The Cannon-Thurston map for punctured-surface groups.

Brian H. Bowditch
Faculty of Mathematical Studies, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
B.H.Bowditch@maths.soton.ac.uk

[First draft: December 2001. Revised: February 2002]

0. Introduction.

In this paper, we describe the Cannon-Thurston map for a punctured-surface group of finite type having a strictly type-preserving action on hyperbolic 3-space, where we assume a positive lower bound on the translation distances of loxodromics. The Cannon-Thurston map was originally described for closed surfaces in [CannT]. Various generalisations are described in [Min1,Mit1,Mit2,AlDP,Mc,CannD,Cann].

Let Σ be a compact surface with boundary components, $(C^m)_{m\in\mathcal{P}}$ indexed by a finite set $\mathcal{P}$ which we shall assume to be non-empty. We assume that the Euler characteristic, $\chi = \chi(\Sigma)$ is negative. Let $\Gamma = \pi_1(\Sigma)$. We refer to the subgroups of Γ supported on boundary curves as *peripheral*.

We write $(\mathbf{H}^\nu, \rho)$ for ν-dimensional hyperbolic space, and $\partial\mathbf{H}^\nu \cong S^{\nu-1}$ for its boundary. A discrete and faithful action of Γ on $\mathbf{H}^\nu$ is *strictly type-preserving* if the maximal parabolic subgroups are precisely the peripheral subgroups of Γ. (This is commonly referred to as "type-preserving with no accidental parabolics".) We write $\Pi = \Pi(\Gamma)$ for the set of parabolic points, and $\Lambda = \Lambda(\Gamma)$ for the limit set — the closure of Π in $\partial\mathbf{H}^\nu$. Let $N = N(\Gamma) = \mathbf{H}^\nu/\Gamma$. We write $\mathrm{inj}(N)$ for half the length of the shortest closed geodesic in N, or, equivalently, half the shortest translation distance of a loxodromic element of Γ. This is the "injectivity radius away from cusps".

If $\nu = 2$, we can always embed Σ in N so that each boundary curve in a horocycle of fixed length. Thus N retracts onto Σ, and $\mathrm{inj}(N) > 0$. In this case $\Lambda = \partial\mathbf{H}^2$ is a circle. Up to equivariant homeomorphism, this "circle at infinity" is well defined independently of the choice of N. In fact, $\partial\mathbf{H}^2$ can be canonically identified with the boundary, $\partial\Gamma$, of Γ as a relatively hyperbolic group (see [Bow2]). We need not worry about the formal definition of such a boundary here, but merely regard this as a convenient notation. We can identify Π as a dense subset of $\partial\Gamma \equiv \partial\mathbf{H}^2$.

The case $\nu = 3$ is considerably more subtle and much studied. It is known that N is always topologically finite (the "tameness conjecture" [Bon]). Indeed one can properly embed $\Sigma\times\mathbf{R}$ in N so that each boundary component $C^m\times\mathbf{R}$ gets mapped to a horocyclic cylinder. However, the geometry may be complicated, and the restriction that $\mathrm{inj}(N) > 0$ is non-trivial, indeed in some sense "non-generic".

The structure on either of the two ends of $\Sigma\times\mathbf{R}$ may be either "geometrically finite" or "simply degenerate". Geometrically finite ends will be trivial from our point of view, so only the case where at least one end is simply degenerate is of interest. The "doubly degenerate" case arises precisely when $\Lambda = \partial\mathbf{H}^3$. Important examples of this case are infinite cyclic covers of hyperbolic 3-manifolds fibring over the circle, see [T2,O].

In this paper, we show:

Theorem 0.1 : *If* $\mathrm{inj}(N) > 0$, *then there is a continuous* Γ*-equivariant map from* $\partial \mathbf{H}^2$ *to* $\partial \mathbf{H}^3$.

This map is necessarily the identity on Π and hence unique. We refer to it as the *Cannon-Thurston map* $\omega : S^1 \longrightarrow S^2 \equiv \partial \mathbf{H}^3$. Note that $\Lambda = \omega(S^1)$, and so we deduce immediately that the limit set is locally connected in this case.

We can also give a topological description of the map ω. Note that each simply degenerate end has associated to it an ending lamination as in [Bon,T1]. Such a lamination determines a closed equivalence relation on $S^1 \equiv \partial \mathbf{H}^2$. In the doubly degenerate case, we take the transitive closure of the two equivalence relations arising in this way. We shall see:

Theorem 0.2 : *The Cannon-Thurston map is the quotient of the circle by the equivalence relation arising from the ending laminations.*

A more precise statement of this is given by Theorem 9.1.

Here we have assumed that $\Pi \neq \emptyset$. The case where $\Pi = \emptyset$ is dealt with in [CannT,Min1]. Another approach is described in [Mit1,Mit2]. The case where $\Pi \neq \emptyset$ would seem to call for additional techniques. In particular we use combinatorial arguments of the type developed in [Bow2]. (It is stated in [CannT] that the results proven there "are true when [the manifold that fibres over the circle] has cusps", though no reason is supplied.)

Of particular interest are cyclic covers of manifolds fibring over the circle. In this case, we are able to give a simplified proof (though at the cost of making appeal to other results) — see Section 6.

In [AlDP], the special case of the figure-8-knot complement — a punctured-torus bundle over the circle — was analysed in some detail. In particular, they prove Theorems 0.1 and 0.2 for this manifold. Moreover a number of interesting questions relating to the Cannon-Thurston map are raised. More general once-punctured-torus bundles are considered in [CannD], and a proof of the above results in this situation will appear in [Cann].

We note that a corollary of Theorem 0.1 and the version without parabolics [CannT,Min1] together a result of Anderson and Maskit [AnM] is the following:

Theorem 0.3 : *Let* G *be a finitely generated kleinian group and suppose that there is a positive lower bound on the translation lengths of loxodromic elements of* G. *If the limit set* $\Lambda(G)$ *is connected, then it is locally connected.*

Here a "kleinian group" is any group acting properly discontinuously on $\mathbf{H}^3$. The "translation length" of a loxodromic is the distance it translates its axis. Note that Theorem 0.3 includes the case of a geometrically finite group [AnM]. It is an open question as to whether Theorem 0.3 remains true if the condition on loxodromics is dropped. It is known in full generality for once-punctured-torus groups [Min3,Mc]. We explain how

Theorem 0.3 is deduced in Section 10.

Although we shall be referring to hyperbolic manifolds, apart from Theorem 0.3, our arguments do not make essential use of constant curvature. Thus, the results of Section 7 appropriately interpreted, apply to any Gromov hyperbolic space. Moreover the essential points of Section 8 use only bounded local geometry of our manifold M. We shall be quoting results of Thurston and Bonahon, but as noted in [Cana] for example, these also hold more generally. Thus, although we shall not do so explicitly, we could state and prove Theorems 0.1 and 0.2 with hyperbolic 3-space replaced by a 3-dimensional Hadamard manifold of pinched negative curvature.

We remark that these results fit into a much broader project of understanding the ends of hyperbolic 3-manifolds. In view of the result of Minsky [Min2] on the ending lamination conjecture, much effort now concentrates on understanding the case where $\mathrm{inj}(N) = 0$, or where one allows for cusps. One goal already mentioned concerns the local connectivity of limit sets. It is hoped that a continuation of the present work will deal with more general questions such as that discussed in [K], again with $\mathrm{inj}(N) > 0$, but where one allows for parabolics.

The structure of this paper is as follows. In Section 1, we give a simple criterion for extending maps between compacta. In sections 2 and 3, we discuss and develop the theory of "fine" hyperbolic graphs from [Bow2]. In Sections 4 and 5, we discuss "stacks" of such graphs and define a kind of Cannon-Thurston map in this context. In order to apply this, we associate such a stack to our manifold N. Most of the remaining work involves showing that a stack arising in this way is hyperbolic. For manifolds fibring over the circle, this can be achieved relatively simply (Section 6). For the general case, one needs some technical facts relating to systems of horoballs (Section 7) and the Thurston-Bonahon-Minsky et al. theory of ends of 3-manifolds (Section 8). The proof of Theorem 0.1 is completed in Section 8, and in Section 9, we prove Theorem 0.2. Finally we briefly discuss Theorem 0.3 in Section 10.

We make a brief remark about a terminological convention. Throughout the paper, we refer to various constants or functions as being "uniform". By this we mean that they depend only on other constants introduced in the hypotheses or implicit in the set-up. Ultimately there are only two constants to be considered — the lower bound, η, on $\mathrm{inj}(N)$, and the complexity of our surface Σ, conveniently measured by $\chi(\Sigma)$. (There is also the Margulis constant, η_0, though since we are only interested in dimensions 2 and 3, this can be fixed once and for all.) All the other constants can thus, in principle, be expressed in terms of η and χ. There is one point where we use a limiting argument, or more precisely, precompactness (Proposition 8.8). However, again, this can in principle be made explicit by reducing to a combinatorial set-up as we shall explain.

I thank Warren Dicks for drawing my attention to the problem of constructing the Cannon-Thurston map in this context, and for his helpful comments on this paper. I was further inspired by the results of the paper [AlDP].

1. A simple lemma.

We shall define the Cannon-Thurston map using the following lemma.

Lemma 1.1 : *Let X, Y be compact metrisable spaces. Let $f : \Pi \longrightarrow Y$ be a function from a subset $\Pi \subseteq X$ to Y. Suppose that $\mathcal{O}$ is a base of open sets in X with the property that if $(O_n)_{n \in \mathbf{N}}$ is any decreasing sequence in $\mathcal{O}$ with $|\bigcap_n O_n| = 1$ then $\left|\bigcap_n \overline{f(\Pi \cap O_n)}\right| = 1$. Then f extends to a continuous map $f : X \longrightarrow Y$.*

(Here $\overline{A}$ denotes the closure of A.) Note that Π is necessarily dense in X, so the extension is unique.

Proof : Given $x \in X$, choose any decreasing sequence, $(O_n(x))_n$, in $\mathcal{O}$ with $\bigcap_n O_n(x) = \{x\}$. Define $f(x) \in Y$ by $\{f(x)\} = \bigcap_n \overline{f(\Pi \cap O_n(x))}$. Clearly this is consistent on Π. To see that f is continuous (hence well-defined) let U be an open neighbourhood of $f(x)$ in Y. Since Y is compact, there is some $n \in \mathbf{N}$ such that $\overline{f(\Pi \cap O_n(x))} \subseteq U$. We claim that $f(O_n(x)) \subseteq U$, for if $y \in O_n(x)$, there is some $m \in \mathbf{N}$ such that $O_m(y) \subseteq O_n(x)$, and so $f(y) \in \overline{f(\Pi \cap O_m(y))} \subseteq \overline{f(\Pi \cap O_n(x))} \subseteq U$ as required. $\diamondsuit$

We shall ultimately apply this in the case where $X = S^1$ and $\mathcal{O}$ is the set of open intervals.

2. Fine hyperbolic graphs.

In this section, we review some of the theory of fine hyperbolic graphs as developed in [Bow2], and their connections with systems of horoballs in $\mathbf{H}^\nu$.

Let K be a connected graph. We write $V = V(K)$ and $E = E(K)$ for the sets of vertices and edges respectively. We write $d = d_K$ for the combinatorial path metric on K (that assigns unit length to each edge).

Given $x \in V$, write $E(x) = E(K, x)$ for the set of edges incident on x. If $e, f \in E(x)$, write $\angle(e, f) = \angle_x(e, f) = d_{K \setminus \{x\}}(y, z)$ where y and z are the other endpoints of e and f. Thus $\angle_x$ defines a metric on $E(x)$ with values in $\mathbf{N} \cup \{\infty\}$. We refer to $\angle(e, f)$ as the *angle* between e and f.

If α is a path or cycle in K, and x is a vertex of α other than an endpoint, then we write $\angle_x(\alpha)$ for the angle between the incident edges of α. If α and β are two paths emanating from the same point x, we write $\angle_x(\alpha, \beta) = \angle_x(\alpha \cup \beta)$. We abbreviate this to $\angle(\alpha, \beta)$ if there is no ambiguity. We adopt the convention that $\angle_x(\alpha, \beta) = \infty$ if either of α or β is the trivial path $\{x\}$. Note that if γ is a circuit of length n, and x is a vertex thereof, then $\angle_x(\gamma) \leq n - 2$. Similarly, if α, β are arcs of length at most n both connecting x to a different vertex, y, then $\angle_x(\alpha, \beta) \leq 2n - 2$.

Definition : We say that K is *fine* if for all $x \in V$, the metric $\angle_x$ on $E(x)$ is locally finite.

This means that for all $n \in \mathbf{N}$ and all $e \in E(x)$, the set $\{f \in E(x) \mid \angle_x(e, f) \leq n\}$ is finite.

Various equivalent formulations of fineness are given in [Bow2]. In particular, we note that there are only finitely many circuits of a given length containing any given edge. Equivalently, any two points of V are connected by only finitely many arcs of (at most) any given length.

Suppose now that K is both fine and hyperbolic (in the sense of Gromov [Gr1], see also [GhH,Bow1]). Let ∂K be its usual "Gromov" boundary (i.e. parallel classes of geodesic rays). Let ΔK be the formal disjoint union, $\Delta K = V(K) \sqcup \partial K$. In [Bow2], we defined a natural topology on ΔK which is compact and hausdorff (indeed metrisable, at least in the cases of interest to us).

Before describing this topology, we note that if $x, y \in \Delta K$ then x, y are connected by at least one geodesic (compact, bi-infinite, or a ray, depending on whether or not $x, y \in V$). Any two such geodesics, α, β remain a bounded distance apart, and $\angle_x(\alpha, \beta)$ is bounded. In view of this, the choice will not usually matter to us, and we denote by $[x, y]$ some choice of geodesic.

Suppose now that $x \in \Delta K$, and $(y_n)_n$ is a sequence in ΔK.

(1) If $x \in \partial K$, then $y_n \to x$ in ΔK if and only of $d(z, [x, y_n]) \to \infty$ for some (hence every) $z \in V$.

(2) If $x \in V$, then $y_n \to x$ in ΔK if and only if $\angle_x(e, [x, y_n]) \to \infty$ for some (hence every) $e \in E(x)$.

Note that V is dense in ΔK.

Suppose that $L \subseteq K$ is a subgraph with vertex set $V(L) = V(K) = V$. Then L is necessarily fine. The inclusion is a quasi-isometry if any only if any two K-adjacent vertices are a bounded distance apart in L. In this case, L, is also hyperbolic. Moreover, from [Bow2], we have:

Lemma 2.1 : *The identity on V extends to a homeomorphism from ΔL to ΔK.* $\diamondsuit$

We next relate these notions to systems of disjoint horoballs in $\mathbf{H}^\nu$.

The following construction is discussed in a more general context in [Bow2]. Thus a disjoint horoball system in $\mathbf{H}^\nu$ (or a convex subset thereof) is a special case of a separated system of uniformly quasiconvex subsets of a Gromov hyperbolic space.

Suppose that $\Pi \subseteq \partial \mathbf{H}^\nu$ is any subset, and let $\overline{\Pi}$ denote its closure and let $\mathrm{hull}(\overline{\Pi})$ be the convex hull of $\overline{\Pi}$ in $\mathbf{H}^\nu$. Suppose to each $p \in \Pi$, we associate a closed horoball, $B(p)$, about p, such that if $p \neq q$, then $B(p) \cap B(q) = \emptyset$. Let $L(\infty)$ be the complete graph with vertex set Π, and for any $t \geq 0$, let $L(t)$ be the subgraph with vertex set Π where $p, q \in \Pi$ are deemed to be adjacent if $\rho(B(p), B(q)) \leq t$. We refer to $L(t)$ as the *t-nerve* of the system $(B(p))_{p \in \Pi}$.

It is fairly easy to see [Bow2] that:

Lemma 2.2 : *For all $t \geq 0$, the graph $L(t)$ is fine.* $\diamondsuit$

(We note that in the more general context of [Bow2] we imposed a "separation" condition on the system $(B(p))_p$. However, in the present situation, this could always be arranged by replacing each horoball by a uniformly smaller horoball about the same point.)

Definition : Given $r \geq 0$, we say that $(B(p))_p$ is *r-quasidense* if every point of $\mathrm{hull}(\overline{\Pi})$ lies a distance at most r from $\bigcup_{p \in \Pi} B(p)$.

We note [Bow2]:

Lemma 2.3 : *Given $r, u \geq 0$, there is some $k \geq 0$ such that if $(B(p))_{p \in \Pi}$ is an r-quasidense system of horoballs and if L is a graph satisfying $L(2r+1) \subseteq L \subseteq L(u)$, then L is k-hyperbolic. Moreover there is a homeomorphism from $\overline{\Pi}$ to ΔL which is the identity on $\Pi = V(L)$.* ♢

The idea is to associate to any geodesic in L, a uniform quasigeodesic in $\mathbf{H}^\nu$, and apply hyperbolicity of $\mathbf{H}^\nu$. Variations on this construction will be used in Section 7.

3. Escaping sets.

In this section, we develop further some of the theory of fine hyperbolic graphs. In particular, we are interested in systems of uniformly quasiconvex sets wandering off to infinity in some sense or other.

Let K be a fine hyperbolic graph with compact boundary ΔK.

Definition : A sequence of subsets, $(W_n)_n$, of V is *escaping* if $d(x, W_n) \to \infty$ for some (hence every) $x \in V$.

We shall also apply this definition to a subgraph, L, of K with respect to the metric, d_L (possibly taking infinite values).

Suppose that $(\alpha_n)_n$ is a sequence of paths emanating from the same point $x \in V$. We say that $(\alpha_n)_n$ is *edge-escaping* if only finitely many α_n start with a given edge of $E(x)$. By fineness, this is equivalent to asserting that $\angle_x(e, \alpha_n) \to \infty$ for some (hence every) $e \in E(x)$.

In fact, we shall only be applying this definition either in the case where the α_n are all geodesic or the case where all are arcs of bounded length (in such cases, the condition turns out to be equivalent to "evading" as defined later). Note that in this situation, we can always pass to a subsequence $(\alpha_{n_i})_i$ such that $\alpha_{n_i} \cap \alpha_{n_j} = \{x\}$ whenever $i \neq j$. Note also that if $(\alpha_n)_n$ is an edge-escaping sequence of arcs of bounded length, and β_n is a geodesic with the same endpoints, then $(\beta_n)_n$ is also edge-escaping.

Let $W \subseteq V$. Write $J(W)$ for the union of all geodesics connecting two points of W. We say that W is *r-quasiconvex* if $J(W)$ is contained in an r-neighbourhood of W. We write $\overline{W}$ for the closure of W in ΔK. The following is a simple consequence of hyperbolicity and the description of the topology on ΔK.

Lemma 3.1 : *Suppose that W is quasiconvex and that $x \in \partial K \subseteq \Delta K$. Let α be a geodesic ray tending to x. Then $x \in \overline{W}$ if and only if α remains a bounded distance from W. Moreover, in this case, from some point onward, α remains a uniformly bounded distance from W.* ♢

Here, and henceforth, "uniformly" means depending only on the parameters of the given situation — hyperbolicity and quasiconvexity in this case. (We did not introduce any explicit parameters in our definition of fineness.)

Lemma 3.2 : *Let $W \subseteq V$ be quasiconvex, and suppose $x \in V \setminus W$. The following are equivalent:*

(1) $x \in \overline{W}$.

(2) There is an edge-escaping sequence of geodesics from x to W.

(3) There is an edge-escaping sequence of arcs of (uniformly) bounded length from x to W.

(4) There is an edge-escaping sequence of geodesics of (uniformly) bounded length from x to W.

Corollary 3.3 : *If W is quasiconvex and $x \in V \cap \overline{W}$, then $d(x, W)$ is uniformly bounded.*

Proof : By conditions (3) or (4) of Lemma 3.2. In fact a direct proof of Corollary 3.3 will form part of that of Lemma 3.2. $\diamondsuit$

Proof of Lemma 3.2 :

(1) $\Leftrightarrow$ (2) : If $y_n \in W$, then $y_n \to x$ in ΔK if and only if $[x, y_n]$ is edge-escaping.

(2) $\Rightarrow$ (3) : First we show that $d(x, W)$ is unformly bounded. To this end, choose $y \in W$ with $d(x, y) = d(x, W)$. A simple consequence of quasiconvexity is that y lies a uniformly bounded distance from any geodesic connecting x to W. Let $(\alpha_n)_n$ be a sequence of such geodesics given by hypothesis (2). Let β_n be a shortest path from y to α_n, meeting α_n at z_n, say. Since $(\alpha_n)_n$ is edge-escaping, we must have $z_n = y$ for all but finitely many n. Thus, $d(x, y) \leq \text{length}(\beta_n)$ which is uniformly bounded.

Suppose now that $a_n \in \alpha_n$. Since $d(x, y)$ is uniformly bounded, by hyperbolicity, a_n lies a bounded distance from $[y, y_n]$. Hence, by quasiconvexity, there is some $b_n \in W$ with $d(a_n, b_n)$ uniformly bounded, by some constant, r, say. Let c_n be the first point at which $[b_n, c_n]$ meets α_n, and let γ_n be the subsegment from b_n to c_n. Let δ_n be the arc that runs along α_n from x to c_n, then along γ_n from c_n to b_n. Note that if $d(x, z_n) > r$, then $c_n \neq x$ and so α_n and δ_n have a common initial edge.

We now construct a sequence of arcs, $(\epsilon_n)_n$, as follows. If $\text{length}(\alpha_n) \leq r$, set $\epsilon_n = \alpha_n$. If $\text{length}(\alpha_n) > r$, let $z_n \in \alpha_n$ with $d(x, z_n) = r + 1$, and set $\epsilon_n = \delta_n$ as above. We see that $(\epsilon_n)_n$ is edge-escaping, and that $\text{length}(\epsilon_n) \leq 2r + 1$ is uniformly bounded.

(3) $\Rightarrow$ (4) : Let $(\alpha_n)_n$ be the sequence of arcs given by (3), and let β_n be a geodesic with the same endpoints as α_n. Since $\angle_x(\alpha_n, \beta_n)$ is bounded, $(\beta_n)_n$ is also edge-escaping.

(4) $\Rightarrow$ (2) : Trivial. $\diamondsuit$

Corollary 3.4 : *If $W \subseteq V$ is quasiconvex, then any geodesic connecting any two points of $\overline{W}$ lies a bounded distance from W.*

Proof : This is an easy consequence of hyperbolicity and quasiconvexity, after applying Lemma 3.1 or Corollary 3.3, depending on whether the respective endpoints lie in ∂K or in V. $\diamondsuit$

In what follows, we shall be assuming that $(W_n)_n$ is a decreasing sequence of uniformly quasiconvex subsets of V, i.e. $W_{n+1} \subseteq W_n$ for all n.

Lemma 3.5 : *Suppose that $x \in V \setminus \bigcup_n W_n$. Then $x \in \bigcap_n \overline{W}_n$ if and only if there is an edge-escaping sequence of arcs, $(\alpha_n)_n$, of uniformly bounded length with α_n connecting x to W_n.*

Proof : By a diagonal sequence argument, applying Lemma 3.2 part (3). $\diamondsuit$

Lemma 3.6 : *Suppose that there is some $r \geq 0$ and $x \in V$ such that $d(x, W_n) \leq r$ for all n. Then there is some $y \in V \cap \bigcap_n \overline{W}_n$ with $d(x,y) \leq r$.*

Proof : Let α_n be a shortest path from x to W_n. If $(\alpha_n)_n$ is edge-escaping, then set $y = x$. If not, there is some edge $e \in E(x)$ contained in an infinite subsequence of α_n. We now move from x to the other endpoint of e and repeat. After at most r steps, we arrive at a point $y \in V \cap \bigcap_{n_i} \overline{W}_{n_i}$, for an infinite sequence n_i. Since $(W_n)_n$ is decreasing, it follows that $y \in V \cap \bigcap_n \overline{W}_n$ as required. $\diamondsuit$

We remark that an immediate consequence is that $(W_n)_n$ is escaping if and only if $\bigcap_n \overline{W}_n \subseteq \partial K$.

We next move on to consider a weaker property of "evasion".

If $W \subseteq V$ and $r \geq 0$ write $J_r(W)$ for the union of all arcs of length at most r with both endpoints in W. We write $EJ_r(W) \subseteq E(K)$ for the set of edges in $J_r(W)$.

Lemma 3.7 : *Let $(W_n)_n$ be a decreasing sequence of uniformly quasiconvex subsets of V. The following are equivalent:*
(1) $\bigcap_n EJ(W_n) = \emptyset$.
(2) For all $r \in \mathbf{N}$, $\bigcap_n EJ_r(W_n) = \emptyset$.
(3) $\left|V \cap \bigcap_n \overline{W}_n\right| \leq 1$.

Proof :
(1) $\Rightarrow$ (2) : Suppose that $e \in \bigcap_n EJ_r(W_n)$ for some $r \geq 0$. Let α_n be an arc of length at most r, containing e and with both endpoints in W_n. Let β_n be a geodesic with the same endpoints. If e lies in β_n for all n, then $e \in \bigcap_n EJ(W_n)$ and we are done. If not, then after passing to a subsequence, we can suppose that e does not lie in any β_n. Now there is a circuit containing e contained in the subgraph $\alpha_n \cup \beta_n$ of K. By fineness, we can suppose that this circuit is constant. Thus, there is some edge contained in each β_n, and hence in $\bigcap_n EJ(W_n)$.
(2) $\Rightarrow$ (1) : Suppose $e \in \bigcap_n EJ(W_n)$. Let α_n be a geodesic containing e with both endpoints in W_n. Consider the two components, β_n, γ_n, of $\alpha_n \setminus e$. By an argument similar

to that of (2) $\Rightarrow$ (3) in Lemma 3.2, we can replace β_n and γ_n by disjoint arcs, β'_n and γ'_n, of bounded length, connecting e to W_n. Thus $\beta'_n \cup e \cup \gamma'_n$ is an arc of bounded length with endpoints in W_n, so $e \in \bigcap_n EJ(W_n)$.

(2) $\Rightarrow$ (3) : Suppose that $x, y \in V \cap \bigcap_n \overline{W}_n$ are distinct. Let γ be any arc connecting x to y. By Lemma 3.2, there are edge-escaping sequences, $(\alpha_n)_n$ and $(\beta_n)_n$, of arcs of bounded length respectively connecting x to W_n and y to W_n. By fineness, $\alpha_n \cap \beta_n$ can be non-empty for only finitely many n. Moreover, $\alpha_n \cup \beta_n$ can contain an edge of γ for only finitely many n. For all other n, $\alpha_n \cup \gamma \cup \beta_n$ is an arc, giving the contradiction that any edge of γ lies in $\bigcap_n EJ_r(W_n)$ for some sufficiently large r.

(3) $\Rightarrow$ (2) : Let α_n be an arc containing e of length at most r with endpoints in W_n, and let β_n and γ_n be the components of $\alpha_n \setminus e$. Applying the argument of Lemma 3.6, after passing to a subsequence, we can find constant subpaths, β' and γ', of β_n and γ_n respectively, connecting e to points of $V \cap \bigcap_n \overline{W}_n$. It follows that these points must be distinct, showing that $\left|V \cap \bigcap_n \overline{W}_n\right| \geq 2$. ◇

(Note that in (2) it is enough to insist that $\bigcap_n EJ_r(W_n) = \emptyset$ for some r sufficiently large in relation to the constants of hyperbolicity and quasiconvexity. Similarly in (1), we can restrict to geodesics of uniformly bounded length.)

Definition : We say that a decreasing sequence, $(W_n)_n$, is *evading* if any, hence all, the conditions of Lemma 3.7 hold.

Lemma 3.8 : *If $x \in \bigcap_n W_n$, then $(W_n)_n$ is evading if and only if $(W_n \setminus \{x\})_n$ is escaping in $K \setminus \{x\}$.*

Proof : If $(\alpha_n)_n$ were a sequence of arcs of length at most r say, connecting e to W_n in $K \setminus \{x\}$, then $e \in \bigcap_n EJ_{r+1}(W_n)$, so $(W_n)_n$ cannot be evading. Conversely, if $(W_n)_n$ is not evading, let α_n be a sequence of arcs of bounded length containing e and with endpoints in W_n. At least one of the components of $\alpha_n \setminus e$ connects e to W_n in $K \setminus \{x\}$, showing that $W_n \setminus \{x\}$ is not escaping in $K \setminus \{x\}$. ◇

(In view of the fact that geodesics and arcs of bounded length are quasiconvex, Lemma 3.8 shows the property of being edge-escaping is equivalent to that of evading for a sequence of such arcs emanating from a given point.)

Proposition 3.9 : *Suppose that $(W_n)_n$ is an evading decreasing sequence of uniformly quasiconvex subsets of V. Then $\left|\bigcap_n \overline{W}_n\right| = 1$.*

Proof : By compactness of ΔK, $\bigcap_n \overline{W}_n \neq \emptyset$, so we must show that it contains at most one point.

Suppose, to the contrary, that $x, y \in \bigcap_n \overline{W}_n$ are distinct. Let $\alpha = [x, y]$. By Corollary 3.4, α lies a bounded distance, say r, for each W_n. By Lemma 3.7 part (3), x and y cannot both lie in V. Thus α is a ray or bi-infinite geodesic. We can thus find $z_0, z_1 \in \alpha$ with $d(z_0, z_1) \geq 2r + 1$. Now $d(x, W_n) \leq r$ for all n, so by Lemma 3.6, we can find

$w_0, w_1 \in V \cap \bigcap_n \overline{W}_n$ with $d(z_i, w_i) \leq r$. Thus $w_1 \neq w_2$, contradicting Lemma 3.7 part (3). $\diamondsuit$

In section 9, we will need the following observation.

Lemma 3.10 : *Suppose $W, W' \subseteq V(K)$ are quasiconvex. If there is an edge disjoint (or evading) sequence of arcs of bounded length connecting W to W' in K then $\overline{W} \cap \overline{W'} \neq \emptyset$.*

Proof : Let $(\alpha_n)_n$ be an edge-disjoint sequence of arcs of bounded length connecting $x_n \in W$ to $x'_n \in W'$.

If α_n is escaping, we can pass to a subsequence so that x_n converges in ΔK to some $x \in \partial K$. Thus, $x \in \overline{W} \cap \overline{W'}$.

Suppose then that α_n is not escaping. As in Lemma 3.6, after passing to a subsequence, we can find some $a \in V(K)$ and a sequence of edge-disjoint arcs, β_n, of bounded length with one endpoint at a and meeting α_n in precisely the other endpoint. Let γ_n and γ'_n be the subarcs of β_n connecting y_n respectively to x_n and x'_n. Now $(\beta_n \cup \gamma_n)_n$ and $(\beta_n \cup \gamma'_n)_n$ are edge-disjoint arcs of bounded length connecting a to W and W'. It follows that $a \in \overline{W} \cap \overline{W'}$. $\diamondsuit$

Corollary 3.11 : *Suppose $(W_n)_n$ and $(W'_n)_n$ are decreasing sequences of uniformly quasiconvex subsets of $V(K)$. Suppose that $(\alpha_n)_n$ is an edge-disjoint sequence of paths of bounded length connecting W_n and W'_n in K. Then $\bigcap_n \overline{W_n} \cap \bigcap_n \overline{W'_n} \neq \emptyset$.*

Proof : By Lemma 3.10, $\overline{W_n} \cap \overline{W'_n} \neq \emptyset$ for all n. But $\overline{W_n} \cap \overline{W'_n}$ is decreasing, so by compactness $\bigcap_n (\overline{W_n} \cap \overline{W'_n}) \neq \emptyset$. $\diamondsuit$

4. Stacks.

In this section, we introduce the notion of a "stack" of graphs. A more detailed discussion of stacks of metric spaces can be found in [Bow3]. Many of the underlying ideas can be found in Mitra's approach to the Cannon-Thurston map [Mit1,Mit2].

Let V be a set, and let $\mathcal{I} \subseteq \mathbf{Z}$ be a set of consecutive integers. Suppose that for each $i \in \mathcal{I}$, we have a graph K_i with vertex set V. Let $K = \bigcup_{i \in \mathcal{I}} K_i$ be the graph with vertex set V and edge set $\bigcup_i E(K_i)$. For notational convenience, we regard the subgraphs K_i as forming part of the structure of K, and refer to K, or to $(K_i)_{i \in \mathcal{I}}$, as a *stratified graph*. We write $d = d_K$ for the metric on K and $d_i = d_{K_i}$ for the metric on K_i.

We shall associate to a stratified graph, K, another graph, $Z = Z(K)$, as follows. We set $V(Z) = \mathcal{I} \times V$. For all $x \in V$ and $i \in \mathcal{I}$ we connect (i, x) to $(i+1, x)$ by a "vertical" edge (assuming $i + 1 \in \mathcal{I}$). If $i \in \mathcal{I}$ and $x, y \in V$ are adjacent in K_i, we connect (i, x) to (i, y) by a "horizontal" edge. We write $E^{\updownarrow}(Z)$ and $E^{\leftrightarrow}(Z)$ for the sets of vertical and horizontal edges respectively. Thus $E(Z) = E^{\updownarrow}(Z) \sqcup E^{\leftrightarrow}(Z)$.

Given $i \in \mathcal{I}$, write $V_i = \{i\} \times V \subseteq V(Z)$. Let Z_i be the full subgraph of Z with vertex set V_i. Thus $E(Z_i) \subseteq E^{\leftrightarrow}(Z)$. We refer to Z_i as the *sheet* of Z at *level* i. Note that Z_i is naturally isomorphic to K_i. We shall often identify Z_i with K_i.

Given $x \in V$, let $l(x)$ denote the full subgraph of Z with vertex set $\{x\} \times \mathcal{I}$. We refer to $l(x)$ as a *vertical line* of Z. It is isometric to a real interval, and geodesically embedded in Z.

There is a natural surjective projection, proj, from Z onto K. In particular, any path in α in Z projects to a path $\mathrm{proj}(\alpha)$ in K. Conversely, any path β in K lifts to a path $\mathrm{lift}(\beta)$ in Z with $\mathrm{proj}\,\mathrm{lift}(\beta) = \beta$: first lift each edge and then interpolate with vertical segments. In general this lift will not be unique (unless the $E(K_i)$ are disjoint). Note that the lift of an arc is an arc. Also, the projection of any geodesic in Z is an arc in K.

We shall be imposing various conditions on our stratified graph, K. In particular:

(S1) There is a constant $k \geq 0$ such that if $x, y \in V$ are adjacent in some K_i, then $d_{i+1}(x, y) \leq k$ and $d_{i-1}(x, y) \leq k$.

(S2) $Z(K)$ is hyperbolic.

(S3) K is fine.

(S4) There is a function, $F : \mathbf{N} \longrightarrow \mathbf{N}$ such that if $x, y \in V$, $r \in \mathbf{N}$ and $i, j \in \mathcal{I}$ with $d_i(x, y) \leq r$ and $d_j(x, y) \leq r$, then $|i - j| \leq F(r)$.

(S5) The sheets K_i are uniformly hyperbolic.

Remarks :

(1) Property (S1) is equivalent to asserting that for all i the identity on V is a uniform quasi-isometry from K_i to K_{i+1}. Thus the K_i are all quasi-isometric. We can also deduce that each Z_i is uniformly properly embedded in Z (Lemma 4.1).

(2) Given (S1) and (S5), property (S2) is equivalent to a certain "flaring" condition expressible in terms of the metrics d_i. This is the Bestvina-Feighn hyperbolicity criterion [BeF] — see condition (S2′) given in Section 6.

(3) Property (S3) implies that each of the K_i is fine. Indeed any union of K_i is fine. Moreover, $(S1)$ tells us that for all $i, j \in \mathcal{I}$, the inclusions of K_i and K_j into $K_i \cup K_j$ are both quasi-isometries. Thus, by (S5), $K_i \cup K_j$ is hyperbolic, and so Lemma 3.1 gives us homeomorphisms of ΔK_i and ΔK_j to $\Delta(K_i \cup K_j)$. In particular, the identity on V extends to a homeomorphism of ΔK_i to ΔK_j. We can thus define a compact hausdorff space, $\Delta^0 K$ with V embedded as a dense subset, to which all the ΔK_i are canonically homeomorphic.

(4) Given (S1), property (S4) tells us that if $x, y \in V$ are distinct, then the lines $l(x)$ and $l(y)$ cannot remain "close over a large distance" (Lemma 4.2). Thus, in view of hyperbolicity (S2), they can be thought of as "diverging uniformly". We also note that property (S4) is implied by the property (S4′) that appears in Section 7.

(5) Given (S1), (S2) and (S4), (S3) is equivalent to asserting that any finite union of K_i is fine. The idea is as follows. Suppose that γ is a circuit in K of length n. Lift γ to a circuit, α, in Z. Thus α has n horizontal edges and n vertical segments. By condering the nearest point retractions of α to vertical lines (see below) we find that the length of each

vertical segment, hence the total length of α, is bounded in terms of n. Projecting back to K, we see that γ lies in a subgraph of K of the form $\bigcup_{i \le k \le j} K_k$, where $j - i$ is bounded in terms of n. Since graphs of this type are assumed to be fine, we see that any edge of K lies in only finitely many circuits of a given length in K.

(6) In the special case of stacks of Farey graphs, discussed in Section 5, we shall be imposing an additional condition (S6).

Let us now assume that K satisfies properties (S1)–(S5).

Lemma 4.1 : *There is a function, $F_1 : \mathbf{N} \longrightarrow \mathbf{N}$ such that if $i \in \mathcal{I}$ and $x, y \in V(Z_i)$, then $d_i(x, y) \le F_1(d_Z(x, y))$.*

Proof : This is a simple exercise using (S1) (see [Bow3]). We could take, for example, $F_1(n) = k^n$. ♢

Lemma 4.2 : *There is a function, $F_2 : \mathbf{N} \longrightarrow \mathbf{N}$ such that if $x, y \in V$ are distinct, and $x_0, x_1 \in l(x)$, $y_0, y_1 \in l(y)$ and $r \in \mathbf{N}$, with $d_Z(x_0, y_0) \le r$ and $d_Z(x_1, y_1) \le r$, then $d_Z(x_0, x_1) \le F_2(r)$.*

Proof : Note that x_0, y_0 can be assumed to lie in the same sheet of Z, and similarly for x_1, y_1. Apply (S4) and Lemma 4.1. ♢

Given a subset, $Q \subseteq Z$, we can define $\pi_Q : Z \longrightarrow Q$, by choosing $\pi_Q(x) \in Q$ so that $d_Z(x, \pi_Q(x)) = d_Z(x, Q)$. This is well-defined up to a bounded distance in Z, and we refer to it as "the" *nearest point retraction*. (The choice involved is usually not significant.) An easy consequence of Lemma 4.2 is that if $x \ne y$, then $\pi_{l(x)}(l(y))$ has bounded diameter. This is the key observation in the following result (see [Bow2]).

Lemma 4.3 : *Suppose that α is a geodesic in K and that β is a lift of α to Z. Then β is uniformly quasigeodesic in Z.* ♢

Here we allow for there to be non-trivial vertical segments which project to the endpoints of α.

It follows easily that:

Lemma 4.4 : *K is uniformly hyperbolic.* ♢

The proof is essentially the same as that of Lemma 2.3. Indeed in applications, we will know directly by Lemma 2.3 that K is hyperbolic.

Lemma 4.5 : *Suppose $W \subseteq V$ is uniformly quasiconvex in K_i for all $i \in \mathcal{I}$. Then W is uniformly quasiconvex in K.*

Proof : Let $Q = \bigcup_{x \in W} l(x) = \text{proj}^{-1} W \subseteq Z$. If Q is quasiconvex in Z then the result follows; for if α is a geodesic segment in K connecting any two points of W, then $\text{lift}(\alpha)$ is a uniform quasigeodesic in Z connecting two points of Q. Since any quasigeodesic in a hyperbolic space remains a bounded distance from a geodesic, it follows that $\text{lift}(\alpha)$ remains a bounded distance from Q. Projecting back to K, the result follows.

To show that Q is indeed quasiconvex, we use an argument that appears in [Mit1,Mit2] and is reproduced in [Bow3]. Briefly the idea is as follows.

Given $i \in \mathcal{I}$, let π_i be the nearest point retraction of V onto W in the sheet Z_i. If $x, y \in V$ are adjacent in Z_i, then $d_Z(\pi_i(x), \pi_i(y)) \leq d_i(\pi_i(x), \pi_i(y))$ is bounded. Moreover, since the identity on V is a quasi-isometry from K_i to K_{i+1}, it follows that if $x \in Z_i$ and $y \in Z_{i+1}$ are connected by a vertical edge, then $d_Z(\pi_i(x), \pi_{i+1}(y))$ is bounded.

We now assemble the retractions π_i into a single map $\pi : V(Z) \longrightarrow Q$. The above observations show that π can increase distances in Z by at most a linearly bounded amount. From this it follows that Q is quasiconvex, as required. ◊

(We should remark that the reason given in [Mit2] for the last assertion, namely the existence of a linearly bounded retraction implying quasiconvexity, appears to be incomplete. However, this is easily rectified, see for example [Bow3].)

In fact, the argument we have given shows more:

Lemma 4.6 : *Suppose that $W \subseteq V$ is uniformly quasiconvex in K_i for each i. Suppose that α is a geodesic in K connecting two points of W, and that $e \in E(K_i)$ is an edge of α. Then $d_i(e, W)$ is uniformly bounded.*

Proof : Let f be the edge corresponding to e in $\text{lift}(\alpha) \subseteq Z$. Since Q is quasiconvex and $\text{lift}(\alpha)$ is quasigeodesic, there is some $x \in W$ with $d_Z(e, l(x))$ bounded. It follows that $d_Z(e, y)$ is bounded, where y is the vertex of $l(x)$ at level i. Thus, by Lemma 3.1, $d_i(e, W) \leq d_i(e, y)$ is bounded as claimed. ◊

Lemma 4.7 : *Suppose that $(W_n)_n$ is a decreasing sequence of subsets of V. Suppose that for all $i \in \mathcal{I}$ and $n \in \mathbf{N}$, W_n is uniformly quasiconvex in K_i. Suppose that for some (hence all) $i \in \mathcal{I}$, $(W_n)_n$ is escaping in K_i. Then $(W_n)_n$ is evading in K.*

Proof : If not, then there is some edge $e \in E(K)$, and a sequence of geodesics, α_n, in K containing e and with both endpoints in W_n. Suppose $e \in E(K_i)$. By Lemma 4.6, $d_i(e, W_n)$ is bounded, contradicting the assumption that $(W_n)_n$ is escaping in K_i. ◊

5. Farey graphs.

In this section, we consider a stack, Z, of Farey graphs. We show (Proposition 5.5) that under hypotheses (S1)–(S5) together with hypothesis (S6) below, we can define a Cannon-Thurston map from $S^1 \cong \Delta^0 K$ to ΔK.

We begin by describing the Farey graph. It is easily verified that up to isomorphism there is a unique simply connected 2-dimensional simplicial complex, Ω, with the property that each 1-simplex is incident on exactly two 2-simplices and each 0-simplex is incident on infinitely many 1-simplices.

Definition : We refer to the complex Ω described above as the *Farey complex* and to its 1-skeleton as the *Farey graph*.

One can easily verify that the Farey graph A is fine and hyperbolic. (Indeed A is quasi-isometric to an infinite valence tree, though not in any natural way.) Moreover, ΔA is homeomorphic to a circle. The link of each vertex of Ω is the real line. Every edge of A separates Ω. Also, if $x \in V(A)$, then the metric $\angle_x$ on $E(x)$ is isometric to the standard metric on the integers, $\mathbf{Z}$. Note that Ω is determined by the combinatorics of A, so we may write $\Omega = \Omega(A)$.

Combinatorially, the Farey complex arises naturally as a regular tessellation of the hyperbolic plane by ideal triangles, where we have included all the rational ideal points as vertices. (We can assume that the triangulation is invariant under the action of $PSL(2, \mathbf{Z})$.) We can thus think of the Farey graph as dual to the Apollonian packing of $\mathbf{H}^2$ by horodiscs. A simple consequence of Lemmas 2.1 and 2.3 is that ΔA is homeomorphic to the circle $S^1 \cong \partial \mathbf{H}^2$. (This can also be easily verified directly.)

In the discussion that follows, we shall put a complete metric on Ω by giving each 2-simplex the structure of a euclidean equilateral triangle of unit side-length. Thus, $\Omega \setminus V(\Omega)$ is euclidean.

Suppose $\Phi \subseteq \Omega$ is a subcomplex such that $\Phi \setminus V$ is connected. Then, Φ is convex. (It is enough to note that its boundary in Ω is locally convex.) In particular, suppose $I \subseteq S^1 \cong \Delta A$ is an interval. Let $\Phi(I)$ denote the union of all 2-simplices of Ω with at least one vertex in I. It is easily verified that $\Phi(I) \setminus V$ is connected. Thus $\Phi(I)$ is convex. We conclude:

Lemma 5.1 : *Suppose that $I \subseteq S^1 \equiv \Delta A$ is an interval. Then $V \cap I$ is uniformly quasiconvex in A.*

Proof : The complex $\Phi(I)$ is convex and lies in a 1-neighbourhood of $V \cap I$. Thus $V \cap I$ is quasiconvex in Ω and hence in A. ◊

We also note that if $(I_n)_n$ is a decreasing sequence of intervals with $\bigcap_n I_n = \{x\}$, then $(V \cap I_n)_n$ is escaping if $x \in \partial A$, and evading if $x \in V$.

Now consider a stratified graph, $K = \bigcup_{i \in \mathcal{I}} K_i$ where each K_i is a Farey graph. Let $Z = Z(K)$ be the associated stack. We suppose hypotheses (S1)–(S4) hold ((S5) being automatic). Thus, $\Delta^0 K \cong \Delta K_i$ (as defined by Remark (3)) is homeomorphic to a circle with V as a dense subset.

Putting together Lemma 5.1 and Lemma 4.5, we immediately deduce:

Lemma 5.2 : *If $I \subseteq \Delta^0 K$ is an interval, then $V \cap I$ is uniformly quasiconvex in K.* $\diamondsuit$

To progress further, we shall impose one further condition on our stack, namely:

(S6) There is a constant $\theta \in \mathbf{N}$ such that if $i, j \in \mathcal{I}$, and $x, y, z \in \mathcal{I}$ with y, z both adjacent to x in K_i and with $d_{K_i \setminus \{x\}}(y, z) \geq \theta$, then there is some $w \in V$ adjacent to x in K_j such that $\{x, w\}$ separates $\{y, z\}$ in $\Delta^0 K$.

Another way of thinking of this is as follows. Let $E_i(x) = E(K_i) \cap E(x)$ where $E(x)$ is the set of edges incident on x in K. If $\iota(e)$ denotes the other endpoint of $e \in E(x)$, then since $V \subseteq \Delta^0 K$, the map ι gives us an embedding of $E(x)$ in $\Delta^0 K \setminus \{x\}$, which we know to be homeomorphic to $\mathbf{R}$. Moreover, the angular metric, $\angle_x$, on each $E_i(x)$ is isometric to $\mathbf{Z}$. Condition (S6) tells us that we can, in fact, identify $\Delta^0 K \setminus \{x\}$ with $\mathbf{R}$ in such a way that each inclusion $\iota | E_i(x)$ is a uniform quasi-isometry from $\mathbf{Z}$ to $\mathbf{R}$ in their standard metrics.

Suppose then that K satisfies (S1)–(S6).

Lemma 5.3 : *Suppose $x \in V$ and that $(I_n)_n$ is a decreasing sequence of intervals in $\Delta^0 K$ with $\bigcap_n I_n = \{x\}$. Then $(V \cap I_n)_n$ is evading in K.*

Proof : By Lemma 3.8, this is the same as asserting that $((V \cap I_n)_n \setminus \{x\})_n$ is escaping in $K \setminus \{x\}$.

Fix any $i \in \mathcal{I}$. Let $\iota : E_i(x) \equiv \mathbf{Z} \longrightarrow \Delta^0 K$ be the embedding referred to above, and write $y_m = \iota(m)$ (so that $\{x, y_m\}$ separates $\{y_{m-1}, y_{m+1}\}$ for all m.) Now given any $p \in \mathbf{N}$, we can find $n \in \mathbf{N}$ such that $\{y_{-\theta p}, y_{\theta p}\}$ separates I_n from y_0 in $\Delta^0 K$.

Suppose $y_0 = z_0, z_1, \ldots, z_m$ is a path in K connecting y_0 to $V \cap I_n$. Let J_k denote the open interval between z_{k-1} and z_k in $\Delta^0 K \setminus \{x\}$. Now z_{k-1} and z_k are adjacent in some K_j for $j \in \mathcal{I}$. Thus $J_k \cap E_j(x) = \emptyset$, and so by property (S6), J_k contains at most $\theta - 1$ points of $E_i(x)$. Now the intervals $(J_k)_{k=1}^m$ together get us from y_0 to $V \cap I_n$, and therefore must eventually cross $\{y_{-\theta p}, y_{\theta p}\}$. Thus $\theta p \leq \theta m$ and so $p \leq m$. In other words, we have shown that $d_{K \setminus \{x\}}(y_0, V \cap I_n \setminus \{x\}) \geq p$. Letting $p \to \infty$, the result follows. $\diamondsuit$

Lemma 5.4 : *Suppose that $(I_n)_n$ is a decreasing sequence of intervals in $\Delta^0 K$ with $|\bigcap_n I_n| = 1$. Then $(V \cap I_n)_n$ is evading in K.*

Proof : Let $\bigcap_n I_n = \{x\}$, where $x \in \Delta^0 K$.

If $x \in V$, then the result follows by Lemma 5.3.

If $x \notin V$, then $(V \cap I_n)_n$ is escaping in each K_i. In this case the result follows from Lemma 4.7. $\diamondsuit$

We can now define our Cannon-Thurston map for stacks.

Proposition 5.5 : *Suppose that $K = \bigcup_i K_i$ is a stratified graph with each K_i isomorphic to the Farey graph. Suppose that K satisfies (S1)–(S6). Then there is a continuous map from $\Delta^0 K$ to ΔK which is the identity on V.*

Proof : The set, $\mathcal{O}$, of open intervals in $\Delta^0 K$ is a base for the topology. The result therefore follows by putting together Lemma 5.1, Lemma 5.4, Proposition 3.9 and Lemma 1.1. $\diamondsuit$

We finish this section by discussing a way in which stacks of Farey graphs arise and showing that property (S6) is automatic in this case. First, we note:

Lemma 5.6 : *Let $K = \bigcup_i K_i$ be a stratified graph of Farey graphs satisfying (S1)–(S4). Suppose that a group, Γ, acts on K preserving each subgraph K_i. Suppose that there is some $h \in \mathbf{N}$ such that for all i, $E(K_i)$ is the union of at most h Γ-orbits. Then K satisfies (S6).*

Proof : Let $x \in V$. The Γ-stabiliser, $\Gamma(x)$, of x acts by homeomorphism on $\Delta^0 K \setminus \{x\} \cong \mathbf{R}$, preserving each discrete subset $\iota(E_i(x))$. Since $E_i(x)/\Gamma(x)$ is finite, it follows that $(\Delta^0 K \setminus \{x\})/\Gamma(x)$ is compact. Let $\Gamma_0(x)$ be the end-preserving subgroup of $\Gamma(x)$. Thus $\Gamma_0(x)$ is infinite cyclic generated by some element g. Now for any $y \in \Delta^0 K \setminus \{x\}$, the closed interval between y and gy is a fundamental domain for the action of $\Gamma_0(x)$. This contains at most $4h+1$ elements of $E_j(x)$ for any $j \in \mathcal{I}$. The result follows easily. $\diamondsuit$

Our stack of Farey graphs will arise from the following construction. Let Θ be a closed surface, and $P \subseteq \Theta$ a non-empty finite subset. For example, Θ might be obtained by collapsing to a point each boundary component of the surface Σ described in the Introduction. In this case we have a natural bijection between P and the indexing set $\mathcal{P}$.

Suppose we put a path-metric on Θ (in practice, this will be a singular, non-positively curved riemannian metric). Let X be the completion of the universal cover, $\widetilde{\Theta \setminus P}$, of $\Theta \setminus P$. Thus, $\Gamma = \pi_1(\Theta \setminus P) \equiv \pi_1(\Sigma)$ acts on X with quotient Θ. We may identify the preimage of P with the set of parabolic points, Π, described in the Introduction. Thus $\Pi \subseteq X$ is precisely the set of points at which X is not locally compact.

Now consider any triangulation of Θ with vertex set P. (Such a triangulation may be singular, in the sense that a simplex may have identifications around its boundary.) We give each simplex the structure of a euclidean equilateral triangle with unit side-length. With this metric, we see that the space X obtained above is a Farey complex with vertex set Π. Any other space space obtained in similar fashion starting with a singular riemannian metric on Θ will be equivariantly bilipschitz equivalent to this model space.

Let $\mathcal{M}$ be the mapping class group of (Θ, P), i.e. the group of homotopy classes of self-homeomorphisms of Θ relative to P, which are the identity on P. If $\psi \in \mathcal{M}$, then by taking the image of our triangulation under (a representative of) ψ, and lifting to X, we get another Farey graph with vertex set Π. We denote this by $K(\psi)$.

If $\mathcal{I} \subseteq \mathbf{Z}$ is a set of consecutive integers, and $(\psi_i)_i$ is a set of elements of $\mathcal{M}$ indexed by $\mathcal{I}$, then we get a stratified graph, namely $K = \bigcup_i K_i$, where $K_i = K(\psi_i)$. Now Γ acts on K satisfying the hypotheses of Lemma 4.6. Thus, if K satisfies (S1)–(S4), then it also satisfies (S6).

6. Punctured-surface bundles.

In this section we give a proof of Theorem 0.1 in the special case of manifolds fibring over the circle. Our argument uses the Bestvina-Feighn flaring condition, as well as the existence of Teichmüller differentials for pseudo-anosov mapping classes. In Sections 7 and 8, we will give a different proof in the general case that does not make use of these particular results.

Let Σ be a compact surface with boundary components $(C^m)_{m\in\mathcal{P}}$, and let Θ be the surface obtained by collapsing each boundary component to a point as in Section 5. Any mapping class, ψ, in $\mathcal{M}$ also determines a mapping class of Σ. We can thus form the mapping torus $M(\psi)$ with toroidal boundary components in bijective correspondence with $\mathcal{P}$ or with P. Let $\operatorname{int} M(\psi)$ be its interior. We recall the result of Thurston [T2], see also [O]:

Theorem 6.1 : $\operatorname{int} M(\psi)$ *admits a complete hyperbolic structure if and only if* ψ *is pseudo-anosov.* $\diamondsuit$

Such a manifold has finite volume, and the structure is unique by Mostow rigidity.

We remark that, as a consequence, the fundamental group, $G = \pi_1(M(\psi))$, is hyperbolic relative to its peripheral subgroups, in the sense of Gromov [Gr1]. In fact, we shall give another proof of this much weaker assertion below.

In the case of a compact manifold, $M(\psi)$, the hyperbolicity of $\pi_1(M(\psi))$ is shown directly in [BeF], and the existence of the Cannon-Thurston map was shown in the original paper [CannT].

In summary, we shall prove:

Proposition 6.2 : *Let* Σ *be a compact surface with non-empty boundary, and let* ψ *be a pseudo-anosov mapping class. Let* $M(\psi)$ *be the mapping torus thus defined. Then* $G = \pi_1(M(\psi))$ *is hyperbolic relative to the peripheral subgroups. Moreover, there is a* Γ*-equivariant map from* S^1 *to* ∂G*, where* $\Gamma = \pi_1(\Sigma) \triangleleft G$*.*

Here, S^1 is the circle at infinity associated to Σ (by choosing any finite-area hyperbolic structure on our surface), and ∂G is the boundary of G as a relatively hyperbolic group as defined in [Bow2]. In fact, given Theorem 6.1, we do not need to worry about the latter point, since we shall see directly that the target space can be G-equivariantly indentified with $\partial \mathbf{H}^3$.

To prove Proposition 6.2, we construct the stack of Farey graphs, $Z = Z(K)$ where $K = \bigcup_{i\in\mathbf{Z}} K_i$ and K_i is defined as $K(\psi^i)$ as at the end of Section 5. Here ψ^i is the ith power of ψ.

Note that G acts on Z with finite quotient, and that $\Gamma \triangleleft G$ preserves each sheet. Now, properties (S1) and (S5) are automatic, and (S6) will be taken care of by Lemma 6.5. We thus need to verify (S2), (S3) and (S4). The key to these is (S2) — the hyperbolicity of Z.

Here is the equivalent "flaring condition" (stated for any indexing set, $\mathcal{I}$, of consecutive integers):

(S2′) There exist $k \in \mathbf{N}$, $0 < \lambda < \frac{1}{2}$ and $c > 0$ such that if $i-k, i+k \in \mathcal{I}$, then for all $x, y \in V$ we have

$$d_i(x,y) \leq \lambda(d_{i-k}(x,y) + d_{i+k}(x,y)) + c.$$

Indeed, we can make λ arbitrarily small at the cost of increasing k and c. Note that (S2′) makes sense for an arbitrary stratified graph. It follows from [BeF] (as made explicit in [Bow3]) that:

Theorem 6.3 : *If $K = \bigcup_i K_i$ is a stratified graph satisfying (S1) and (S5), then (S2) is equivalent to (S2′).* ◊

In order to verify (S2′) in our situation, we make use of Teichmüller differentials, interpreted geometrically as singular euclidean stretch maps. Indeed, this was the idea behind the original proof in the compact case [CannT].

To this end, we return to our closed surface, Θ, with $P \subseteq \Theta$ finite. By a *singular euclidean* structure on Θ, we mean a metric that is euclidean away from a finite set $Q \supseteq P$. At each point of Q, we have a cone singularity. Such a metric is *good* if the cone angle at each point of Q is an integer multiple of π, and in addition, is at least 3π at every point of $Q \setminus P$.

Given a good metric, we can construct X as the completion of $\widetilde{\Theta \setminus P}$ as in Section 5. This is a CAT(0) space. We denote the preimage of Q in X by $\tilde{Q}$. The preimage of P can be identified with $\Pi \subseteq \tilde{Q}$. We note that a path α in X is geodesic if and only if each component of $\alpha \setminus \tilde{Q}$ is a euclidean geodesic segment and each of the exterior angles at each point of $\tilde{Q}$ is at least π. (There is only one exterior angle at a point of Π.)

We can define a *stretch map* on Θ as follows. At each point of $\Theta \setminus Q$, we choose local euclidean coordinates (ζ, ξ) such that all transition functions are translations, possibly composed with a rotation through π. (This is possible since all the cone angles are multiples of π.) The euclidean metric is given infinitesimally by $ds^2 = d\zeta^2 + d\xi^2$. We also have a $\sqrt{2}$-bilipschitz equivalent L^1-metric given by $ds = d\zeta + d\xi$. Given $t \in \mathbf{R}$, we define the "stretched metric" by $ds^2 = e^{2t}d\zeta^2 + e^{-2t}d\xi^2$. This is also a good singular euclidean metric on X, and we denote the induced path metric by ρ_t. From the earlier description, we see that the property of being a geodesic is invariant under stretching.

In the metric spaces (X, ρ_t), the equivalent of property (S2′) is easily seen to be satisfied. (Note that in the bilipschitz equivalent L^1-metric, the distance at time t between any two given points has the form $[t \mapsto Ae^{\mu t} + Be^{-\mu t}]$ for non-negative constants A and B.) Indeed, we can make λ as small as we want. Moreover, if we assume a positive lower bound on $\rho_t(x,y)$ for $t \in \mathbf{R}$, then it follows that (the equivalent of) property (S4) holds, namely that $|t-u|$ is bounded above as a function of $\max\{\rho_t(x,y), \rho_u(x,y)\}$.

We now return to the set-up of Proposition 6.2. Let $\psi \in \mathcal{M}$ be pseudo-anosov. The existence of a Teichmüller differential (see for example [O]) tells us that there is a good singular euclidean metric, ρ_0, on Θ, a constant $\mu > 0$, and a representative, $\hat{\psi}$ of ψ which is an isometry from (Θ, ρ_0) to (Θ, ρ_μ). Since the local coordinate system is determined (from the eigenvalues) by the stretch map, it follows that all iterates of $\hat{\psi}$ are stretch maps, and that $\hat{\psi}^i$ is an isometry from (Θ, ρ_0) to $(\Theta, \rho_{\mu i})$ for all $i \in \mathbf{Z}$. The same applies on lifting $\hat{\psi}$ to a Γ-equivariant map $\tilde{\psi}$ from X to itself.

Now there are constants $c_0, c_1 > 0$ such that if $p, q \in \Pi$ are adjacent in K_0 then $c_0 \leq \rho_0(p,q) \leq c_1$. Pushing forward under $\tilde{\psi}^i$, we see that $c_0 \leq \rho_i(p,q) \leq c_1$. Thus the identity on Π is uniformly bilipschitz with respect to the metrics d_i and $\rho_{\mu i}$.

Since there is a lower bound on $\rho_t(p,q)$ for $p,q \in \Pi$ and $t \in \mathbf{R}$ (given the cocompactness of the action of G), we see that properties (S2′) and (S4) hold for the stratified graph $K = \bigcup_i K_i$.

Note that Γ acts on any finite union of K_i with finite quotient. From this, it follows easily that any such finite union is fine (see [Bow2]). Thus, by Remark (5) of Section 5, we see that K is fine. Thus property (S3) holds. (In fact, given Theorem 6.1, the fineness of K follows from Lemma 2.2, as we discuss below.)

In summary, we know that K satisfies all the properties (S1)–(S6). In particular, it is fine and hyperbolic. Now G acts on K with finite quotient. It follows (by Definition 2 of [Bow2]) that G is hyperbolic relative to the set of vertex stabilisers. By the construction, the vertex stabilisers are precisely the peripheral subgroups of $G = \pi_1(M)$. Moreover, we can identify ΔK with the boundary, ∂G, of G as a relatively hyperbolic group. Similarly we can identify $\Delta^0 K$ with the boundary, $\partial\Gamma$, of Γ as a relatively hyperbolic group. Thus Proposition 5.5 gives us a continuous equivariant map of $\partial\Gamma \equiv \Delta K$ to ∂G. This proves Proposition 6.2.

If we admit Theorem 6.1, then we don't need to know anything about relatively hyperbolic groups. We have an action of G on $\mathbf{H}^3$, and the Margulis Lemma gives us a G-invariant set of disjoint horoballs $(B(p))_{p\in\Pi}$, where we identify Π with the set of parabolic points. Since G acts on K with finite quotient, there is an upper bound, say t_0 on $d(B(p), B(q))$ for $p, q \in \Pi$ adjacent in K. Moreover, $(B(p))_{p\in\Pi}$ is r-quasidense for some $r \geq 0$. Let $t = \max\{t_0, 2r+1\}$. By Lemmas 2.2 and 2.3, the t-nerve $L(t)$ is fine and hyperbolic, and so $\Delta L(t)$ is equivariantly homeomorphic to $\partial\mathbf{H}^3$. Moreover, since G acts cofinitely on $L(t)$, the inclusion of K in $L(t)$ is a quasi-isometry. Thus, by Lemma 2.1, we can also identify ΔK with $\partial\mathbf{H}^3$. Again the existence of the Cannon-Thurston map follows from Proposition 5.5.

7. Horoball systems.

As a step towards proving Theorem 0.1 in the general case, we show that a stratified graph associated to a horoball system will satisfy conditions (S2), (S3) and (S4) of Section 4. For applications, we will need a statement intrinsic to the complement of the set of horoballs (Proposition 7.12), though for most of this section we will be working with the hyperbolic metric, ρ, on $\mathbf{H}^\nu$. Given $x, y \in \mathbf{H}^\nu \cup \partial\mathbf{H}^\nu$, we write $[x,y]$ for the geodesic from x to y.

Recall the set-up of Section 2. Thus, we have $\Pi \subseteq \partial\mathbf{H}^\nu$, and we assume that $(B(p))_{p\in\Pi}$ is r_0-quasidence (in hull($\overline{\Pi}$)) for some $r_0 \geq 0$. Let $L(\infty)$ be the complete graph on Π, and let $L(t) \subseteq L(\infty)$ denote the t-nerve. Given $e \in L(\infty)$ with endpoints $p, q \in \Pi$, we write $\mu(e) = \mu(p,q)$ for the midpoint of the shortest path in $\mathbf{H}^\nu$ from $B(p)$ to $B(q)$. Thus

$\mu(e) \in [p, q]$.

Suppose now that $\mathcal{I} \subseteq \mathbf{Z}$ is a set of consecutive integers. Suppose that for each $i \in \mathbf{Z}$, we have a subgraph $L_i \subseteq L(\infty)$. Write $L = \bigcup_i L_i$. We shall suppose:

(P1) For some $t_0 \geq 2r_0 + 1$, we have $L(2r_0 + 1) \subseteq L \subseteq L(t_0)$.

(P2) There is an increasing function $\Psi : [0, \infty) \longrightarrow [0, \infty)$ such that if $e \in E(L_i)$ and $f \in E(L_j)$, then $|i - j| \leq \Psi(\rho(\mu(e), \mu(f)))$.

Note that $L = \bigcup_i L_i$ is a stratified graph on the vertex set Π. We may thus define the stack $Z = Z(L)$ as in Section 4.

We shall show:

Proposition 7.1 : *If $(B(p))_{p \in \Pi}$ is r_0-quasidense and $L = \bigcup_i L_i$ satisfies (P1) and (P2), then $Z(L)$ is hyperbolic.*

The constant of hyperbolicity depends only on r_0, t_0 and Ψ. We shall go on to show that L satisfies also hypotheses (S3) and (S4).

We begin with some preliminary observations.

Lemma 7.2 : *There is an increasing function g_0 such that if $p, q \in \Pi$ are distinct and β is any path from $\mathbf{H}^\nu$ from $B(p)$ to $B(q)$, then β lies in a $g_0(\mathrm{length}(\beta))$-neighbourhood of $\mu(p, q)$.*

Proof : Set $g_0(x) = \sinh(x)$. $\diamondsuit$

We shall say that a point x is a distance r *inside* a horoball B if $x \in B$ and $\rho(x, \partial B) = r$. A path *enters* B a distance at least r if there is a point on the path a distance (at least) r inside B. We note:

Lemma 7.3 : *There is a constant, t_1, such that the following holds. Suppose $p \in \Pi$ and that $e, f \in L(t_0)$ are incident on p. Let $\alpha = [\mu(e), \mu(f)]$. We have:*

(1) $\alpha \setminus B(p)$ lies in a t_1-neighbourhood of $\{\mu(e), \mu(f)\}$.

(2) If $x \in \alpha$ lies a distance r inside $B(p)$, then $\rho(x, \{B(p), B(q)\}) \leq r + t_1$.

(3) If α enters $B(p)$ a distance at most r then $\mathrm{length}(\alpha) \leq 2r + t_1$.

(4) If $q \in \Pi \setminus \{p\}$, then α enters at most t_1 into $B(q)$. $\diamondsuit$

Lemma 7.4 : *Supppose that $p_0, p_1 \in \Pi$ are distinct. Suppose that $e_0, f_0 \in E(L(t_0))$ are incident on p_0 and $e_1, f_1 \in E(L(t_0))$ are incident on p_1. Suppose $x_0 \in [\mu(e_0), \mu(f_0)]$ and $x_1 \in [\mu(e_1), \mu(f_1)]$. Then for $i = 0, 1$, we have $\rho(x_i, \{\mu(e_i), \mu(f_i)\}) \leq \rho(x_0, x_1) + 2t_1$.*

Proof : By 7.3(4), $[\mu(e_1), \mu(f_1)]$ can enter at most t_1 into $B(p_0)$. Thus x_0 lies at most $\rho(x_0, x_1) + t_1$ inside $B(p_0)$. By 7.3(2), $\rho(x_0, \{\mu(e_0), \mu(f_0)\}) \leq \rho(x_0, x_1) + 2t_1$. Similarly for x_1. $\diamondsuit$

Suppose π is a path in L, and let $e_1, e_2, \ldots, e_n \in E(L)$ be the sequence of edges of π. We shall write $\sigma(\pi)$ for the broken geodesic path

$$\sigma(\pi) = [\mu(e_1), \mu(e_2)] \cup [\mu(e_2), \mu(e_3)] \cup \cdots \cup [\mu(e_{n-1}), \mu(e_n)].$$

Thus each segment $[\mu(e_{i-1}, \mu(e_i)]$ of $\sigma(\pi)$ is associated to an interior vertex p_i of the path π, and lies mostly (all but a bounded length) in the horoball $B(p_i)$. Note that if $q \in \Pi \setminus \pi$, then $\sigma(\pi)$ enters $B(q)$ at most a distance t_1. Given a path π in Z, we shall similarly write $\sigma(\pi) = \sigma(\text{proj}(\pi))$, where $\text{proj} : Z \longrightarrow L$ is the natural projection map.

The strategy for proving Proposition 7.1 will be as follows. Given any $\epsilon > 0$, we define a path metric, d_Z^ϵ on Z by assigning to each horizontal edge unit length, and to each vertical edge a length ϵ. Thus the standard combinatorial metric on Z is $d_Z = d_Z^1$. We write Z^ϵ for (Z, d_Z^ϵ).

Now Z^ϵ is quasi-isometric to Z, so it is sufficient to show that Z^ϵ is hyperbolic for some (uniform) $\epsilon > 0$. (We remark that as $\epsilon \rightarrow 0$, Z^ϵ collapses to L, which we know to be hyperbolic by Lemma 2.3.) We shall in turn prove this by finding some $\epsilon > 0$, so that if π is geodesic in Z^ϵ, then $\sigma(\pi)$ is uniformly quasigeodesic in $\mathbf{H}^\nu$.

Before embarking on this, we shall need to recall some general facts about quasi-geodesics. The following discussion applies to any Gromov hyperbolic space, though we shall refer only to $\mathbf{H}^\nu$.

Given a path α in $\mathbf{H}^\nu$, write $\text{straight}(\alpha)$ for the geodesic segment with the same endpoints. If Q is a linear function, we say that α is a *Q-quasigeodesic* if every subpath, β, of α satisfies $\text{length}(\beta) \leq Q(\text{length}(\text{straight}(\beta)))$. If $r \geq 0$, we say that α is *r-locally* Q-quasigeodesic if this relation holds for all subpaths β of length at most r.

If $\alpha = \beta_1 \cup \cdots \cup \beta_n$ is a broken geodesic with geodesic segments $\beta_1, \ldots, \beta_n$, then we write $\text{back}(\alpha) = \max\{\text{diam}(N_1(\beta_i) \cap N_1(\beta_j)) \mid 1 \leq i < j \leq n\}$, where N_1 denotes the 1-neighbourhood. We refer to $\text{back}(\alpha)$ as the *backtracking* of α. This is a convenient, if somewhat artificial, way of expressing the maximal distance over which distinct segments of α can remain "close".

We shall write $\text{hd}(A, B)$ for the Hausdorff distance between two subsets $A, B \subseteq \mathbf{H}^\nu$.

We note the following facts:

Proposition 7.5 :

(1) Given a linear function Q, there is some $r \geq 0$ such that if α is a Q-quasigeodesic and $\beta = \text{straight}(\alpha)$, then $\text{hd}(\alpha, \beta) \leq r$.

(2) Given a linear function Q, there is another linear function Q' and some $r \geq 0$ such that any r-locally Q-quasigeodesic path is (globally) Q'-quasigeodesic.

(3) Given any increasing subexponential function S, there is a linear function Q such that if α is any path with the property that any subpath β satisfies $\text{length}(\beta) \leq S(\text{length}(\text{straight}(\beta)))$, then α is Q-quasigeodesic.

(4) Given any $r \geq 0$, there is a subexponential function S such that if α is a broken geodesic with at most n segments and $\mathrm{back}(\alpha) \leq r$, *then* $\mathrm{length}(\alpha) \leq \mathrm{length}(\mathrm{straight}(\alpha)) + S(n)$.
$\diamondsuit$

Remarks : Properties (1) and (2) are well known (see for example, [GhH]). Property (3) follows directly from the standard arguments (e.g. that in [GhH]) for proving (1) (see for example [Bow2] for an explicit account). Properly (4) can be proven easily so as to give S quadratic. In fact, one can take S to be linear, though this appears to be less obvious (see [Bow1]). The last observation would allow us to dispense with part (3) in applications, though overall an argument via (3) using S quadratic would seem to be simpler.

We now continue with the discussion of our horoball system $(B(p))_{p\in\Pi}$.

Lemma 7.6 : *There is a constant, t_2, such that if π is an arc in L, then* $\mathrm{back}(\sigma(\pi)) \leq t_2$.

Proof : This is an easy consequence of the fact that each segment of $\sigma(\pi)$ lies mostly inside the associated horoball (7.3(1)). Since we are assuming π to be an arc, distinct arcs correspond to distinct horoballs. $\diamondsuit$

Lemma 7.7 : *There is a linear function T_0 and an increasing function F_0 such that if $r \geq 0$ and e, f are distinct horizontal edges of Z with $\rho(\mu(e), \mu(f)) \leq r$, then $d_Z^{\epsilon}(e, f) \leq T_0(r) + \epsilon F_0(r)$.*

Proof : Let $p, q \in \Pi$ be either of the endpoints of e, f respectively. There is a path γ in $\mathrm{hull}(\overline{\Pi})$ connecting $B(p)$ to $B(q)$ of length at most $r + t_0$ and passing though the points $\mu(e)$ and $\mu(f)$ (namely the segment $[\mu(e), \mu(f)]$ together with two geodesic segments, each of length at most $t_0/2$, connecting respectively $\mu(e)$ to $B(p)$ and $\mu(f)$ to $B(q)$). Choose points $x_0, x_1, \ldots, x_n$ on γ with $x_0 \in B(p)$, $x_n \in B(q)$, $n \leq \mathrm{length}(\gamma) + 1 \leq r + t_0 + 1$ and $\rho(x_m, x_{m+1}) \leq 1$ for all m. For each m, choose $p_m \in \Pi$ with $\rho(x_m, B(p_m)) \leq r_0$. We can take $p_0 = p$ and $p_n = q$.

Now if $p_m \neq p_{m+1}$, then since $\rho(B(p_m), B(p_{m+1})) \leq 2r_0 + 1$ and $L(2r_0 + 1) \subseteq L$ (hypothesis (P1)) we see that p_m and p_{m+1} are adjacent in L, and hence connected by a horizontal edge, e_m, in Z.

We now construct a path β by taking as horizontal edges e and f together with the set of edges of the form e_m arising above whenever $p_m \neq p_{m+1}$, and by interpolating by vertical segments (lying in $l(p_m)$). Clearly the horizontal length of β is at most $n + 2$. To bound the vertical length, note that if $p', q' \in \Pi$ correspond to the endpoints of a horizontal edge a of β, then there is a path of length at most t_0 from $B(p)$ to $B(q)$ in $\mathbf{H}^{\nu}$ meeting the path γ. Thus, by Lemma 7.2, we have $\rho(\mu(a), \gamma) \leq g_0(t_0)$. If b is another horizontal edge of β, then we also have $\rho(\mu(b), \gamma) \leq g_0(t_0)$ and so $\rho(\mu(a), \mu(b)) \leq \mathrm{length}(\gamma) + 2g_0(t_0) \leq r + t_0 + 2g_0(t_0)$. If a, b lie in the sheets Z_i, Z_j respectively, then by hypothesis (P2), we have $|i - j| \leq \Psi(r + t_0 + 2g_0(t_0))$. In particular, each vertical segment comprising β has at most $\Psi(r + t_0 + 2g_0(t_0))$ edges. Since there are at most $n + 1$ such segments, β has at most $(n+1)\Psi(r+t_0+2g_0(t_0))$ vertical edges. We deduce that

$$\begin{aligned} d_Z^\epsilon(e,f) &\le n+\epsilon(n+1)\Psi(r+t_0+2g_0(t_0)) \\ &\le r+t_0+1+\epsilon(r+t_0+2)\Psi(r+t_0+2g_0(t_0)) \\ &= T_0(r)+\epsilon F_0(r), \end{aligned}$$

where we set $T_0(r)=r+t_0+1$ and $F_0(r)=(r+t_0+2)\Psi(r+t_0+2g_0(t_0))$. ♢

Lemma 7.8 : *There is a linear function, T_1 and an increasing function F_1 such that if π is a geodesic segment in Z^ϵ and α is a subpath of $\sigma(\pi)$, then α contains at most $T_1(r)+\epsilon F_1(r)$ geodesic segments, where $r=\mathrm{length}(\mathrm{straight}(\alpha))$.*

Here, and in what follows, we can interpret "geodesic" in Z^ϵ broadly to mean a path, π, of length n, such that the distance between the midpoints of the initial and final edges is exactly $n-1$. (We allow the distance between the endpoints of π to be $n-1$.) This makes no difference to the arguments that follow.

Proof : Let x,x' be the endpoints of α so that $r=\rho(x,x')$. Let β,β' be the segments of $\sigma(\pi)$ containing x,x' respectively. Applying Lemma 7.4, we can find endpoints, y,y' of β,β' respectively with $\rho(x,y)\le r+2t_1$ and $\rho(x',y')\le r+2t_1$. (For this, we note that the projection of any geodesics in Z^ϵ is an arc in K. Moreover, we can assume that $\beta\ne\beta'$ and so the horoballs associated to β and β' are distinct.) We thus have $\rho(y,y')\le 3r+4t_1$. Now y,y' are breakpoints of $\sigma(\pi)$, in other words, $y=\mu(e)$ and $y'=\mu(e')$, where e,e' are horizontal edges of π. By Lemma 7.6, we have $d_Z^\epsilon(e,e')\le T_0(3r+4t_1)+\epsilon F_0(3r+4t_1)$. But π is assumed to be geodesic in Z^ϵ. Thus, in particular, the number of horizontal segments in π between e and e' is at most $d_Z^\epsilon(e,e')$. Now any breakpoint of α corresponds either to such an edge or to e or e'. Thus the total number of segments of α is at most $d_Z^\epsilon(e,e')+3\le T_0(3r+4t_1)+3+\epsilon F_0(3r+4t_1)$. We thus set $T_1(r)=T_0(3r+4t_1)+3$ and $F_1(r)=F_0(3r+4t_1)$. ♢

Lemma 7.9 : *There is some $\epsilon>0$ and a linear function Q such that if π is a geodesic in Z^ϵ, then $\sigma(\pi)$ is Q-quasigeodesic in $\mathbf{H}^\nu$.*

Proof : We shall show that there is some linear function Q_0 such that for all $r\ge0$, there is some $\epsilon>0$ such that if π is geodesic in Z^ϵ, then $\sigma(\pi)$ is r-locally Q_0-quasigeodesic. We can then choose r_1 sufficiently large so that any r_1-locally Q_0-quasigeodesic is globally Q-quasigeodesic, and then fix ϵ appropriately.

Suppose then that $\alpha\subseteq\sigma(\pi)$ is a subpath of length at most r, and suppose, in turn, that β is a subpath of α. Let $\gamma=\mathrm{straight}(\beta)$ and set $s=\mathrm{length}(\gamma)$. Thus $s\le r$. Let n be the number of geodesic segments in γ.

Since α is geodesic in Z^ϵ, its projection to L is an arc. Thus, by Lemma 7.6, $\mathrm{back}(\beta)\le\mathrm{back}(\sigma(\pi))\le t_2$. By 7.5(4), it follows that there is a subexponential function, S, depending only on t_2, such that $\mathrm{length}(\beta)\le s+S(n)$. Moreover, by Lemma 7.8, we have $n\le T_1(s)+\epsilon F_1(s)\le T_1(s)+\epsilon F_1(r)$. Combining these facts, we have $\mathrm{length}(\beta)\le s+S(T_1(s)+\epsilon F_1(r))$.

Now define a function R by $R(x) = S(T_1(x)+1)$. Thus R is subexponential, and so, by 7.5(3) determines a linear function Q_0 so that any path satisfying the relation of 7.5(3) with respect to R is Q_0-quasigeodesic. By 7.5(2), there is some r_1, and a linear function Q, such that any r_1-locally Q_0-quasigeodesic is Q-quasigeodesic.

Thus if $\epsilon \leq 1/F_1(r_1)$, then $\epsilon F_1(r) \leq \epsilon F_1(r_1) \leq 1$, and so $\text{length}(\beta) \leq s + S(T_1(s)+1) = R(s)$. Since this applies to any subpath of β of α, by 7.5(3), β is Q_0-quasigeodesic, and so $\sigma(\pi)$ is r_1-locally Q_0-quasigeodesic and hence Q-quasigeodesic as required.

To summarise the logic, t_2 is given by Lemma 7.6, and T_1, F_1 by Lemma 7.8. By 7.5(4), t_2 gives us S, and we define R in terms of S and T_1. By 7.5(3), this gives Q_0, and 7.5(2) in turn gives r_1 and Q. Finally choose any $\epsilon \leq 1/F_1(r_1)$. ♢

We now fix ϵ and Q as given by Lemma 7.9 (each depending only on r_0, t_0 and Ψ). Let h_1 be the constant given by 7.5(1) so that any Q-quasigeodesic α is within a Hausdorff distance h_1 of straight(α).

Lemma 7.10 : *There is a constant k_1 with the following property. Suppose that $p \in \Pi$ and that e is a horizontal edge of Z. Suppose that $x, x' \in l(p)$ and that π, π' are geodesics in Z^ϵ, each with initial edge e and terminating at x and x' respectively. Suppose that $\pi \cap l(p) = \{x\}$ and that $\pi' \cap l(p) = \{x'\}$. Then $d^\epsilon_Z(x, x') \leq k_1$.*

Proof : Let f, f' be the final edges of π, π' respectively. Thus f and f' are horizontal. Their projections to K are both incident on p. Let $\beta = \text{straight}(\sigma(\pi)) = [\mu(e), \mu(f)]$ and $\beta' = \text{straight}(\sigma(\pi')) = [\mu(e), \mu(f')]$. Let $\gamma = [\mu(f), \mu(f')]$. Thus (β, β', γ) is a geodesic triangle in $\mathbf{H}^\nu$, and so every point of γ lies in a $\log(1+\sqrt{2})$-neighbourhood of $\beta \cup \beta'$ and thus within a distance $h_1 + \log(1+\sqrt{2})$ of $\sigma(\pi) \cup \sigma(\pi')$. Now p is not an interior vertex of either π or π', and so by Lemma 7.3(4), $\sigma(\pi) \cup \sigma(\pi')$ can enter at most a distance t_1 into $B(p)$. Thus, γ enters at most $h_2 = h_1 + \log(1+\sqrt{2}) + t_1$ into $B(p)$. By Lemma 7.3(3), $\rho(\mu(f), \mu(f')) = \text{length}(\gamma) \leq 2h_2 + t_1$. Thus, by Property (P2), the vertical distance between f and f', i.e. $d'(x, x')$ is at most $\Psi(2h_2 + t_1)$ and so $d^\epsilon_Z(x, x') \leq \epsilon\Psi(2h_2 + t_1)$. We therefore set $k_1 = \epsilon\Psi(2h_2 + t_1)$. ♢

Proof of Proposition 7.1 : In order to show that Z^ϵ is hyperbolic, we show that every geodesic triangle has a "centre", that is, a point a bounded distance from all three edges. We can assume that the verices of the triangle at midpoints of horizontal edges (since the set of such midpoints is quasidense in Z^ϵ).

Suppose then that (π_1, π_2, π_3) is a geodesic triangle in Z^ϵ, i.e. geodesics cyclically connecting three horizontal edges. (By this we mean that there are horizontal edges, a_1, a_2, a_3 such that the initial and final edges of π_i are a_{i+1} and a_{i+2}, taking subscripts mod 3, and that π_i is a geodesic in the sense described after the statement of Lemma 7.8.) We distinguish two cases.

Case (1) : There is some point $p \in \Pi$ such that for each i, $\pi_i \cap l(p) \neq \emptyset$.

Let $l_i = \pi_i \cap l(p)$. Applying Lemma 7.10, we see that $d^\epsilon_Z(l_i, l_j) \leq k_1$ for all i, j. Thus, there is some $x \in l(p)$ with $d^\epsilon_Z(x, l_i) \leq k_1$ for each i. We see that x is a centre of (π_1, π_2, π_3) as required.

Case (2) : Each vertical line is missed by at least one π_i.

Let $\beta_i = \text{straight}(\sigma(\pi_i))$. Thus $\text{hd}(\beta_i, \sigma(\pi_i)) \leq h_1$. Now $(\beta_1, \beta_2, \beta_3)$ is a geodesic triangle in $\mathbf{H}^\nu$, so there are points $x_i \in \sigma(\pi_i)$ such that $\rho(x_i, x_j) \leq h_3 = 2h_1 + \log(1+\sqrt{2})$.

Suppose $p \in \Pi$. At least one of π_1, π_2, π_3, say π_1, misses $l(p)$. In other words, p does not lie in the projection of π_1 to L. Thus by Lemma 7.3(4), $\sigma(\pi_1)$ enters at most a distance t_1 into $B(p)$. It follows that each of x_1, x_2, x_3 lies a distance at most $h_3 + t_1$ inside $B(p)$.

In particular, suppose x_i lies in the segment α_i of $\sigma(\pi_i)$. Now α is associated to some $p_i \in \Pi$. Since x_i lies at most a distance $h_3 + t_1$ inside $b(p_i)$, Lemma 7.3(2) tells us that some endpoint, $\mu(e_i)$, say, of α_i satisfies $\rho(x_i, \mu(e_i)) \leq h_3 + 2t_1$. Thus $\rho(\mu(e_i), \mu(e_j)) \leq h_3 + 2(h_3 + 2t_1) = 3h_3 + 4t_1$. By Lemma 7.7, $d^\epsilon_Z(e_i, e_j) \leq T_0(3h_3 + 4t_1) + \epsilon F_0(3h_3 + 4t_1)$. Since e_i is an edge of π_i, we have found a centre for (π_1, π_2, π_3) as required.

We have shown that Z^ϵ is hyperbolic. Since Z is quasi-isometric to Z^ϵ, it follows that Z is hyperbolic. ♢

Put another way, we have shown that the stratified graph, $L = \bigcup_i L_i$ satisfies (S2). We know by Lemma 2.2 that it also satisfies (S3) (fineness). We can also verify a strong version of (S4) namely:

(S4′) There is a function, F', such that if $p \in \Pi$ and $e \in E(L_i)$ and $f \in E(L_j)$ are both incident on p, then $|i - j| \leq F'(\angle_p(e, f))$.

(Here $\angle_p$ denotes angle in L.)

It is easy to see that this implies (S4), for if $p, q \in \Pi$ and $r \in \mathbf{N}$ with $d_i(p, q) \leq r$ and $d_j(p, q) \leq r$, we connect p to q by geodesics α, β in L_i and L_j. Thus $\angle_p(\alpha, \beta) \leq 2r - 2$ and so $|i - j| \leq F'(2r - 2)$.

Lemma 7.11 : *If L satisfies properties (P1) and (P2), then it satisfies (S4′).*

Proof : Let p, e, f be as in the statement of (S4′). Let α be an arc of length $n = \angle_p(e, f)$ connecting e to f in $L \setminus \{x\}$. Thus, $\gamma = e \cup \alpha \cup f$ is circuit of length $n + 2$. Now $\sigma(\gamma)$ is a closed broken geodesic in $\mathbf{H}^3$. By Lemma 7.6, $\text{back}(\sigma(\gamma)) \leq t_2$. Applying 7.5(4) to γ, viewed as a path with both endpoints at a breakpoint, we see that $\text{length}(\gamma)$ and hence $\rho(\mu(e), \mu(f))$ is bounded in terms of n. By property (P2) it follows that $|i - j|$ is bounded in terms of n. ♢

We shall be applying the results of this section in a form intrinsic to the complement of the horoball system.

Let $(B(p))_{p \in \Pi}$ be a system of disjoint horoballs. Let $Y = \mathbf{H}^\nu \setminus \bigcup_{p \in \Pi} \text{int}\, B(p)$. Thus Y is closed, and we denote the induced path metric by ρ_Y. (We remark that Y is a CAT(0) space, though we make no explicit use of this fact. Indeed this is a special feature of constant curvature.) The embedding of Y in $\mathbf{H}^\nu$ is uniformly proper. Indeed, if $x, y \in Y$, then $\rho(x, y) \leq \rho_Y(x, y) \leq e^{\rho(x,y)}$.

Suppose that $\mathcal{C}_i$ is a set of arcs in Y, each connecting a pair of distinct horoballs. Define the graph L_i with vertex set Π by deeming $p, q \in \Pi$ to be adjacent if $\partial B(p)$ and $\partial B(q)$ are connected by an arc in $\mathcal{C}_i$. Let $\mathcal{C} = \bigcup_i \mathcal{C}_i$ and let $L = \bigcup_i L_i$, where i varies over a set, $\mathcal{I}$, of consecutive integers. Suppose that $r, t, u \geq 0$ and that Φ is an increasing function. We introduce the following hypotheses.

(A1)(r) $(B(p))_p$ is r-quasidense (with respect to the metric ρ_Y).

(A2)(t) If $\rho_Y(\partial B(p), \partial B(q)) \leq t$, then there is an arc in $\mathcal{C}$ connecting $\partial B(p)$ to $\partial B(q)$.

(A3)(u) Every arc of $\mathcal{C}$ has length at most u.

(A4)(Φ) If $\alpha \in \mathcal{C}_i$ and $\beta \in \mathcal{C}_j$, then $|i - j| \leq \Phi(\rho_Y(\alpha, \beta))$.

Proposition 7.12 : *Given r, there exists t such that for all u and Φ, if $(B(p))_p$ and $(\mathcal{C}_i)_i$ satisfy (A1)(r), (A2)(t), (A3)(u) and (A4)(Φ), then the stratified graph $(L_i)_i$ satisfies (S2), (S3) and (S4).*

Proof : In view of Proposition 7.1, Lemma 2.2 and Proposition 7.11, it is sufficient to verify properties (P1) and (P2).

We choose $t \geq e^{2r+1}$. Now $(B(p))_p$ is certainly r-quasidense in $\mathbf{H}^\nu$. Moreover, if $\rho(B(p), B(q)) \leq 2r+1$, then $\rho_Y(B(p), B(q)) \leq e^{2r+1} \leq t$, so by (A2), p and q are adjacent in L. In other words, $L(2r+1) \subseteq L$. Moreover, (A3) tells us that $L \subseteq L(u)$, so property (P1) holds.

For (P2), suppose that $e \in E(L_i)$, $f \in E(L_j)$ and let $\alpha \in \mathcal{C}_i$ and $\beta \in \mathcal{C}_j$ be the corresponding arcs. Let $s = \rho(\mu(e), \mu(f))$ Let x, y be any points on α, β respectively. By Lemma 7.2, we have $\rho(x, \mu(e)) \leq g_0(u)$ and $\rho(y, \mu(f)) \leq g_0(u)$. Thus $\rho(x, y) \leq s + 2g_0(u)$ and so $\rho_Y(\alpha, \beta) \leq \rho_Y(x, y) \leq e^{s+2g_0(u)}$.

By (A4) it follows that $|i - j| \leq \Phi(e^{s+2g_0(u)})$ and so we can set $\Psi(s) = \Phi(e^{s+2g_0(u)})$. $\diamondsuit$

8. Simply degenerate ends.

In this section, we consider the geometry of a simply degenerate end of a hyperbolic 3-manifold, and explain how such an end gives rise to a stack of Farey graphs.

Given a hyperbolic manifold $N = \mathbf{H}^\nu/\Gamma$ we write $\text{inj}(N)$ for the injectivity radius away from the cusps, and $\Pi \subseteq \partial \mathbf{H}^\nu$ for the set of parabolic points. Here we are interested in surface groups, and so the stabiliser, $\Gamma(p)$ of each $p \in \Pi$ is infinite cyclic.

In dimensions $n = 2$ and $n = 3$, the Margulis Lemma gives us a universal $\eta_0 > 0$ such that we can always find a Γ-invariant system $(B(p))_{p \in \Pi}$ of horoballs such that $\rho(B(p), B(q)) \geq 1$ for $p \neq q$ and such that $\partial B(p)$ is a circle of length η_0 if $n = 2$, and a euclidean product of such a circle with the real line if $n = 3$. Let $Y = \mathbf{H}^\nu \setminus \bigcup_{p \in \Pi} \text{int}(B(p))$ and $M = M(\Gamma) = Y/\Gamma \subseteq N$.

Let Σ be our surface, with boundary components $(C^m)_{m \in \mathcal{P}}$. Let $\mathcal{S}(\eta)$ be the moduli space of unmarked hyperbolic structures on Σ such that each boundary curve C^m is horocycle of length at least η_0 and $\text{inj}(\Sigma) \geq \eta$. (We can always choose $\eta < \eta_0$ so that the notion of injectivity radius is unambiguous.) It is well known that $\mathcal{S}(\eta)$ is compact. Thus there is some $D(\eta, \chi)$ so that $\text{diam}(\Sigma) \leq D(\eta, \chi)$ for all $\Sigma \in \mathcal{S}(\eta)$. We also note:

Lemma 8.1 : *There is some $\lambda = \lambda(\eta, \chi)$ such that if $\Sigma, \Sigma' \in \mathcal{S}(\eta)$, then there is a λ-bilipschiz homeomorphism from Σ to Σ'.* $\diamondsuit$

For this reason it will eventually be convenient to restrict to some "favourite" hyperbolic structure on Σ.

One particular structure is constructed as follows. Let Θ be the surface obtained by collapsing each boundary component of Σ to a point. Let $P \subseteq \Theta$ be the set of such points (cf. Section 6). Now take a triangulation of Θ with vertex set P. For each simplex of this triangulation, we take a hyperbolic right-angled hexagon three of whose sides are geodesics of length 1, alternating with three horocyclic sides (each of length $\sqrt{e}$). Such a hexagon is well-defined up to isometry. We now glue together these hexagons along their geodesic edges to obtain a hyperbolic surface, $\Sigma \in \mathcal{S}(\eta_0)$. (We may as well suppose that $\eta_0 < \sqrt{e}$.)

We write $\mathcal{A}_0$ for the set of (homotopy classes of) geodesic arcs in Σ arising from the geodesic edges of our triangulation. Given $t \geq 1$, we shall write $\mathcal{B}_0(t) \supseteq \mathcal{A}_0$ for the set of homotopy classes of arcs connecting boundary components which admit representatives of length at most t. Thus, $\mathcal{B}_0(t)$ is finite. Note that in the universal cover, $\tilde{\Sigma}$, $\mathcal{A}_0$ lifts to a set of arcs connecting distinct horocycles. Collapsing each horocycle to a point, we get a Farey graph, K_0, with vertex set identified with Π (cf. Section 6). Similarly, $\mathcal{B}_0(t)$ lifts to a graph $L_0(t)$ in which K_0 is quasi-isometrically embedded.

Now let us suppose that we have a strictly type-preserving action of $\Gamma = \pi_1(\Sigma)$ on $\mathbf{H}^3$. Let $N = N(\Gamma)$ and let $M = M(\Gamma)$ be defined as above. Write $(S^m)_{m \in \mathcal{P}}$ for the set of boundary components of M. Note that there is a natural homotopy class of maps from Σ into M sending C^m to S^m. For homological reasons, M has two ends.

Now the work of Bonahon [Bon] tells us that each end of M is either geometrically finite or simply degenerate. To simplify the exposition, we shall assume for the time being that both ends are simply degenerate. The case with geometrically finite ends only calls for simple modification which we shall explain at the end. In the doubly degenerate case, $\Lambda\Gamma = \partial\mathbf{H}^3$.

We have the following result of Thurston-Bonahon:

Theorem 8.2 : *There is a uniform constant $\mu > 0$ such that for all $x \in M$, there is some $\epsilon > 0$ and a hyperbolic surface $\Sigma_x \in \mathcal{S}(\epsilon)$ together with a μ-lipschitz map $\phi_x : \Sigma_x \longrightarrow M$ such that $x \in \phi_x(\Sigma_x)$ and $\phi_x(C^m) \subseteq S^m$ for all $m \in \mathcal{P}$, and which is a relative homotopy equivalence of $(\Sigma, \bigcup_m C^m)$ to $(M, \bigcup_m S^m)$.* $\diamondsuit$

In fact, Bonahon's tameness theorem [Bon] together with Thurston's interpolation of pleated surfaces [T1], gives us a proper 1-lipschitz map of a complete hyperbolic surface into N. It is not hard to see that such a map behaves nicely inside horoballs, and so restricts to give us a lipschitz map from Σ into M. We could thus take $\eta = 1$ in this case, though later we will want to allow for other lipschitz constants. (Note that we have no need to assume ϕ_x depends continuously on x.)

Suppose now that $\mathrm{inj}(N) \geq \eta > 0$. It follows that each $\Sigma_x \in \mathcal{S}(\eta/\mu)$. Thus $\mathrm{diam}(\Sigma) \leq D(\eta/\mu, \chi)$, and so every point of M is at most $\mu D(\eta/\mu, \chi)$ from $\bigcup_m S^m$. In other words, $\bigcup_m S^m$ is uniformly quasidense in M.

Let $N^m = N(S^m, \eta/3)$ denote the $(\eta/3)$-neighbourhood of S^m in M. Note that N^m is a topological product of S^m with an interval. There is natural nearest point retraction,

$\pi_{S^m} : N^m \longrightarrow S^m$. If $A \subseteq N^m$ then $\mathrm{vol}(\pi_{S^m}^{-1}(A)) = c\,\mathrm{area}(A)$, where c is a constant depending only on η.

Now the result of Freedman-Hass-Scott [FHS] (see also [Bon]) tells us that in any neighbourhood of $\phi_x(\Sigma_x)$, we can find an embedded surface, T_x, in M whose inclusion into M is also a relative homotopy equivalence. We can assume that $x \in T_x$, and that $\mathrm{hd}(T_x, \phi_x(\Sigma_x)) \leq 1$. Thus, $\mathrm{diam}(T_x) \leq D = \mu D(\eta/\mu, \chi) + 2$.

After local adjustments, we can assume that for each $m \in \mathcal{P}$, both T_x and $\phi_x(\Sigma_x)$ meet C^m in a closed geodesic σ^m_x (of length μ). Moreover, we can assume that $T_x \cap N^m = \phi_x(\Sigma_x) \cap N^m = \pi_{N^m}^{-1}(\sigma^m_x)$ is a geodesically embedded annulus in N^m.

We note that if $T_x \cap T_y = \emptyset$, then $T_x \cup T_y$ bounds a compact subset, $R_{x,y}$, of M. Waldhausen's h-cobordism theorem tells us that $R_{x,y}$ is homeomorphic to $\Sigma \times [0,1]$. From this one deduces that M is homeomorphic to $\Sigma \times \mathbf{R}$ as remarked in the Introduction.

We now choose a sequence, $(x_i)_{i \in \mathbf{Z}}$ in M, so that writing $T_i = T_{x_i}$, we have $\rho(T_i, T_{i+1}) = D$ for all i, and T_i separates T_{i-1} for T_{i+1} in M. We write $\Sigma_i = \Sigma_{x_i}$, $\phi_i = \phi_{x_i}$ and $\sigma^m_i = \sigma^m_{x_i}$. We write $R_i = R_{x_i, x_{i+1}}$ for the region bounded by T_i and T_{i+1}, and $A^m_i = S^m \cap R_i$. Thus, A^m_i is an annulus in S^m bounded by the euclidean geodesic curves σ^m_i and σ^m_{i+1}.

Lemma 8.3 : $\mathrm{diam}(R_i) \leq 5D$.

Proof : Let α be an arc of length D connecting T_i to T_{i+1}. Thus $\mathrm{diam}(T_i \cup \alpha \cup T_{i+1}) \leq 3D$. If $x \in R_i$, then $T_x \cap (T_i \cup \alpha \cup T_{i+1}) \neq \emptyset$ and $\mathrm{diam}(T_x) \leq D$. The result follows. ♢

Lemma 8.4 : *There is a function, F, such that if $Q \subseteq M$ is a compact subset then there is a compact submanifold $R \subseteq M$ with $\mathrm{diam}\, R \leq F(\mathrm{diam}(Q))$ such that the inclusion of $(R, R \cap \bigcup_m C^m)$ into $(M, \bigcup_m C^m)$ is a relative homotopy equivalence.*

Proof : Q lies in a subset of the form $\bigcup_{i \leq k \leq j} R_k$ with $j - i$ bounded in terms of $\mathrm{diam}(Q)$. ♢

Lemma 8.5 : *There is a constant, D_1, such the euclidean length of each A^m_i is at most D_1.*

By the "euclidean length" we mean the distance between the boundary curves σ^m_i and σ^m_i as measured in A^m_i.

Proof : The uniform neighbourhood, $N^m \cap \pi_{S^m}^{-1}(A^m_i)$ lies in R_i. Moreover, $\mathrm{vol}(\pi_{S^m}^{-1}(A^m_i)) = c\,\mathrm{area}(A^m_i) = c\eta_0 l$, where l denotes the length of A^m_i. Since the diameter of R_i is bounded, so is its volume. This puts a bound on l as required. ♢

Although we shall not explicitly be using the fact, we remark that an easy consequence is the following:

Proposition 8.6 : *The inclusion of S^m into M is a uniform quasi-isometry.* $\diamondsuit$

As usual, the constants depend only on η and χ.

In the above discussion, we allowed the domain surfaces, Σ_x to vary in moduli space $\mathcal{S}(\eta)$. In what follows, it will be convenient to assume that Σ_x has a fixed structure (modulo the action of $\mathcal{M}$.) In particular, we can take it to be our "favourite" surface described above. After adjusting the lipschitz constant, μ, there is no loss of generality in this. In fact, by Lemma 8.1, we can simply precompose each map ϕ_x by a with a bilipschitz homeomorphism of Σ to Σ_x. This increases the bilipschitz constant by at most a factor of $\lambda(\eta/\mu, \xi)$. We shall therefore henceforth assume we have a sequence of μ-bilipschitz maps, $\phi_i : \Sigma \longrightarrow M$.

The following is a key fact in the analysis of the structure of ends:

Proposition 8.7 : *Giving μ, η, χ, there is a function H with the following property. Suppose that $\phi : \Sigma \longrightarrow M$ is a μ-lipschitz map which is a relative homotopy equivalence from $(\Sigma, \bigcup_m C^m)$ to $(M, \bigcup_m S^m)$. Suppose that $x, y \in \Sigma$ and that β is a path connecting $\phi(x)$ to $\phi(y)$. There is a path α in Σ connecting x to y such that* $\mathrm{length}(\alpha) \leq H(\mathrm{length}\,\beta)$ *and $\phi \circ \alpha$ is homotopic in M relative to $\{\phi(x), \phi(y)\}$.*

The above result is sometimes expressed by saying that the lift of ϕ to the universal covers is uniformly proper (cf. [Min1]).

Note that the relative homotopy class of α is determined, so we can always take α to be the unique geodesic representative thereof. Also, since the structure on Σ is fixed, we can assume without loss of generality that $x = y$.

To prove this result, we shall view it as a statement about the set of all manifolds, M, arising in this way. That is, M is obtained by removing horoballs from a complete hyperbolic 3-manifold N with $\mathrm{inj}(N)$ bounded below. In fact, the essential points are that M is 3-dimensional, and has bounded local geometry. The latter arises from a control of curvature, a lower bound on injetivity radii and a lower bound on the distance between boundary components. In this situation, we have:

Proposition 8.8 : *Let (Υ, d) be a compact path-metric space, and $a \in \Upsilon$. Suppose we have a sequence of 3-manifolds, M_n, as above, together with a sequence, $f_n : \Upsilon \longrightarrow M$ of uniformly lipschitz maps. Then we can find a subsequence, f_{n_i}, such that for all $\epsilon > 0$ and $r \geq 0$, there is some constant k such that for all $i, j \geq k$, there is a continuous map $h : N(f_{n_i}(a), r) \longrightarrow M_{n_j}$ such that for all $x \in \Upsilon$, $\rho_{n_j}(\phi_{n_j}(x), h \circ \phi_{n_i}(x)) \leq \epsilon$.*

Here, $N(b, r)$ denotes the uniform r-ball about b (in M_{n_i}).

In fact, one can say a lot more. By Gromov's $C^{1,1}$ convergence theorem [Gr2, Section 8.20], one can pass to a subsequence that converges in the bilipschitz sense on larger and larger balls. One can then apply compactness of uniformly bilipschitz (hence equicontinuous) maps into the limit manifold.

For our purposes, we need much less. Indeed, precompactness of the space of such manifolds in the Hausdorff topology suffices (see [Gr2, Proposition 5.2]). This does not require a lower bound on injectivity radius (only an upper bound on the volume growth of balls). The lower bound on injectivity radius allows us to subsequently construct a continuous map between Hausdorff-close manifolds. For this one can use a partition of unity argument. Finally, it is a simple exercise to incorporate information about the maps ϕ_i into the Hausdorff precompactness, so that the images of these maps are also arbitrarily close.

We note that although we have phrased it in terms of a limiting argument, the reasoning is essentially combinatorial using uniform nets in the various spaces (as in the proof of [Gr2, Proposition 55.2]). It can thus, in principle, be used to give explicit bounds on the various constants involved.

Proof of Proposition 8.7 : Suppose the result fails. Fix a basepoint $a \in \Sigma$. We can find a sequence of μ-lipschitz maps $\phi_n : \Sigma \longrightarrow M_n$ together with loops β_n based at $b_n = \phi_n(a) \in M_n$ such that if α_n is the geodesic in Σ based at a with $\phi_n \circ \beta_n$ homotopic to β_n, then $\text{length}(\alpha_n) \longrightarrow \infty$. In particular, we can assume that the α_n are all distinct.

Let Υ be the metric space obtained by taking the unit interval, $[0,1]$ and identifying both its endpoints with $a \in \Sigma$. (Thus, Υ is a homeomorphic to the wedge of Σ and a circle.) We can realise β_n as a constant speed path $\beta_n : [0,1] \longrightarrow M_n$. Combining the maps ϕ_n and β_n, we get a sequence of uniformly lipschitz maps $f_n : \Upsilon \longrightarrow M_n$. Now pass to the subsequence given by Proposition 8.8.

Choose ϵ less that the lower bound, η, on $\text{inj}(M_n)$, and choose $r_0 \geq 0$ to be as determined shortly. Thus, for sufficiently large $m < n$, we a have a map $h : N(b_m, r_0) \longrightarrow M_n$ such that for all $x \in \Upsilon$, $\rho_n(f_n(x), h \circ f_m(x)) \leq \epsilon$. In particular, $\rho_n(b_n, h(b_m)) \leq \epsilon$. For notational convenience in what follows, let us assume that $h(b_m) = b_n$. This can be achieved by pussing f_n along the geodesic segment connecting them. We shall use $\simeq$ to denote homotopy relative to the points a, b_m of b_n, according to context.

For each $x \in \Upsilon$ there is a unique shortest geodesic from $h \circ f_m(x)$ to $f_n(x)$ which varies continuously in x. Thus $h \circ \beta_m \simeq \beta_n$ and $h \circ \phi_m \simeq \phi_n$ in M_n. Also, by hypothesis, $\beta_n \simeq \phi_n \circ \alpha_n$ in M_n and $\beta_m \simeq \phi_m \circ \alpha_m$ in M_m. If we can show that in fact

$(*)$ $\beta_m \simeq \phi_m \circ \alpha_m$ in $N(b_m, r_0)$,

then it follows that $h \circ \beta_m \simeq h \circ \phi_m \circ \alpha_m$. Thus,

$$\phi_n \circ \alpha_n \simeq \beta_n \simeq h \circ \beta_m \simeq h \circ \phi_m \circ \alpha_m \simeq \phi_n \circ \alpha_m.$$

By hypothesis, $\phi_n : \Sigma \longrightarrow M_n$ is a homotopy equivalence, so $\alpha_m \simeq \alpha_n$ in Σ. But we have arranged that the α_n all lie in distinct homotopy classes, so this is a contradiction.

So far, we have only made essential use of bounded local geometry of our manifolds, M_n. However, it remains to justify $(*)$, and it is here that we use the fact that we are dealing with 3-manifolds. Note that since we are free to choose any fixed r_0, this amounts to asserting that we can bound the diameters of the homotopies from $\phi_n \circ \alpha_n$ to β_n. But for this it is sufficient to note that the sets $\phi_n(\Upsilon)$ have bounded diameter in M_n. Thus, we can use Lemma 8.4 to push the homotopies in M_n into subsets of bounded diameter as required. $\diamondsuit$

Remarks : Our argument runs parallel to that of Minsky [Min1] in the closed surface case. In place of Gromov-Hausdorff convergence, Minsky uses Thurston's compactness of pleated surfaces in the geometric topology (cf. [CanaEG]). The argument about bounding homotopies corresponds to showing convergence in the algebraic topology in [Min1]. It is interesting to speculate to what extent Proposition 8.8 could be generalised by finding substitutes for Thurston-Bonahon in Lemma 8.4. For example is it true in $\mathbf{H}^\nu$ for $n > 3$?

We now return to our fixed manifold, M.

Lemma 8.9 : *There is a uniform function H_0 such that if β is a path in M connecting S^m to S^n with $\beta \cap R_i \neq \emptyset$, then there is a path α is Σ connecting C^m to C^n with $\phi_i \circ \alpha$ homotopic to β in M relative to $S^m \cup S^n$, and with* $\mathrm{length}(\alpha) \leq H_0(\mathrm{length}(\beta))$.

Proof : The successive surfaces, T_i, are distance D apart. Thus, β lies in a set of the form $\bigcup_{j\in\mathcal{J}} R_j$ where $\mathcal{J} \subseteq \mathcal{I}$ is a finite subset of length at most $2 + (\mathrm{length}(\beta))/D$ and containing i. In particular, the endpoints of β lie in $\bigcup_{j\in\mathcal{J}} A^m_j$ and $\bigcup_{j\in\mathcal{J}} A^n_j$ respectively. Applying Lemma 8.6, we can find paths γ and δ in S^m and S^n respectively connecting these endpoints to σ^m_i and σ^n_i. By Proposition 8.7, there is a path α in Σ from C^m to C^n such that $\phi_i \circ \alpha$ is homotopic relative to the endpoints of $\gamma \cup \alpha \cup \delta$ in M and with $\mathrm{length}(\alpha) \leq H(\mathrm{length}(\gamma \cup \alpha \cup \delta))$. Thus $\mathrm{length}(\alpha)$ is bounded in terms of $\mathrm{length}(\beta)$. $\diamondsuit$

We have observed that $\bigcup_m S^m$ is $(D/2)$-dense in M. Letting $r_0 = D/2$, we see that the horoball system $(B(p))_{p\in\Pi}$ is r_0-dense in $\mathbf{H}^3$. Let t_0 be the constant given by Proposition 7.12, and let $\tau_0 = H_0(t_0)$, where H_0 is the function given by Lemma 8.9. Note that each of these depends only on η and χ.

We shall construct a stack of Farey graphs as follows. Let $[\phi_i]$ be the class of maps from $(\Sigma, \bigcup_m C^m)$ to $(M, \bigcup_m S^m)$ homotopic to ϕ_i. Thus $[\phi_0]^{-1}[\phi_i]$ determines an element, ψ_i, of $\mathcal{M}$, the mapping class group of $(\Sigma, \bigcup_m C^m)$. Let $K = \bigcup_i K_i$, where $K_i = K(\psi_i)$ as defined in Section 5. We need to verify that K satisfies properties (S1)–(S4) (hence also (S5) and (S6)).

We can also construct $L_0 = L_0(\tau_0)$ from the set of arcs $\mathcal{B}_0 = \mathcal{B}_0(\tau_0)$ as described earlier. Pushing forward under ψ_i, we similarly obtain a stratified graph $L = \bigcup_i L_i$. Each K_i is uniformly quasi-isometrically embedded in L_i and so $Z(K)$ is quasi-isometrically embedded in $Z(L)$.

We can also describe L as follows. Let $\tilde{\phi}_i : \tilde{\Sigma} \longrightarrow \tilde{M} \subseteq Y \subseteq \mathbf{H}^3$ be a lift of ϕ_i, and let $\mathcal{C}_i$ be the set of arcs of the form $\tilde{\phi}_i \circ \alpha$ for $\alpha \in \mathcal{B}_0$. Let $\mathcal{C} = \bigcup_i \mathcal{C}_i$. Thus L_i can be alternatively thought of as obtained from $\mathcal{C}_i$ where $p, q \in \Pi$ are adjacent if $B(p)$ and $B(q)$ are connected by an arc in $\mathcal{C}_i$.

We need to verify properties (A1)–(A4). We know that $(B(p))_p$ is r_0-quasidense, and so (A1)(r_0) is satisfied. For (A2), suppose $\rho_Y(B(p), B(q)) \leq t_0$. In other words, there is a path of length at most t_0 connecting $B(p)$ to $B(q)$ in Y. Let β be its projection to M. By Lemma 7.9, β is homotopic relative $\bigcup_m S^m$ to an arc of the form $\phi_i \circ \alpha$, where $\mathrm{length}(\alpha) \leq H_0(\tau_0)$, and so $\alpha \in \mathcal{B}_0(\tau_0)$. Thus, α lifts to an arc in $\mathcal{C}_i \subseteq \mathcal{C}$ as required. This shows that (A2)(t_0) holds. Since each of the maps ϕ_i is μ-lipschitz, (A3)(u) holds with $u = \mu\tau_0$. Finally, for (A4), suppose $\alpha \in \mathcal{C}_i$ and $\beta \in \mathcal{C}_j$. Let γ be a path $\rho_Y(\alpha, \beta)$ connecting.

We see that γ can cross at most $\rho(\alpha,\beta)/D$ regions R_k and so $j-i \leq \Psi(\rho_Y(\alpha,\beta))$ where $\Psi(x) = x/D$.

By Proposition 7.12, we deduce that L satisfies (S2), (S3) and (S4). It follows immediately that K does also.

We still need to verify that K satifies (S1). But the argument for (A2) above shows that if two points are adjacent in K_i they are also adjacent in L_{i-1} and in L_{i+1}. They are thus a bounded distance apart in K_{i-1} and K_{i+1}.

Putting this together we Lemma 5.6, conclude:

Proposition 8.10 : *The stratified graph K satisfies properties (S1)–(S6).* $\diamondsuit$

Now Proposition 5.5 gives us a Γ-equivariant map from $\Delta^0 K$ to K. Moreover, applying Lemmas 2.1 and 2.3 (cf. the discussion at the end of Section 6), we can equivariantly identify ΔK with $\partial \mathbf{H}^3$.

We have thus proven Theorem 0.1 in the doubly degenerate case.

We still have to consider the case where at least one of the ends is geometrically finite.

If both ends are geometrically finite, then the group is quasifuchsian, and it well-known that there is an equivariant homeomorphism of $\Delta\Sigma$ to the limit set.

It remains to consider the singly degenerate case. Let N' be the convex core, i.e. $N' = (\operatorname{hull}\Lambda)/\Gamma$ and let $M' = (Y \cap \operatorname{hull}\Lambda)/\Gamma$. Now N' has one boundary component which is a finite-area surface homeomorphic to $\operatorname{int}\Sigma$. It thus determines a surface T in the boundary of M', with horocyclic boundary curves $(\sigma^m)_{m\in\mathcal{P}}$. The remainder of the boundary of M' consists of euclidean cylinders of the form $\sigma^m \times [0,\infty)$. Since a geometrically finite end retracts, by nearest point retraction, onto the convex core boundary, we see that the inclusion of M' into M is a homotopy equivalence. In this case, the Thurston-Bonahon theorem, 8.2, applies and we can assume that the image of each ϕ_x lies in M'. Moreover, if $x \in T$, then we can take ϕ_x to be the incusion of T into M'. We can similarly assume that each T_x lies in M'. Thus, this time, we get sequences, ϕ_i and T_i, indexed by the natural numbers, $\mathbf{N}$, and we can take $T_0 = T$. Now the proof of Theorem 0.1 proceeds as before, replacing M by M' and the indexing set $\mathbf{Z}$ by $\mathbf{N}$.

We should comment however, that in the proof Proposition 8.7, we should continue to work with M. There is no change to the essential hypothesis (bounded local geometry). The fact that we can bound diameters of homotopies comes from Lemma 8.4. However, given that the inclusion of M' into M is a homotopy equivalence, this works equally well with M or with M'.

9. The structure of the Cannon-Thurston map.

In this section, we show that the Cannon-Thurston map is what one would expect, namely the quotient of the circle by the closed equivalence relations that arise from the ending laminations (Theorem 0.2).

Let us assume the hypotheses of Theorem 0.1. Thus, $\mathbf{H}^2/\Gamma$ is a finite-area hyperbolic surface, and we write Σ for $\mathbf{H}^2/\Gamma$ with standard open horodiscs removed. This, $\partial\Gamma$ is canonicaly identified with $\partial\mathbf{H}^2$. We have a strictly type-preserving action of $\Gamma = \pi_1(\Sigma)$

on $\mathbf{H}^3$ so that $N = \mathbf{H}^3/\Gamma$ is a hyperbolic 3-manifold with $\mathrm{inj}(N) > 0$. Let Π be the set of parabolic points, and let $\omega : S^1 \equiv \partial\mathbf{H}^2 \longrightarrow S^2 \equiv \partial\mathbf{H}^3$ be the Cannon-Thurston map (which restricts to the identity on Π). There are three cases to consider. If both ends are geometrically finite, then Γ is quasifuchsian, and ω is injective. The interesting cases are therefore the singly degenerate and doubly degenerate cases, where there are respectively one or two simply degenerate ends, each giving rise to an ending lamination as in [T1,Bon]. One characterisation of an ending lamination is described by Proposition 9.18. Lifting to $\mathbf{H}^2$, such a lamination gives rise to a closed Γ-invariant equivalence relation on $\partial\Gamma$, as will be discussed shortly.

If N is singly degenerate, we have one such relation, $\sim$, on $\partial\Gamma$. If N is doubly degenerate, we have two relations, $\sim^+$ and $\sim^-$, each arising from one of the ends. In this case, we let $\sim$ be the transitive closure of $\sim^+ \cup \sim^-$. Since transitivity is achieved in two steps (see below) it follows that $\sim$ is also a closed relation. We shall show:

Theorem 9.1 : *If $a, b \in \partial\Gamma$, then $a \sim b$ if and only if $\omega(a) = \omega(b)$.*

Note that the limit set, $\Lambda\Gamma = \omega(\partial\Gamma)$, can be identified with $\partial\Gamma/\sim$. In the singly degenerate case, this is a dendrite. If N is doubly degenerate, then $\Lambda\Gamma = S^2$, and so ω can be thought of as defining a Γ-invariant Peano curve. In this case, $\sim^+$ and $\sim^-$ are "almost independent" in the sense that if $a \sim b$ then $a \sim^+ b$ or $a \sim^- b$ or else there is some $p \in \Pi$ such that either ($a \sim^+ p$ and $b \sim^- p$) or ($a \sim^- p$ and $b \sim^+ p$). We also note that distinct elements of Π cannot be equivalent under $\sim$. (This justifies the remark about transitivity before Theorem 9.1.) We remark that under these circumstances, one can see directly that $\partial\Gamma/\sim$ has to be a topological 2-sphere via Moore's Theorem (cf. [CannD]).

To prove Theorem 9.1, we shall use again the stack, Z, of Farey graphs, and another variation on the Cannon-Thustion map, denoted τ, from $\partial\Gamma$ to ∂Z. We shall first deal with the singly degenerate case and then discuss what modifications are necessary for the doubly degenerate case. In the singly degenerate case, we shall see that two points of $\partial\Gamma$ are identified under ω if and only if they are identified under τ. We shall relate the latter condition to a relation $\approx$ arising from the stack Z. We shall see that $\approx$ is defined by some lamination L on Σ. It then remains to check that L is the same as the ending lamination.

We begin with some discussion about laminations and equivalence relations arising from them. (For more details of laminations, see for example [CasB].)

We fix, for the moment, a finite-area hyperbolic surface, $\mathbf{H}^2/\Gamma$ as above. We can suppose that any simple closed geodesic on $\mathbf{H}^2/\Gamma$ either lies in Σ or has an end that runs out a cusp. Thus, Γ acts on the circle, $S^1 \equiv \partial\Gamma \equiv \partial\mathbf{H}^2$, and hence on the open Möbius band, $\mathcal{B}$, of distinct unordered pairs of S^1. We say that two pairs, $\{x, y\}$ and $\{z, w\}$ are *linked* if $\{x, y\}$ separates z from w in S^1. They are *unlinked* otherwise (for example if they intersect). We can formally define an (abstract) lamination as a non-empty closed Γ-invariant set of pairwise unlinked pairs in $\mathcal{B}$. We shall usually work with the *realisation*, $\tilde{L}$, of such a lamination as a set of disjoint bi-infinite geodesics in $\mathbf{H}^2$. This descends to a set, L, of geodesics (bi-infinite or closed) on the surface $\mathbf{H}^2/\Gamma$, whose *support*, $\bigcup L \subseteq \mathbf{H}^2/\Gamma$, is closed. A leaf (i.e. element) of L is *isolated* if it arises from a pair in $\mathcal{B}$ that is an isolated point of the abstract lamination. A lamination is *perfect* if it has no isolated leaves. In this case, its support is compact and lies in the interior of Σ. If we remove the set of isolated leaves from a laminiation with no closed leaves, then we obtain a perfect lamination. A lamination is *complete* if each complementary region in Σ is either a topological disc or an annulus containing a boundary curve of Σ. A complete perfect lamination is minimal in the sense that it contains no proper sublamination. If we change to another structure, Σ', we get another realisation of the same abstract lamination, which is topologically conjugate: there is a homeomorphism of Σ to Σ' sending one realisation to the other. In particular, the property of completeness is well-defined.

Suppose L is a perfect lamination on Σ. We define a relation $\sim = \sim_L$ on $\partial\Gamma$ by writing $x \not\sim y$ if there is some leaf of $\tilde{L}$ whose endpoints are linked with $\{x, y\}$. It is easily verified that this is a closed Γ-invariant equivalence relation on $\partial\Gamma$. If L is complete, then one can equivalently define $\sim$ as the closed equivalence relation generated by the set of pairs of endpoints of leaves of $\tilde{L}$ (i.e. the abstract lamination). Note that, in this case, no pair of endpoints of a loxodromic can be identified under $\sim$. We first aim to characterise equivalence relations arising in this way.

Definition : Two disjoint subsets, $P, Q \subseteq S^1$ are *linked* if there exist linked pairs, $\{x, y\} \subseteq P$ and $\{z, w\} \subseteq Q$.
An equivalence relation on S^1 is *unlinked* if distinct equivalence classes are unlinked.

One verifies easily that an equivalence relation arising from a perfect lamination as above is unlinked. We have the following converse:

Lemma 9.2 : *Let $\sim$ be a non-empty closed unlinked Γ-invariant equivalence relation on S^1. Suppose that no pair of fixed points of any loxodromics are identified under $\sim$. Then there is a unique complete perfect lamination, L, on Σ such that $\sim = \sim_L$.*

In fact it is enough to consider only those loxodromics that correspond to simple closed curves on Σ.

Proof : We define an abstract lamination as follows. We deem a point, $\{x, y\}$ in $\mathcal{B}$ to belong to the lamination if $x \sim y$ and there do not exist $z, w \in S^1$ with $z \sim w$ and with $\{z, w\}$ linked with $\{x, y\}$ (so that x, y, z, w are all equivalent). It is easily seen that this is indeed an abstract lamination, and thus corresponds to a lamination, L, in Σ, with no closed leaves. (There are no leaves with endpoints in Π.) Let L_0 be the sublamination obtained by removing all isolated leaves from L. (We shall see retrospectively that $L = L_0$.) Thus, L_0 is perfect, and it is easily checked that $\sim \subseteq \sim_{L_0}$. Let $\approx$ be the closed equivalence relation generated by the pairs of endpoints of leaves of $\tilde{L}_0$. Thus, $\approx \subseteq \sim_{L_0}$. Moreover, we see easily that $\approx \subseteq \sim$.

We claim that L_0 is complete. For if not, there must be a complementary region, R, that is neither of the types permitted in a complete lamination. Such a region must contain a simple closed geodesic, β, that is either a boundary curve of R or else lies in the interior of R, and bounds an annulus in R which is bounded on the other side by a finite set of bi-infinite leaves of Σ. Let $\tilde{\beta}$ be a lift of β to $\mathbf{H}^2$. Now it is easily seen that the endpoints of $\tilde{\beta}$ are identified under $\approx$, and hence under $\sim$, contrary to our hypotheses. This proves the claim.

As observed above, since L_0 is complete and perfect, we must have $\approx = \sim_{L_0}$, and so we conclude that $\sim_{L_0} = \sim$ as required. (In retrospect, we see that $L = L_0$.)

The uniqueness of the lamination is easily verified, since any other lamination must contain a leaf that crosses every leaf of L. $\diamondsuit$

Next we introduce again a hyperbolic stack of graphs. First we recall some standard definitions. Let Z be any hyperbolic graph (not necessarily fine). We define its Gromov boundary, ∂Z, as the set of parallel classes of quasigeodesic rays, where "parallel" means at finite Hausdorff distance. (One can restrict to a set of uniform quasigeodesics, but not necessarily to geodesics.) Here we are only interested in ∂Z as a set. If $Q \subseteq Z$ is quasiconvex, we write ∂Q for the subset of ∂Z corresponding to quasigeodesics remaining a bounded distance from Q. (Note that for some uniform r, the r-neighbourhood, $N(Q, r)$, of Q is intrinsically hyperbolic, and we can identify ∂Q with the intrinsically defined boundary $\partial N(Q, r)$.)

Now consider a stratified graph, $(K_i)_{i \in \mathbf{N}}$, indexed by $\mathbf{N}$. Let $K = \bigcup_i K_i$ and let $\Pi = V(K) = V(K_i)$. Let Z be the corresponding stack, with sheets $(Z_i)_i$, so that Z_i is isomorphic to K_i. We write d_i, d_Z and d_K for the induced combinatorial metrics respectively on Z_i (or K_i), Z and K.

Now suppose that $(K_i)_i$ satisfies (S1)–(S5). Thus $\Delta^0 K$ is identified with $V(Z_i) \cup \partial Z_i$ for each i, and we have the Cannon-Thurston map $\omega : \Delta^0 K \longrightarrow \Delta K$. Given $x, y \in \Delta^0 K$, write $[x, y]_i$ for some choice of geodesic in Z_i connecting x to y. (Any two such geodesics are a uniformly bounded d_i-distance apart.) It will be convenient to assume that $[x, y]_i = [y, x]_i$ and that $[z, w]_i \subseteq [x, y]_i$ whenever $z, w \in [x, y]_i$. (One cannot necessarily make such choices globally, but there will be no problem in the situations that arise here.) We shall write $\Upsilon(x, y) = \bigcup_i [x, y]_i \subseteq Z$.

By an *r-centre* of three points, $x, y, z \in \Delta^0 K$, at level i, we mean a vertex of Z_i within a distance r of each of the geodesics $[x, y]_i$, $[y, z]_i$ and $[z, x]_i$. Any two such centres are a bounded d_i-distance apart. We can uniformly choose r such that an r-centre always exists on $[x, y]_i$. We shall write $m_i(x, y, z) \in [x, y]_i$ for some choice of r-centre.

Again, we fix $r \in \mathbf{N}$. By an *r-chain* in Z, we mean a sequence, $\underline{x} = (x_i)_i$, of vertices of Z such that $x_i \in V(Z_i)$ and $d_Z(x_i, x_{i+1}) \le r$ for all $i \in \mathbf{N}$. We refer to r as the *chain constant*. Thus, a chain is quasigeodesic with parameters depending only on r. It thus determines a point $x_\infty \in \partial Z$. Note that a 1-chain is just the vertex set of a vertical line, $l(x)$, for some $x \in \Pi$. We refer to this as the *constant* chain at x, and denote its endpoint by $\tau(x) \in \partial Z$. We shall write $\partial^+_r Z$ for the set of points x_∞ as $\underline{x}$ varies over all r-chains in Z.

Given $a, b, c \in \Delta^0 K$, we write $\underline{m}(a, b, c) = (m_i(a, b, c))_i$. Thus, $\underline{m}(a, b, c)$ is a uniform chain, i.e. with chain constant, r_0, say, depending only on the parameters. It determines a point $m_\infty(a, b, c) \in \partial Z$. We can choose r_0 so that $\underline{m}(a, b, c)$ lies in $\Upsilon(a, b)$, and such that every point of $\Upsilon(a, b)$ for every $a, b \in \Delta^0 K$ lies in an r_0-chain. We shall eventually fix such an r_0, and refer to r_0-chains simply as "chains". We shall abbreviate $\partial^+ Z = \partial^+_{r_0} Z$.

Given $a, b \in \Delta^0 K$, not both equal to the same ideal point, it follows exactly as in Lemma 4.5 that $\Upsilon(a, b)$ is a uniformly quasiconvex subset of Z. It thus determines a subset $I(a,b) = \partial \Upsilon(a,b)$ of ∂Z. Note that $\Upsilon(a,b)$ depends on the choice of geodesics $[a,b]_i$. However any two such sets arising from different choices are at finite Hausdorff distance and hence give rise to the same subset, $I(a,b)$. Note that if $a \in \Pi$, then $\tau(a) \in I(a,b)$. Also, if $a,b,c \in \Delta^0 K$, then $\Upsilon(a,b)$ lies in a bounded neighbourhood of $\Upsilon(a,c) \cup \Upsilon(c,b)$, and so $I(a,b) \subseteq I(a,c) \cup I(c,b)$.

Given $r \geq 0$, we can embed $\Upsilon(a,b)$ in a graph $\Upsilon_r(a,b)$ by connecting $x \in [a,b]_i$ to $y \in [a,b]_{i+1}$ with an edge whenever $d_Z(x,y) \leq r$. If $r \geq r_0$, then the inclusion of $\Upsilon(a,b)$ with the metric d_Z into $\Upsilon_r(a,b)$ with the intrinsic path-metric is a quasi-isometry. We see that $\Upsilon_r(a,b)$ is intrinsically hyperbolic with $I(a,b)$ identified with $\partial \Upsilon_r(a,b)$ and we can topologise $I(a,b)$ accordingly. Note that $\Upsilon_r(a,b)$ is locally finite and one-ended, and so $I(a,b)$ is a continuum. Indeed, it's not hard to see that $I(a,b)$ is a (possibly degenerate) interval. Moreover, if $a \in \Pi$, then $\tau(a)$ is an endpoint of this interval. (See [Bow3] for details.)

We shall define a relation, $\approx$, on $\Delta^0 K$ by writing $a \approx b$ if $a = b$ or if $I(a,b)$ consists of a single point. From the above observations, we see that $\approx$ is an equivalence relation. The following gives a criterion for recognising this equivalence (see [Bow3]).

Lemma 9.3 :
(1) If $a,b \in \Pi$, then $a \approx b$ if and only if $a = b$.
(2) If $a,b \in \Delta^0 K \setminus \Pi$, and there exist sequences of uniform chains $(\underline{x}^n)_{n \in \mathbf{N}}$ and $(\underline{y}^n)_{n \in \mathbf{N}}$ in $\Upsilon(a,b)$ with $x_0^n \to a$ and $y_0^n \to b$ in $[a,b]_0$, and with $\min\{d_i(x_i^n, y_i^n) \mid i \in \mathbf{N}\}$ bounded above (independently of n), then $a \approx b$.
(3) If $a \in \Pi$, $b \in \Delta^0 K \setminus \Pi$ and there is a sequence of uniform chains $(\underline{y}^n)_{n \in \mathbf{N}}$ in $\Upsilon(a,b)$ with $y_0^n \to b$ in $[a,b]_0$ and with $\min\{d_i(a, y_i^n) \mid i \in \mathbf{N}\}$ bounded above, then $a \approx b$. $\diamondsuit$

Next, we restrict to the case where each K_i is a isomorphic to the Farey graph, A, and that $(K_i)_i$ satisfies (S1)–(S6). In this case, $\Delta^0 K \cong S^1$.

Lemma 9.4 : *If $\{x,y\}$ and $\{z,w\}$ are disjoint and linked in $\Delta^0 K$, then $I(x,y) \cap I(z,w) \neq \emptyset$.*

Proof : The chain $\underline{m}(x,y,z)$ lies a bounded distance from both $\Upsilon(x,y)$ and $\Upsilon(z,w)$, and so $m_\infty(x,y,z) \in I(x,y) \cap I(z,w)$. $\diamondsuit$

Corollary 9.5 : *The relation $\approx$ is unlinked.*

Proof : If $x \approx y$ and $z \approx w$ and $\{x,y\}$ is linked with $\{z,w\}$, then by Lemma 9.4, $I(x,y) = I(z,w)$ and so $x \approx y \approx z \approx w$. $\diamondsuit$

We next want to prove that the relation $\approx$ is closed. This can be bypassed in the singly degenerate case as it follows from Proposition 9.16. However it will be needed to deal with the doubly degenerate case.

First, we need the following observation concerning the Farey graph, A. If $a \in V(A)$ and $x,y \in \Delta A \setminus \{a\}$ are "far enough apart" then any geodesic from x to y in A must pass through a. To be far enough apart, it is sufficient that at least three distinct points of

$V_a(A)$ separate x from y in $\Delta A \setminus \{a\}$, where $V_a(A) \subseteq V(A)$ denotes the set of vertices adjacent to a.

Now applying property (S6) and the above observation, we deduce:

Lemma 9.6 : *There is some $\theta_0 \in \mathbf{N}$ such that if $a \in \Pi$, $x, y \in \Delta^0 K$ and at least θ_0 elements of $V_a(K_0) \subseteq \Pi$ separate x from y in $\Delta^0 K \setminus \{a\}$, then for all $i \in \mathbf{N}$, $a \in [x,y]_i$.* $\diamondsuit$

Lemma 9.7 : *The relation $\approx$ on $\Delta^0 K$ is closed.*

Proof : Suppose $a, b \in \Delta^0 K$ are distinct, and $a^n \to a$ and $b^n \to b$ in $\Delta^0 K$ with $a^n \approx b^n$ for all n. We want to prove that $a \approx b$. There are three cases.

Case (1) $a, b \notin \Pi$.

Let $\underline{x}^n = \underline{m}(a, b, a^n)$ and $\underline{y}^n = \underline{m}(a, b, b^n)$. Thus, $\underline{x}^n$ and $\underline{y}^n$ are chains in $\Upsilon(a,b)$ with $x_0^n \to a$ and $y_0^n \to b$ in $[a,b]_0$. Also $x_\infty^n, y_\infty^n \in I(a^n, b^n)$, so since $a^n \approx b^n$, we have $x_\infty^n = y_\infty^n$. In particular, $\min\{d_i(x_i^n, y_i^n) \mid i \in \mathbf{N}\}$ is (uniformly) bounded. By Lemma 9.3(1) we see that $a \approx b$ as required.

Case (2) $a \in \Pi$, $b \notin \Pi$.

Since $a^n \to a$, for all sufficiently large n, a^n and b are separated in $\Delta^0 K \setminus \{a\}$ by at least $\theta_0 + 1$ points of $V_a(K_0)$. Since $b^n \to b$, the same goes for a^n and b^n. It follows by Lemma 9.6 that $a \in [a^n, b^n]_i$ for all i. We now apply the argument of Case (1), with $\underline{y}^n$ as before, and with each $\underline{x}^n$ replaced by the constant chain, $\underline{a}$, based at a. Since $a^n \approx b^n$ for all n, the chains $\underline{y}^n$ and $\underline{a}$ have the same endpoints in ∂Z. Thus, applying Lemma 9.3(2), we see that $a \approx b$.

Case (3) $a, b \in \Pi$.

We argue as in Case (2), this time also replacing $\underline{y}^n$ by the constant chain $\underline{b}$. We again deduce that $a \approx b$ (so that, in fact, $a = b$). $\diamondsuit$

Now let us suppose that we are in the situation described at the end of Section 5. Thus, we have a closed surface, Θ, and a finite set $P \subseteq \Theta$. We write Σ for the surface obtained by removing small open discs around each of the points of P. Let $\Gamma = \pi_1(\Sigma)$. We can also view Σ as a complete finite-area hyperbolic surface, $\mathbf{H}^2/\Gamma$, with a set of horodiscs removed. Any triangulation of Θ with vertex set P gives rise to an action of Γ on the Farey graph, A, with $P = \Pi/\Gamma$ where $\Pi = V(A)$. Thus, a sequence of such triangulations gives us a stratified graph, $(K_i)_i$. We assume this satisfies (S1), (S2) and (S4), and hence all of (S1)–(S6). We have actions of Γ on K and on Z. We note:

Lemma 9.8 : *If $a, b \in \Delta^0 K$ are the endpoints of a loxodromic, then $a \not\approx b$.*

Before proving this, we need an observation, that will be used again later. Let A be a Farey graph. Suppose we fix an orientation on $S^1 \equiv \Delta A$. Given an ordered pair of distinct points, a and b, in ΔA we can label the complementary intervals, J_L and J_R as "left" and "right" respectively. One can show that there is a unique *rightmost* geodesic, α, from a to b in A. This is characterised by the property that if β is any other geodesic connecting a to b, then $\beta \cap J_R \subseteq \alpha \cap J_R$. If $g \in \Gamma$ is an orientation-preserving loxodromic, then the rightmost geodesic connecting the repelling fixed point to the attracting fixed point is $\langle g\rangle$-invariant (the *rightmost axis* of g). Let $D(g) \in \mathbf{N}$ be the distance g translates the axis α. One shows that $D(g) = \min\{d_A(x, gx) \mid x \in V(A)\}$. Note also that $D(g^n) = nD(g)$. In particular, $D(g^n) \geq n$. Applying this to our stack, we write $D_i(g)$ for the translation distance of g in K_i.

Proof of Lemma 9.8 : Suppose a and b are the fixed points of a loxodromic $g \in \Gamma$. Replacing g by g^2 if necessary, we can suppose that g is orientation-preserving. Construct $\Upsilon(a, b)$ by taking $[a, b]_i$ to be the rightmost axis of g in Z_i. Thus $\Upsilon(a, b)$ is $\langle g\rangle$-invariant. Let $\underline{x}$ be any chain in $\Upsilon(a, b)$. If $a \approx b$ then for all n, $g^n x_\infty = x_\infty$. In other words, the chains $\underline{x}$ and $g^n\underline{x}$ are parallel. It follows that $\min\{d_i(x_i, g^n x_i) \mid i \in \mathbf{N}\}$ is bounded independently of n. But $d_i(x_i, g^n x_i) \geq D_i(g^n) \geq n$ giving a contradiction. $\diamondsuit$

Lemma 9.9 : *There is a unique complete perfect lamination, L, on Σ such that $\approx = \sim_L$.*

Proof : The relation $\approx$ is closed by Lemma 9.6 and unlinked by Corollary 9.5. Combining this with Lemma 9.8, we see that the hypotheses of Lemma 9.2 are satisfied, and the result follows. $\diamondsuit$

We shall later need the following further observation regarding Farey graphs. We suppose that the Farey graph A arises from the triangulation of the surface (Θ, P) as above, so that $P = V(A)/\Gamma$.

Lemma 9.10 : *There is a constant, $\theta_1 \in \mathbf{N}$, depending only on the topological type of (Θ, P) such that the following holds. Suppose that $a, b \in \Delta A$ are such that $\{ga, gb\}$ is unlinked with $\{a, b\}$ for all (non-trivial) $g \in \Gamma$. If α is any geodesic in A connecting a and b, and x is an interior vertex of α, then $\angle_x(\alpha) \leq \theta_1$.*

Proof : Write $\theta = \angle_x(\alpha)$. The stabiliser of x in Γ is infinite cyclic, generated by some $g \in \Gamma$. The set, $V_x(A)$, of adjacent vertices is isometric to $\mathbf{Z}$ in the metric $d_{A\setminus\{x\}}$, and is translated some distance, say ϕ, by g. Note that ϕ is bounded above by the number of edges of the triangulation of Θ, and hence in terms of the topological type of (Θ, P).

Let $y, z \in V_x(A)$ be the vertices of α adjacent to x. If $\theta \geq \phi + 2$, then, modulo interchanging y and z, the vertices, y, gy, z, gz occur on this order along $V_x(A) \cong \mathbf{Z}$, and $d_{A\setminus\{a\}}(z, gy) \geq 2$. From the geometry of the Farey graph, it is easily verified that pairs $\{a, b\}$ and $\{ga, gb\}$ must be linked in ΔA. This contradiction shows that $\theta \leq \phi + 1$, which is bounded in terms of the type of (Θ, P) as required. $\diamondsuit$

Note that this applies to the leaves of any lamination on Σ, realised as geodesics in A.

The next step is to associate to any $a \in \Delta^0 K$, a point $\tau(a) \in \partial Z$. If $a \in \Pi$, we have already defined $\tau(a)$ as the endpoint of the constant chain, $\underline{a}$, at a. The restriction $\tau : \Delta^0 K \setminus \Pi \longrightarrow \partial Z$ can be thought of a special case of the Cannon-Thurston map for stacks of spaces (as discussed in [Bow3]). Here we give only a particular case of this construction. A similar argument can be found in [Mit1,Mit2].

Suppose $I \subseteq \Delta^0 K$ is an interval. Now $I \cap \Pi$ is uniformly quasiconvex in each graph K_i, or equivalently, Z_i. Let $Q_i(I) \subseteq V(Z_i)$ be the set of vertices projecting to $I \cap \Pi$, and let $Q(I) = \bigcup_i Q_i(I)$. In other words, $Q(I) = V(Z) \cap \mathrm{proj}^{-1}(I \cap \Pi)$, where $\mathrm{proj} : Z \longrightarrow K$ is the projection map. As in Lemma 4.5, we see that $Q(I)$ is a quasiconvex subset of Z.

By a *standard base* for a point $a \in \Delta^0 K$, we mean a decreasing sequence of intervals, $(I^n)_{n \in \mathbf{N}}$, containing a in their interiors and with $\bigcap_n I^n = \{a\}$. If $a \notin \Pi$, then the sets $Q(I^n)$ are escaping in Z, in the usual sense that $d_Z(x, Q(I^n)) \to \infty$ for some (hence all) $x \in Z$. Since they are uniformly quasiconvex, it follows that $\bigcap_n \partial Q(I^n)$ consists of a single point of ∂Z. This point is easily seen to be independent of the choice of standard base, I^n. We denote it by $\tau(a)$.

Lemma 9.11 : *If $a, b \in \Delta^0 K$, then $a \approx b$ if and only if $\tau(a) = \tau(b)$.*

Proof : One way to see this is to use the fact that $\tau(a)$ and $\tau(b)$ are precisely the endpoints of the interval $I(a, b)$ (see [Bow3]).

A more direct way to see this, which will be relevant to the doubly degenerate case, is as follows. We deal with the case where $a, b \notin \Pi$. The other case where $a \in \Pi$ and $b \notin \Pi$ is similar.

We first show that $a \approx b$ implies $\tau(a) = \tau(b)$. To this end, let $(I^n)_n$ and $(J^n)_n$ be standard bases of a and b. Now $\Upsilon(a, b)$ meets each $Q(I^n)$ and each $Q(J^n)$ in an unbounded set. Thus $I(a, b) = \partial \Upsilon(a, b)$ meets both $\partial Q(I^n)$ and $\partial Q(J^n)$, and hence their respective intersections over all n, namely $\{\tau(a)\}$ and $\{\tau(b)\}$. Since $a \approx b$, it follows that $I(a, b) = \{\tau(a)\} = \{\tau(b)\}$ as required.

For the converse, suppose $\tau(a) = \tau(b)$. Now $\partial Q(I^n) \cap \partial Q(J^n) \neq \emptyset$ for all n, and so $d_Z(Q(I^n), Q(J^n))$ is uniformly bounded, and so therefore is $\min\{d_i(Q_i(I^n), Q_i(J^n) \mid i \in \mathbf{N}\}$. Choose $p^n \in \Pi \cap I^n$ and $q^n \in \Pi \cap J^n$, so that $d_{i(n)}(p^n, q^n)$ is uniformly bounded for suitable $i(n)$. Let $\underline{x}^n = \underline{m}(a, b, p^n)$ and $\underline{y}^n = \underline{m}(a, b, q^n)$. It follows that $d_{i(n)}(x_i^n, y_i^n)$ is uniformly bounded. Moreover, since $Q_0(I^n)$ and $Q_0(J^n)$ escape respectively to a and b, we have $x_0^n \to a$ and $y_0^n \to b$ in $[a, b]_0$. Since $\underline{x}^n$ and $\underline{y}^n$ are uniform chains in $\Upsilon(a, b)$, by Lemma 9.3, it follows that $a \approx b$. ◊

We have defined $\tau : \Delta^0 K \longrightarrow \partial Z$ and $\omega : \Delta^0 K \longrightarrow \Delta K$. The next step will be to show that $\tau(a) = \tau(b)$ if and only if $\omega(a) = \omega(b)$. First we deal with the easy direction:

Lemma 9.12 : *If $a, b \in \Delta^0 K$ and $\tau(a) = \tau(b)$ then $\omega(a) = \omega(b)$.*

Proof : First consider the case where $a, b \notin \Pi$. Let $(I^n)_n$ and $(J^n)_n$ be standard bases for a and b respectively in $\Delta^0 K$. Since $\tau(a) = \tau(b)$, it follows that for all n, $\partial Q(I^n) \cap \partial Q(J^n) \neq \emptyset$ and so $d_Z(Q(I^n), Q(J^n))$ is uniformly bounded. Let α^n be a geodesic of bounded length connecting $Q(I^n)$ to $Q(J^n)$ in Z, and let β^n be its projection to K. Now α^n is escaping in Z, and so the β^n can be assumed to be edge-disjoint. Note that β^n is an arc of bounded length connecting $\Pi \cap I^n$ to $\Pi \cap J^n$ in K. By the definition of ω, $\bigcap_n \overline{\Pi \cap I^n} = \{\omega(a)\}$ and $\bigcap_n \overline{\Pi \cap J^n} = \{\omega(b)\}$. Applying Corollary 3.11 it follows that $\omega(a) = \omega(b)$.

The case where $a \in \Pi$ and $b \notin \Pi$ is similar, replacing each $Q(I^n)$ by the constant chain at a. In this way, we get arcs β^n of bounded length connecting a to $\Pi \cap J^n$. Thus, by Lemma 3.5, $\omega(a) = a \in \bigcap_n \overline{\Pi \cap J^n} = \{\omega(b)\}$ and so $\omega(a) = \omega(b)$. ◊

For the converse, we recall the lifting construction described in Section 4. If β is an arc in K, we denote by $\mathrm{lift}\, \alpha$ an arc β in Z such that $\alpha = \mathrm{proj}\, \beta$, and which meets each of the vertical lines $l(a)$ and $l(b)$ in a single point. Thus β consists of a set of horizontal edges alternating with vertical segments, i.e. subarcs of $l(c)$ for each of the interior vertices c of β. If α is a geodesic, then Lemma 4.3 tells us that β is uniformly quasigeodesic in Z. Indeed the same result applies to any arc of the form, $\gamma \cup \beta \cup \delta$, where γ and δ are subarcs of $l(a)$ and $l(b)$ respectively attached to the endpoints of β. From this (applying the hyperbolicity of Z) it follows that β lies a bounded Hausdorff distance from any shortest path connecting $l(a)$ to $l(b)$ in Z.

The proof that $\omega(a) = \omega(b)$ implies $\tau(a) = \tau(b)$ will be in two steps. First we show:

Lemma 9.13 : *If $a, b \in \Delta^0 K$ and $\omega(a) = \omega(b) \notin \Pi$ then $\tau(a) = \tau(b)$.*

Proof : Note that for $c \in \Pi$, $\omega(c) = c \in \Pi$, so we must have $a, b \notin \Pi$.

Suppose for contradiction that $\tau(a) = \tau(b)$. Let γ be a bi-infinite quasigeodesic in Z with endpoints $\tau(a), \tau(b) \in \partial Z$. Let $(I^n)_n$ and $(J^n)_n$ be standard bases for a and b respectively, and let δ^n be a shortest path in Z from $Q(I^n)$ to $Q(J^n)$. Since $\tau(a) \neq \tau(b)$ it follows that $d_Z(Q(I^n), Q(J^n)) \to \infty$, and that the paths δ^n "converge toward" γ. More precisely, on passing to a subsequence there are increasing compact subarcs, γ^n, of γ such that $\gamma = \bigcup_n \gamma^n$, and such that the Hausdorff distance between γ^n and δ^n is uniformly bounded.

By definition, $\bigcap_n \overline{\Pi \cap I^n} = \{\omega(a)\}$. Since $\omega(a) \in \Delta K \setminus \Pi = \partial K$, we see that $\Pi \cap I^n$ is escaping. (See the remark after Lemma 3.6.) The same goes for $\Pi \cap J^n$. Using Corollary 3.4, we see that any quasigeodesic ray in K converging on $\omega(a) = \omega(b) \in \partial K$ remains a uniformly bounded distance from $\Pi \cap I^n$ and $\Pi \cap J^n$. In particular, $d_K(\Pi \cap I^n, \Pi \cap J^n)$ is unformly bounded. We can thus find geodesics α^n of bounded length connecting points of $\Pi \cap I^n$ and $\Pi \cap J^n$. The arcs α^n are escaping in K and can thus be assumed disjoint. Let $\beta^n = \mathrm{lift}\, \alpha^n$. Now each β^n is a uniform quasigeodesic in Z connecting $Q(I^n)$ to $Q(J^n)$. It follows that δ^n remains a uniform bounded distance from β^n, at least for all sufficiently large n. Since δ^n converges towards γ, we see that for all $r \geq 0$, there is some $n(r)$ such that if $m, n \geq n(r)$, then β^m and β^n remain uniformly close over a distance at least r. (More precisely they contain subsegments of length r at a uniformly bounded Hausdorff distance from each other.)

But now each β^n consists of a bounded number of vertical segments connected by horizontal edges. It follows that we can find $m \neq n$ such that β^m and β^n contain vertical segments that remain uniformly close over an arbitrarily large distance, s, say. From the uniform divergence of distinct vertical lines (see Lemma 4.2), it follows that, provided we choose s large enough, these segments must lie in the same vertical line, $l(c)$, say. But now $c \in \alpha^m \cap \alpha^n$ contradicting the fact that the α^n are all disjoint. $\diamondsuit$

Lemma 9.14 : *Suppose $p \in \Pi$ and $a \in \Delta^0 K$. If $\omega(a) = p$ then $\tau(a) = \tau(p)$.* $\diamondsuit$

Proof : If $a \in \Pi$, then $a = p$, so we may suppose that $a \notin \Pi$. Let $(I^n)_n$ be a standard base for a. By definition, $\tau(p)$ is the endpoint of the vertical line $l(p)$. Thus if $\tau(a) \neq \tau(p)$, we have $d_Z(l(p), Q(I^n)) \longrightarrow \infty$. Let δ^n be a shortest path connecting $l(p)$ to $Q(I^n)$. Thus δ^n

converges towards a quasigeodesic ray γ as in the proof of Lemma 9.13, where γ emerges from $l(p)$ and has endpoint $\tau(a)$.

Now since $\omega(a) = p$, we have $\bigcap_n \overline{\Pi \cap I^n} = \{p\}$. Thus we can find a sequence of geodesics, α^n, of bounded length in K connecting p to $a^n \in \Pi \cap I^n$, with $\alpha^n \cap \alpha^m = \{p\}$ whenever $m \neq n$. Let $\beta^n = \operatorname{lift} \alpha^n$. Now β^n is a uniform quasigeodesic, which remains a uniformly bounded distance from any shortest path from $l(p)$ to $l(a^n)$. Since each vertex of $l(a^n)$ lies in $Q(I^n)$, we see that δ^n lies within a bounded distance from β^n, for all sufficiently large n. Again β^n consists of a bounded number of vertical segments and horizontal edges. We thus argue exactly as in Lemma 9.13 to find a vertex $c \neq p$ contained in $\alpha^m \cap \alpha^n$. This contradicts the assumption that $\tau(a) \neq \tau(p)$. $\diamondsuit$

Lemma 9.15 : *If $a, b \in \Delta^0 K$ and $\omega(a) = \omega(b)$ then $\tau(a) = \tau(b)$.*

Proof : If $a, b \in \Pi$ then $a = b$.

If $a \in \Pi$ and $b \notin \Pi$, then $\omega(a) = a = \omega(b)$, and Lemma 9.14 applies.

Finally suppose $a, b \notin \Pi$. If $\omega(a) \in \Pi$, then $\omega(\omega(a)) = \omega(a) = \omega(\omega(b))$, and so applying Lemma 9.14 twice, we get $\tau(a) = \tau(\omega(a)) = \tau(b)$. If $\omega(a) \notin \Pi$, then Lemma 9.13 gives $\tau(a) = \tau(b)$. $\diamondsuit$

Proposition 9.16 : *If $a, b \in \Delta^0 K$, then $a \approx b$ if and only if $\omega(a) = \omega(b)$.*

Proof : Combine Lemma 9.11, Lemma 9.12 and Lemma 9.15. $\diamondsuit$

Now returning to the set-up of Lemma 9.9, where the stratified graph arises from a sequence of triangulated surfaces, we see that $\approx$ is associated to a unique lamination. In summary, we have shown:

Proposition 9.17 : *With the hypotheses introduced before Lemma 9.8, and with $\omega : \Delta^0 K \longrightarrow \Delta K$ defined as in Proposition 5.5, there is a unique complete perfect lamination L on the surface Σ with the property that if $a, b \in \Delta^0 K$, then $\omega(a) = \omega(b)$ if and only if a and b are not separated by any leaf of the lift, $\tilde{L}$, of L to $\mathbf{H}^2$, where we have equivariantly identified $\partial \mathbf{H}^2$ with $\Delta^0 K$.* $\diamondsuit$

It remains to show that if our stratified graph arises from a simply degenerate end as described in Section 8, then L is indeed the ending lamination. To this end, we use the following characterisation of an ending lamination.

Suppose we have a strongly type-preserving action of our surface group, Γ, on $\mathbf{H}^3$, and write $N = \mathbf{H}^3/\Gamma$ and M for the submanifold obtained by removing horoballs. We have a relative homotopy equivalence, ϕ, of our surface, Σ, into M. We fix a hyperbolic hyperbolic structure on Σ, so that each boundary curve is a horocycle of fixed length.

Proposition 9.18 : *Suppose that $(\alpha_i)_i$ is a sequence of simple closed geodesics in Σ and that $(\beta_i)_i$ is a sequence of closed curves of bounded length in M, which go out the end e, and with β_i homotopic in M to $\phi \circ \alpha_i$. Then α_i converges to the ending lamination, $L(e)$, of the end e.* $\diamondsuit$

(Note that we can always find such curves α_i. In fact, for our purposes it would be sufficient to find a convergent subsequence.)

Here, we need only consider the case where $\mathrm{inj}(N) > 0$. To relate this formulation to that given in [Bon], note that in this case each of the curves β_i remains a bounded distance from the closed geodesic in M in the same homotopy class, and so these closed geodesics also go out the end.

Now an ending lamination, such as $L(e)$, is necessarily complete and perfect, and hence minimal in the sense that it contains no proper sublamination. In particular, every leaf is dense. If L is another lamination, then either $L = L(e)$, or else L contains a leaf which crosses some, hence every, leaf of $L(e)$. Applying this to the lamination L, given by Proposition 9.18, we see that in order to show that $L = L(e)$, it suffices to find some leaf of $\tilde{L}(e)$ whose endpoints are identified under the relation $\approx$.

We recall the construction of $(K_i)_i$ given in Section 8. We have a sequence, $\phi_i : \Sigma \longrightarrow M$ of relative homotopy equivalences which are also uniformly lipschitz maps. We can assume that ϕ_0 lies in the same homotopy class as the map ϕ defining the ending lamination on Σ. Collapsing each boundary curve to a point we obtain a surface Θ with finite subset $P \subseteq \Theta$. The lift of the 1-skeleton gives us a Farey graph, K_0, with vertex set Π admitting a Γ-action with $P = \Pi/\Gamma$. Pushing forward under the mapping classes, $\psi_i = [\phi_0]^{-1}[\phi_i]$ and lifting gives us a stratified graph $(K_i)_i$.

Let α_0 be any simple closed geodesic on Σ. Let α_i be the closed geodesic homotopic to $\psi_i\alpha_0$. Let $\beta_i = \phi_i\alpha_0$ so that β_i is homotopy equivalent in M to $\phi_0\alpha_i$. Since the maps ϕ_i are uniformly lipschitz, the β_i have bounded length. Thus, by the above characterisation of ending laminations (Proposition 9.18), α_i converges on $L(e)$. In other words, we can find lifts, $\tilde{\alpha}_i$, of α_i, which converge on a leaf, λ, of $\tilde{L}(e)$.

We need to show:

Lemma 9.19 : *The endpoints of λ are identified under the relation $\approx$.*

Proof : Let $g_i \in \Gamma$ be the element corresponding to the closed curve α_i, and preserving the axis $\tilde{\alpha}_i$. Let a_i and b_i be the fixed points of g_i. Thus $a_i \longrightarrow a$ and $b_i \longrightarrow b$ in $\partial \mathbf{H}^2 \equiv \Delta^0 K$, where a and b are the endpoints of the leaf λ. Let γ_i^n be the rightmost axis of g_n in the Farey graph Z_i (i.e, the rightmost geodesic in Z_i with endpoints a_n and b_n). Thus, γ_i^n is $\langle g_n \rangle$-invariant. (We can assume that g_n is orientation-preserving by taking α_0 and hence each α_n to preserve orientation in Σ.) Now by construction, the actions of $\langle g_n \rangle$ on Z_n are all conjugate. (This is because g_n and K_n arise by pushing forward g_0 and K_0 under a lift of ψ_n.) In particular, the translation distance, $D_n(g_n) = D_0(g_0)$ is independent of n. We can thus find some $x^n \in \Pi$ such that $d_n(x^n, g_n x^n) = D_0(g_0)$ for all n.

On the other hand, we claim that $D_0(g_n) \to \infty$ as $n \to \infty$. To see this, note that by Lemma 9.10, $\angle_x(\gamma_0^n)$ is uniformly bounded for all n and for all $x \in \gamma_0^n$. Thus, up to the action of Γ, there are only finitely many possibilities for a subarc of γ_0^n of any given length. Thus, if $D_0(g_n)$ were bounded, there would be only finitely many possibilities for g_n up to conjugacy in Γ. But g_n corresponds to the geodesics, α_n, whose lengths in Σ tend to infinity. This contradiction shows that $D_0(g_n) \to \infty$ as claimed.

Now let $\delta_i = [a, b]_i$ be any geodesic joining a to b in Z_i. We fix any point $c \in \delta_0 \cap \Pi$.

Let $y^n = m_0(a_n, b_n, x^n) \in \gamma_0^n$.

Now, $a_n \to a$ and $b_n \to b$, and so for sufficiently large n, the geodesic γ_0^n lies a uniformly bounded d_0-distance from $\delta_0 = [a,b]_0$ for an arbitrarily large distance either side of c. (More precisely, there is an arbitrarily long subarc of δ_0, centred on c, that is a bounded distance from a subarc of γ_0^n.) Note that it makes sense to speak of two points of γ_0^n a long way from c as being on "opposite sides" of c.

Now, g_n translates the axis γ_0^n an arbitrarily large distance $D_0(g_n)$. By considering the $\langle g_n \rangle$-orbit of y^n on γ_0^n, we can suppose (after replacing x^n and hence y^n by a suitable $\langle g_n \rangle$-image if necessary) that y^n and $g_n^2 y^n$ lie a long way on opposite sides of c. Thus if $z_0^n = m_0(a,b,x^n) \in \delta_0$, and $w_0^n = m_0(a,b,g_n^2 x^n) \in \delta_0$, then a simple exercise using the hyperbolicity of Z_0 shows that z_0^n and w_0^n lie an arbitrarily large distance on opposite sides of c. Thus (without loss of generality) $z_0^n \to a$ and $w_0^n \to b$.

Now let $\underline{z}^n = \underline{m}(a,b,x^n)$ and $\underline{w}^n = \underline{m}(a,b,g_n^2 x^n)$. Thus $\underline{z}^n$ and $\underline{w}^n$ are chains lying in $\Upsilon(a,b) = \bigcup_i \delta_i$. Since $d_n(x^n, g_n^2 x^n) = 2D_n(g_n) = 2D_0(g_0)$ is constant, we see that $d_n(z_n^n, w_n^n)$ is bounded. Thus by Lemma 9.3, we deduce $a \approx b$. $\diamondsuit$

In summary, we have found a leaf of the lift of the ending lamination whose endpoints are identified under $\approx$. As noted earlier, this is sufficient to prove that $L = L(e)$. Thus, $\approx$ is precisely the equivalence relation defined by the ending lamination as described before the statement of Theorem 9.1. We have thus proven Theorem 9.1 in the singly degenerate case. More precisely:

Proposition 9.20 : *Suppose N is singly degenerate with* $\mathrm{inj}(N) > 0$ *and with ending lamination, L. If $a, b \in \partial \mathbf{H}^2$ then $\omega(a) = \omega(b)$ if and only if $a \sim_L b$.* $\diamondsuit$

We next have to consider the doubly degenerate case.

Let $(K_i)_{i \in \mathbf{Z}}$ be a stratified graph, this time indexed by $\mathbf{Z}$, satisfying (S1)–(S5), and let Z be the corresponding stack. We note that the full subgraphs on $\bigcup_{i \in \mathbf{N}} Z_i$ and $\bigcup_{i \in -\mathbf{N}} Z_i$ can be viewed as stacks corresponding to $(K_i)_{i \in \mathbf{N}}$ and $(K_i)_{i \in -\mathbf{N}}$ respectively. Passing back and forth via the Bestvina-Feighn flaring condition, we see that they are intrinsically hyperbolic, although not usually quasiconvex in $\mathbf{Z}$. We shall not make explicit use of this fact here, though it sets some of the arguments in context.

We shall define *positive*, *negative* and *bi-infinite r-chains* as sequences $\underline{x} = (x_i)_{i \in I}$ indexed respectively by $I = \mathbf{N}$, $I = -\mathbf{N}$ and $I = \mathbf{Z}$, where $x_i \in V(Z_i)$ and $d_Z(x_i, x_{i+1}) \le r$ for all i. We shall fix a suitable r_0 as before and refer to an r_0-chain simply as a *chain*. We write x_∞ and $x_{-\infty}$ in $\mathbf{Z}$ for the positive and negative endpoints of the chain $\underline{x}$. In the case of the constant chain at x, we shall write $\tau^+(x) = x_\infty$ and $\tau^-(x) = x_{-\infty}$. Given $r \in \mathbf{N}$, write $\partial_r^+ Z$ for the set of $x_\infty \in \partial Z$, as $\underline{x}$ varies over all positive r-chains. We similarly define $\partial_r^- Z$. It is easily seen that $\partial_r^\infty Z \cap \partial_r^- Z = \emptyset$.

Given $a, b \in \Delta^0 K$, write $\Upsilon(a,b) = \bigcup_{i \in \mathbf{Z}} [a,b]_i$. Given $r \in \mathbf{N}$, write $I_r^+(a,b)$ (respectively $I_r^-(a,b)$) for the endpoints of positive (respectively negative) r-chains in $\Upsilon(a,b)$. We write $I^\pm(a,b) = \bigcup_{r \in \mathbf{N}} I_r^\pm(a,b)$. This is always non-empty, and independent of the choice of geodesics $[a,b]_i$ defining $\Upsilon(a,b)$. If $a, b, c \in \Delta^0 K$, then $I^\pm(a,b) \subseteq I^\pm(a,c) \cup I^\pm(c,b)$.

We define an equivalence relation $\approx^\pm$ on $\Delta^0 K$ by writing $a \approx^\pm b$ if $|I^\pm(a,b)| = 1$. In fact, this is equivalent to asserting that $|I_r^\pm(a,b)| = 1$ for a certain $r \in \mathbf{N}$ sufficiently large in relation to the parameters. Moreover, one has a criterion for recognising this similar to that of Lemma 9.3, namely:

Lemma 9.21 : *Given $a, b \in \Delta^0 K$ one has the same relation as in Lemma 9.3, with $a \approx^+ b$ (respectively $a \approx^- b$) replacing $a \approx b$, and with "positive chain" (respectively "negative chain") replacing "chain". (Here we are dealing with r-chains for fixed r.)* $\diamondsuit$

Lemma 9.22 : *If $a, b, c \in \Delta^0 K$ are all distinct and $a \approx^+ b$ and $a \approx^- c$, then $a \in \Pi$.*

Proof : Suppose to the contrary that $a \notin \Pi$. Then $a \in \partial Z_0$ and so infinite subrays of $[a, b]_0$ and $[a, c]_0$ remain a bounded distance apart. We can thus find $x, y \in [a, b]_0$ and $x', y' \in [a, c]_0$ with $d_0(x, y)$ arbitrarily large, but with $d_0(x, x')$ and $d_0(y, y')$ uniformly bounded (by r_1 say). Now we can find a uniform positive chain (an r_2-chain say), $(x_i)_{i \in \mathbf{N}}$ in $\Upsilon(a, b)$ with $x_0 = x$, and a uniform negative chain $(x'_i)_{i \in -\mathbf{N}}$ in $\Upsilon(a, c)$ with $x'_0 = x'$. Setting $x_i = x'_i$ for $i < 0$, we obtain a uniform bi-infinite chain (in fact an $(r_2 + r_1)$-chain) $(x_i)_{i \in \mathbf{Z}}$. We similarly construct a bi-infinite chain, $(y_i)_{i \in \mathbf{Z}}$. Moreover, $x_\infty, y_\infty \in I^+(a, b)$ and $x_{-\infty}, y_{-\infty} \in I^-(a, c)$. Since $a \approx^+ b$ and $a \approx^- c$, we have $x_\infty = y_\infty$ and $x_{-\infty} = y_{-\infty}$. Thus, the chains $(x_i)_i$ and $(y_i)_i$ remains a uniformly bounded distance apart. In particular, $d_0(x, y)$ is uniformly bounded, contradicting the fact that it can be chosen arbitrarily large. $\diamondsuit$

Suppose now that $(K_i)_{i \in \mathbf{Z}}$ is a stratified graph of Farey graphs satisfying (S1)–(S6). In this case, we see that if $\{x, y\}$ and $\{z, w\}$ are linked, then $I^\pm(x, y) \cap I^\pm(z, w) \neq \emptyset$. As in the singly degenerate case, we have:

Lemma 9.23 : *There are complete perfect laminations, L^+ and L^-, on Σ such that $\approx^\pm = \sim_{L^\pm}$.* $\diamondsuit$

Given an interval, $I \subseteq \Delta^0 K$, we define the uniformly quasiconvex set $Q(I) = \bigcup_{i \in \mathbf{Z}} Q_i(I)$ exactly as before. This gives rise to a function $\tau : \Delta^0 K \setminus \Pi \longrightarrow \partial Z$. If $a \in \Pi$ we write $\tau^+(a)$ and $\tau^-(a)$ respectively for the positive and negative endpoints of the constant chain at a. We define $\tau^\pm : \Delta^0 K \longrightarrow \partial Z$ by writing $\tau^\pm(a) = \tau(a)$ for all $a \in \Delta^0 K \setminus \Pi$.

Exactly as in the one-ended case (see Lemma 9.11), we have:

Lemma 9.24 : *If $a, b \in \Delta^0 K$ and $a \approx^\pm b$, then $\tau^\pm(a) = \tau^\pm(b)$.* $\diamondsuit$

Lemma 9.25 : *If $a \in \Pi$, $b \in \Delta^0 K \setminus \Pi$ and $\tau^\pm(a) = \tau(b)$, then $a \approx^\pm b$.* $\diamondsuit$

Lemma 9.26 : *If $a, b \in \Delta^0 K \setminus \Pi$ and $\tau(a) = \tau(b)$, then either $a \approx^+ b$ or $a \approx^- b$.*

Proof : The proof proceeds as in Lemma 9.11. This time, we obtain sequences of bi-infinite chains $(\underline{x}^n)_n$ and $(\underline{y}^n)_n$ with $x^n_0 \to a$ and $y^n_0 \to b$ in $[a, b]_0$, and with $\min\{d_i(x^n_i, y^n_i) \mid i \in \mathbf{Z}\}$ bounded. Passing to a subsequence, either $\min\{d_i(x^n_i, y^n_i) \mid i \in \mathbf{N}\}$ or $\min\{d_i(x^n_i, y^n_i) \mid i \in -\mathbf{N}\}$ is bounded, giving respectively $a \approx^+ b$ or $a \approx^- b$ as required. $\diamondsuit$

Putting these facts together, we note:

Lemma 9.27 : *If $a, b \in \Delta^0 K$, then ($a \approx^+ b$ or $a \approx^- b$) if and only if ($\tau^+(a) = \tau^+(b)$ or $\tau^-(a) = \tau^-(b)$).* $\diamondsuit$

Next, we relate this to the Cannon-Thurston map $\omega : \Delta^0 K \longrightarrow \Delta K$. Exactly as in Lemmas 9.12 and 9.13 respectively, we have:

Lemma 9.28 : *If $a, b \in \Delta^0 K$ and $\tau^+(a) = \tau^+(b)$ or $\tau^-(a) = \tau^-(b)$, then $\omega(a) = \omega(b)$.* $\diamondsuit$

Lemma 9.29 : *If $a, b \in \Delta^0 K$ and $\omega(a) = \omega(b) \notin \Pi$ then $\tau(a) = \tau(b)$. (Note that $a, b \notin \Pi$.)* $\diamondsuit$

Lemma 9.30 : *Suppose $p \in \Pi$ and $a \in \Delta^0 K$. If $\omega(a) = p$ then either $\tau^+(a) = \tau^+(p)$ or $\tau^-(a) = \tau^-(p)$.*

Proof : The proof follows that of Lemma 9.14. Again, we can suppose that $a \notin \Pi$, so by definition, $\tau^+(a) = \tau^-(a) = \tau(a)$. If $\tau(a) \notin \{\tau^+(p), \tau^-(p)\}$, then $\tau^\pm(p) \notin \partial Q(I^n)$ for all sufficiently large n. We see that $d_Z(l(p), Q(I^n)) \to \infty$ as $n \to \infty$. We derive a contradiction as before. $\diamondsuit$

Lemma 9.31 : *Suppose $a, b \in \Delta^0 K$. Then $\omega(a) = \omega(b)$ if and only if $\tau^+(a) = \tau^+(b)$ or $\tau^-(a) = \tau^-(b)$ or there is some $p \in \Pi$ such that ($\tau^+(a) = \tau^+(p)$ or $\tau^-(a) = \tau^-(p)$) and ($\tau^+(b) = \tau^+(p)$ or $\tau^-(b) = \tau^-(p)$).*

Proof : The "if" part follows from Lemma 9.28. For the converse, suppose $\omega(a) = \omega(b)$. If $\omega(a) \notin \Pi$, apply Lemma 9.29. Otherwise, set $p = \omega(a) \in \Pi$, so that $p = \omega(p) = \omega(a) = \omega(b)$. The result follows by applying Lemma 9.30. $\diamondsuit$

Lemma 9.32 : *Suppose $a, b \in \Delta^0 K$. Then $\omega(a) = \omega(b)$ if and only if $a \approx^+ b$ or $a \approx^- b$ or there is some $p \in \Pi$ such that ($a \approx^+ p$ and $b \approx^- p$) or ($a \approx^- p$ and $b \approx^+ p$).*

Proof : Lemmas 9.27 and 9.31 together give us the equivalence of $\omega(a) = \omega(b)$ with the statement that $a \approx^+ b$ or $a \approx^- b$ or there exists $p \in \Pi$ such that ($a \approx^+ p$ or $a \approx^- p$) and ($b \approx^+ p$ or $b \approx^- p$). But this is equivalent to the statement in the lemma, given the transitivity of the relations $\approx^+$ and $\approx^-$. $\diamondsuit$

Note that the criterion of Lemma 9.32 defines the transitive closure of $\approx^+ \cup \approx^-$. This follows either from Lemma 9.32, or directly, applying Lemma 9.22.

Now let N be a doubly degenerate manifold with ends e^+ and e^-, and with $\operatorname{inj} N > 0$. To complete the proof of Theorem 9.1, and hence Theorem 0.2, it suffices to show that the laminations L^+ and L^- are precisely $L(e^+)$ and $L(e^-)$, assuming, of course, that we have chosen the maps ϕ_i so that $\phi_i(\Sigma)$ tends out the end e^+ as $i \to \infty$ and the end e^- as $i \to -\infty$. This follows as exactly as in Lemma 9.19 in the singly degenerate case.

10. Comments on Theorem 0.3.

The following is a consequence of the main result proven in [AnM]:

Proposition 10.1 : *Suppose that G is a finitely generated kleinian group. Suppose that every finitely generated strictly type-preserving surface subgroup of G has locally connected limit set. If $\Lambda(G)$ is connected, then it is locally connected.* ◇

To relate this to the statement in [AnM], a few comments are in order. By Selberg's Lemma, G contains a torsion-free finite-index subgroup. This has the same limit set, and so there is no loss in assuming G to be torsion-free. By Ahlfors's Finiteness Theorem, G is "analytically finite" in the sense of [AnM]. Moreover any "structure subgroup" is a finitely generated strictly type-preserving subgroup. Only degenerate structure subgroups are of interest, since non-degenerate ones are quasifuchsian and thus have circle limits sets. Thus, Proposition 2.1 of [AnM] implies Propositon 10.1 as we have stated it.

Now the hypotheses on loxodromics corresponds to positive injectivity radius away from cusps. For type-preserving surface groups this implies that the limit set is locally connected — by [CannT,Min1] in the case without parabolics, and by Theorem 0.1 in the case with parabolics. Putting these facts together now gives Theorem 0.3.

Of course, Theorem 0.3 was already known if we assume also that there are no parabolics, by the above remarks (see also [K]). Also for geometrically finite groups (in dimension 3) it was shown in [AnM]. Thus, Theorem 0.3 generalises both these results. The general case of finitely generated groups remains open.

We remark that there is a closely related question due to Thurston. Suppose G is a finitely generated kleinian group. Suppose that G' is an isomorphic geometrically finite kleinian group with each parabolic in G' parabolic in G. Does there exist a continuous equivariant map from $\Lambda(G')$ to $\Lambda(G)$? The results of [CannT,Min1] answer this affirmatively for strictly type-preserving surface groups. The case where $\Lambda(G)$ is connected and where G has a positive lower bound on the translation lengths of loxodromics and has no parabolics was dealt with in [K]. It is hoped that the methods of the present paper will lead to generalisations of this allowing for parabolics.

References.

[AlDP] R.C.Alperin, W.Dicks, J.Porti, *The boundary of the Gieseking tree in hyperbolic three-space* : Topology Appl. **93** (1999) 219–259.

[AnM] J.W.Anderson, B.Maskit, *On the local connectivity of limit sets of kleinian groups* : Complex Variables **31** (1996) 177–183.

[BeF] M.Bestvina, M.Feighn, *A combination theorem for negatively curved groups* : J. Differential Geom. **35** (1992) 85–101.

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

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

[Bow2] B.H.Bowditch, *Relatively hyperbolic groups* : preprint, Southampton (1997).

[Bow3] B.H.Bowditch, *Stacks of hyperbolic spaces and ends of 3-manifolds* : in preparation.

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

[CanaEG] 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.

[Cann] J.W.Cannon, *The Peano curve defined by a punctured torus bundle* : in preparation.

[CannD] J.W.Cannon, W.Dicks, *On hyperbolic once-punctured-torus bundles* : to appear in Geom. Dedicata.

[CannT] J.W.Cannon, W.P.Thurston, *Group invariant Peano curves* : preprint (1989).

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

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

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

[Gr1] M.Gromov, *Hyperbolic groups* : in "Essays in Group Theory" (ed. S.M.Gersten) M.S.R.I. Publications No. 8, Springer-Verlag (1987) 75–263.

[Gr2] M.Gromov (edited by J.LaFontaine and P.Pansu), *Metric structures for Reimannian and non-Riemannian spaces* : Progress in Mathematics No. 152, Birkhäuser (1998).

[K] E.Klarreich, *Semiconjugacies between Kleinian group actions on the Riemann sphere* : Amer. J. Math. **121** (1999) 1031–1078.

[Mc] C.T.McMullen, *Local connectivity, Kleinian groups and geodesics on the blowup of the torus* : Invent. Math. **146** (2001) 35–91.

[Min1] Y.N.Minsky, *Teichmüller geodesics and ends of hyperbolic 3-manifolds* : Topology **32** (1993) 625-647.

[Min2] Y.N.Minsky, *On rigidity, limit sets, and ends of hyperbolic 3-manifolds* : J. Amer. Math. Soc. **7** (1994) 539–588.

[Min3] Y.N.Minsky, *The classification of punctured torus groups* : Ann. of Math. **149**

(1999) 559-626.

[Mit1] M.Mitra, *Cannon-Thurston maps for trees of hyperbolic metric spaces* : J. Differential Geom. **48** (1998) 135–164.

[Mit2] M.Mitra, *Cannon-Thurston maps for hyperbolic group extenstions* : Topology **37** (1998) 527–538.

[O] J.-P.Otal, *Le théorème d'hyperbolisation pour les variétés fibrées de dimension 3* : Astérisque No. 235, Société Mathématique de France (1996).

[T1] W.P.Thurston, *The geometry and topology of 3-manifolds* : notes, Princeton (1979).

[T2] W.P.Thurston, *Hyperbolic structures on 3-manifolds II: Surface groups and manifolds which fiber over the circle* : preprint, Princeton (1986).