3ϵ is smaller than both the injectivity radius of Θ_C, and the minimal distance between distinct elements of $\mathcal{T}$. For each $T \in \mathcal{T}$ let V_T be an ϵ-net in $\partial T \cap \Theta$ (i.e. a maximal ϵ-separated subset of $\partial T \cap \Theta$). We extend $\bigcup_{T \in \mathcal{T}} V_T$ to an ϵ-net, V, in Θ_C. Let Δ be the graph with vertex set V and with $x, y \in V$ adjacent if $\rho(x, y) \leq 3\epsilon$. We can define a map $\theta : \Delta \longrightarrow \Theta_C$ which extends the inclusion of V into Θ_C, by mapping each edge of Δ to a geodesic segment.

Given $T \in \mathcal{T}$, write $\Omega(T) = \partial T \cap \Theta \setminus \bigcup \mathcal{A}$, and let Υ_T be the complete graph on

$V \cap \Omega(T)$. Note that Property (4) of $\mathcal{A}$ means that the cardinality of $V(\Upsilon_T) = V \cap \Omega(T)$ is uniformly bounded (in terms of L_0 and ϵ). Let Υ be the graph $\Delta \cup \bigcup_{T \in \mathcal{T}} \Upsilon_T$.

Now as discussed in [Bo3], the degree of each vertex of Δ and hence of Υ is bounded.

Let x and a be as in the hypotheses of Lemma 6.1. We can suppose that $x \in V(\Delta)$. Either α lies in Θ_C or it is the core curve of some element of $\mathcal{T}$.

Suppose first that $\alpha \subseteq \Theta_C$. We can find a closed curve, $q(a)$, of length at most t_0/ϵ in Δ such that $\theta(q(a))$ is homotopic to α in Θ_C, hence in M. In particular, $a \in X$ is determined by $q(a)$. We also have a path π connecting x to α with $l(\pi)$ bounded, which we can suppose intersects each $T \in \mathcal{T}$, if at all, only in points of $\Omega(T)$. We therefore see that π similarly gives rise to a path $p = p(\pi)$ in Υ, of bounded length, connecting x to $q(a)$. Since all vertices of Υ have bounded degree, this places a bound on the number of possibilities for p, hence for $q(a)$ and hence for a.

The argument in the case where a is the core curve of some $T \in \mathcal{T}$ is similar. In this case, we choose any point, $q(a) \in V_T$. This determines T and hence a. We can connect x to $q(a)$ by a path of bounded length as before, and this places a bound on the number of such a.

These bounds are entirely definable in terms of the constants C, t_0, $\kappa(\Sigma)$, the Margulis constants, and the constants featuring in the properties of $\mathcal{A}$. This proves Lemma 6.1.

Corollary 6.2 : *Given $n \geq 0$, there is some $N \geq 0$ satisfying the following. Suppose $a \in X_0$ with $\alpha \cap \bigcup \mathcal{A} = \emptyset$. Then there are at most N elements, b, of X_0 with $\beta \cap \bigcup \mathcal{A} = \emptyset$ and with $d(a,b) \leq n$.*

Proof : Let $x \in \alpha$. By Lemma 5.3 any point of β can be connected to x by a path π in Θ with $l(\pi)$ bounded. $\diamondsuit$

7. Proofs.

In this section, we give proofs of the main results of Section 1.

Before starting, there are two additional ingredients we need to cite. The first of these are the "a-priori" bounds on the lengths of closed geodesics corresponding to curves in tight geodesics in the curve complex. These are proven in [Mi] (at least with the stronger condition of tightness given there). More direct proofs are given in [Bo4], and the two results, Theorems 7.1 and 7.2, are quoted directly from there.

As remarked earlier, the bounds given by these results are not constructive. Modulo these, however, all other bounds are constructive, so it would be interesting to attempt to circumvent the geometric limit arguments featuring in [Mi] and [Bo4] that give rise to this ineffectiveness.

Theorem 7.1 : *Given $\epsilon > 0$, there is some $t \geq 0$ depending only on ϵ and $\kappa(\Sigma)$ such that $a, b \in X(\Sigma)$ with $L(M,a) \leq \epsilon$ and $L(M,b) \leq \epsilon$ and if $c \in X(\Sigma)$ is a vertex in a tight geodesic from a to b (i.e. $c \in V(\mathcal{G}) \cap G(a,b)$), then $L(M,c) \leq t$.* $\diamondsuit$

Theorem 7.2 : *There is some $k_2 \geq 0$ depending only on $\kappa(\Sigma)$ such that if $\epsilon > 0$, there is some $t \geq 0$, depending ϵ and $\kappa(\Sigma)$ such that if $a, b \in X(\Sigma)$ with $L(M, a) \leq \epsilon$ and $L(M, b) \leq \epsilon$, if $r \geq 0$ and $c \in X$ lies on a tight geodesic from $N(a; r)$ to $N(b; r)$ with $d(c, \{a, b\}) \geq r + k_2$ then $L(M, c) \leq t$.* ◊

The other ingredient we need is the existence of such a manifold M.

Proposition 7.3 : *Given any $\epsilon > 0$ and any $a, b \in X(\Sigma)$ there is a complete manifold M hyperbolic 3-manifold and a strictly type preserving homotopy equivalence $\chi : M \longrightarrow \Sigma$, such that $L(M, a) \leq \epsilon$ and $L(M, b) \leq \epsilon$. Moreover, we can assume M to be quasifuchsian (geometrically finite without accidental parabolics).*

(In fact, any $\epsilon > 0$ depending only on $\kappa(\Sigma)$ would do for our purposes.)

Proof : There are a number of ways to see this. For example, the deformation space of quasifuchsian groups is homeomorphic to the product of two Teichmüller spaces, where each coordinate corresponds to the conformal structure of one of the surfaces at infinity (i.e. a quotient of the discontinuity domain). Moreover, there is a universal bound on Teichmüller distance between such a surface and the boundary of the corresponding convex core boundary [EM]. By choosing the structure so the extremal lengths of a and b are each sufficiently small in one of the surfaces, the lengths of the corresponding curves in M will be arbitrarily small.

Alternatively, using hyperbolisation, one could find a geometrically finite manifold in which a and b are both parabolic, and then deform slightly. ◊

We can now set about proving Theorems 1.1 and 1.2. We fix any constant $\epsilon > 0$. Let $t_0 = t$ be the constant given by Theorems 7.1 and 7.2 (taking the maximum of the two). This depends only on $\kappa(\Sigma)$. We set $X_0 = X(M, t_0)$. We can now fix the Margulis constants, as in Section 5, so that if $\alpha \in X_0$ and $P \in \mathcal{P}$ then $\alpha \cap P = \emptyset$, and if $T \in \mathcal{T}$ either $\alpha \cap T = \emptyset$ or α is the core curve of T. These Margulis constants now give us a constant W_0, as in the description of the band system $\mathcal{A}$ is Section 4. We also have a constant C_0 described by Lemma 5.1. Fix any $H_0 > 2C_0$. This gives us another constant, L_0, as in property (4) of the band system.

Let k_0 be the constant of Lemma 2.1, given the hyperbolicity constant, k, of the curve graph $\mathcal{G}(\Sigma)$. Let N be the constant of Corollary 6.2 given $n = k_0$.

Proof of Theorem 1.1 : Let $a, b \in X$ and let $c \in X$ lie in a tight geodesic from a to b. If $d(a, b) = 2$, then c is a boundary curve of the subsurface of Σ filled by $\alpha \cup \beta$, and there is clearly a bound on the number of possibilities for such. We can therefore suppose that $d(a, b) \geq 3$, and so c lies in a tight geodesic of length 4. Let M be as given by Proposition 7.3, and let $\mathcal{A}$ be the band system described in Section 4. If $B \in \mathcal{A}$, then by Property (2), there is a collared band, $(B, \hat{B})$ with $D(B, \hat{B}) \geq H_0 > 2C_0$. Thus, by Lemma 5.6, $\gamma \cap B = \emptyset$. In other words, $\gamma \subseteq \Theta \setminus \bigcup \mathcal{A}$.

Now if $c' \in G(a, b) \cap N(c; k_0)$ we similarly have $\gamma' \subseteq \Theta \setminus \bigcup \mathcal{A}$. By Corollary 6.2, there are at most N possibilities for such a c', and so we can set $K_0 = N$. ◊

Proof of Theorem 1.2 : The argument is essentially the same, using Theorem 7.2 in place of Theorem 7.1. By setting $k_1 = k_0 + k_2$ if $d(c, \{a, b\}) \geq r + k_1$ and $d(c, c') \leq k_0$, then $d(c', \{a, b\}) \geq r + k_2$. If $c' \in G(a, b; r)$, then Theorem 7.2 tells us that $L(M, c') \leq t_0$, and the argument proceeds as before. ♢

To deduce Theorem 1.3, we need the following observation:

Lemma 7.4 : *There is some n depending only on $\kappa(\Sigma)$ such that if $a, b \in X$ with $d(a, b) \geq 3$, then $|\operatorname{stab} a \cap \operatorname{stab} b| \leq n$.*

Proof : (cf. [BeF]) If we realise $a, b \in \Sigma$ so as to minimise the number of intersections, then $a \cup b$, as a combinatorial graph embedded in Σ, is determined up to isotopy. Since $d(a, b) \geq 3$, this graph fills Σ. One can construct from it a cononical singular metric on the surface Σ' consisting of Σ with each boundary component collapsed to a point (cf. [Bo2]). Now $|\operatorname{stab} a \cap \operatorname{stab} b|$ acts by isometry on this surface, and hence by conformal automorphism on the corresponding punctured Riemann surface. But it is well-known that the cardinality of such an automorphism group is bounded in terms of $\kappa(\Sigma)$. ♢

Proof of Theorem 1.3 : Let $\mathcal{L}_T(a, b)$ be the set of tight geodesics from a to b as defined in Section 3. Theorem 1.2 tells us that the collection $\mathcal{L}_T = (\mathcal{L}_T(a, b))_{a, b \in V(\mathcal{G})}$ satisfies Property (F2) of Section 3. Lemma 7.4 tells us that the action of $\Gamma(\Sigma)$ on $\mathcal{G}(\Sigma)$ satisfies Property (F3). Lemma 3.3 now tells us that the action of $\Gamma(\Sigma)$ on $\mathcal{G}(\Sigma)$ is acylindrical as required. ♢

Proof of Theorem 1.4 : Since $\mathcal{L}_T$ satisfies (F2), we can apply Lemma 3.4. ♢

Finally Corollary 1.5 follows from Corollary 3.5.

References.

[BeF] M.Bestvina, K.Fujiwara, *Bounded cohomology of subgroups of the mapping class groups* : Geometry & Topology **6** (2002) 69–89.

[Bo1] B.H.Bowditch, *Relatively hyperbolic groups* : preprint, Southampton (1997).

[Bo2] B.H.Bowditch, *Intersection numbers and the hyperbolicity of the curve complex* : preprint, Southampton (2002).

[Bo3] B.H.Bowditch, *Systems of bands in hyperbolic 3-manifolds* : preprint, Southampton (2003).

[Bo4] B.H.Bowditch, *Length bounds on curves arising from tight geodesics* : preprint, Southampton (2003).

[CEG] 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.

[D] T.Delzant, *Sous-groupes distingués et quotients des groupes hyperoliques* : Duke Math. J. **83** (1996) 661–682.

[EM] D.B.A.Epstein, A.Marden, *Convex hulls in hyperbolic space, a theorem of Sullivan, and measured pleated surfaces* : 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) 113-253.

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

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

[Hare] J.L.Harer, *The virtual cohomological dimension of the mapping class groups of orientable surfaces* : Invent. Math. **84** (1986) 157–176.

[Harv] 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.

[I] N.V.Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces* : Internat. Math. Res. Notices (1997) 651–666.

[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.

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

[S] Z.Sela, *Acylindrical accessibility for groups* : Invent. Math. **129** (1997) 527–565.