

STACKS OF HYPERBOLIC SPACES AND ENDS OF 3-MANIFOLDS

B. H. BOWDITCH

ABSTRACT. In this paper we introduce the notion of a “stack” of geodesic spaces. Loosely speaking, this consists of a geodesic space decomposed into a sequence of “sheets” indexed by a set of consecutive integers. A stack is said to be “hyperbolic” if it is Gromov hyperbolic and its sheets are uniformly Gromov hyperbolic. We define a Cannon-Thurston map for such a stack, and show that the boundary of a one-sided proper hyperbolic stack is a dendrite. If the stack arises from a sequence of closed hyperbolic surfaces with a lower bound on injectivity radius, then this allows us to define an “ending lamination” on the surface. We show that the ending lamination has a certain dynamical property that implies unique ergodicity. We also show that such a sequence is a bounded distance from a Teichmüller ray — a result obtained independently by Mosher. This can be reinterpreted in terms of the Bestvina-Feighn flaring condition, and gives a coarse geometrical characterisation of Teichmüller rays. Applying this to a simply degenerate end of a hyperbolic 3-manifold with bounded geometry, we recover Thurston’s ending lamination conjecture, proven by Minsky, in this case. Various related issues are discussed.

2010 Mathematics Subject Classification: 57M50, 20F65

Keywords: hyperbolic space, boundary, Cannon-Thurston map, Teichmüller geodesic, 3-manifold, lamination.

1. INTRODUCTION

There has been much recent work relating the geometry of hyperbolic 3-manifolds to the large-scale geometry of Teichmüller space. One major landmark in this has been the positive resolution of Thurston’s Ending Lamination Conjecture by Minsky, Brock and Canary [Min6, BroCM1, BroCM2]. This work has many consequences. One particular direction which we note here is the construction of the Cannon-Thurston map see [CannT, Mit1, Mit2, Bow3, Mj1, Mj2, Mj3]. The present paper focuses on the bounded geometry case, and sets this in the broader context of coarse geometry. (It thus builds on the work of

Date: February 2012.

[Mit1, Mit2].) In the course of doing so, we obtain a number of results regarding Teichmüller geodesics, ending laminations etc. We also give results about more general stacks.

We will give an outline of the main results in Section 2. In this introduction we will describe some background which set the results in context.

Let Σ be a closed orientable surface, and let $\mathcal{T} = \mathcal{T}(\Sigma)$ be the Teichmüller space of Σ . This can be thought of as the space of marked hyperbolic structures on Σ , or equivalently, the space of conformal structures on Σ (see [Ga, IT]). We fix a small positive constant, η , (less than the Margulis constant) and refer to the closed subset of \mathcal{T} consisting of those hyperbolic metrics of injectivity radius at least η as the “thick part” of \mathcal{T} .

Suppose that M is a complete hyperbolic 3-manifold homeomorphic to $\Sigma \times \mathbb{R}$. We refer to the first and second coordinates as the “horizontal” and “vertical” directions respectively. Associated to each end of M we have an “end invariant” which is either a point of \mathcal{T} , or a lamination (the “ending lamination”) which, very loosely, can be thought of as an ideal point of \mathcal{T} . The Ending Lamination Theorem [Min6, BroCM1, BroCM2] tells us that M is determined up to isometry by these two end invariants. To simplify the discussion in this introduction, we assume that M is doubly degenerate, that is, each end invariant is a lamination. Again very loosely, we can think of the geometry of M as being described by a path in \mathcal{T} , where the vertical direction is the time parameter, and the horizontal direction is, in some sense, represented by a point of \mathcal{T} . In [T, Bon], this is described in terms of pleated surfaces.

In general, the above requires a lot of work to make precise. In the bounded geometry case, however, one can give an outline description more simply. Suppose we have a path β in the thick part of \mathcal{T} , which is Lipschitz with respect to the Teichmüller metric. Up to bilipschitz equivalence, there is a canonical way of putting a Riemannian metric on Σ so that, for all t , the induced metric $\Sigma \times \{t\}$ is uniformly bilipschitz equivalent to the hyperbolic metric Σ given by $\beta(t) \in \mathcal{T}$. (We say more about this construction in Section 22.) We write $P(\beta)$ for the resulting space.

Suppose that M is a doubly degenerate hyperbolic 3-manifold of positive injectivity radius. It follows from the work of [T, Bon] using pleated surfaces, that M is bilipschitz equivalent to $P(\beta)$ for some such path β . Moreover, β is proper, and converges on the two ending laminations of M (necessarily distinct). In fact, we only require

that the universal covers $\tilde{P}(\beta)$ and $\tilde{M} \equiv \mathbb{H}^3$ are Γ -equivariantly quasi-isometric, where $\Gamma = \pi_1(\Sigma)$ and \mathbb{H}^n denotes hyperbolic n -space. (This is much easier to verify.) In [Min1] Minsky shows that, in this case, β is a bounded distance from a Teichmüller geodesic, α . Thus α is determined by the ending laminations. Moreover, $P(\alpha)$ is bilipschitz equivalent to $P(\beta)$ hence to M . It follows that M is determined up to bilipschitz equivalence, hence by a result of Sullivan [Su], up to isometry. This gives us the Ending Lamination Theorem in this case. Again, to apply [Su], we only require a-priori that the universal covers be equivariantly quasi-isometric. In this paper, we will give another approach the result of Minsky above. We will show that bi-infinite (or semi-infinite) path β in the thick part of Teichmüller space is a bounded distance from a Teichmüller geodesic if and only if $\tilde{P}(\beta)$ is Gromov hyperbolic. This result has been obtained independently by Mosher, by different methods, see [Mo]. Note that if β arises from a doubly degenerate 3-manifold, then $\tilde{P}(\beta)$ is quasi-isometric to \mathbb{H}^3 , and hence hyperbolic. It therefore follows that β is a bounded distance from a Teichmüller geodesic, and the argument can be completed as before.

Another consequence of the above is a construction of the Cannon-Thurston map in this case. This gives a Γ -equivariant surjective continuous map of the circle, $\partial\mathbb{H}^2$, to the 2-sphere, $\partial\mathbb{H}^3$. It originates in [CannT]. Many variations and generalisations have been given since, see for example, [AIDP, Mit1, Mit2, Bow3, Mj2, Mj3].

In this paper, we will reset this in the context of coarse geometry. To this end, we will replace the path β by a sequence, $(\Sigma_i)_i$, of surfaces, indexed by \mathbb{Z} . (For example, $\Sigma_i = \beta(i)$.) We will assume that Σ_i lies in the thick part of \mathcal{T} . We also assume that the Teichmüller distance between Σ_i and Σ_{i+1} is bounded above. The latter statement is equivalent to the existence of a uniform equivariant quasi-isometry between consecutive universal covers, X_i and X_{i+1} , where $X_i = \tilde{\Sigma}_i \equiv \mathbb{H}^2$. Out of this we build a “stack”, which is a path metric space, Ξ , containing the X_i as subspaces, which we call the “sheets” of Ξ . The stack Ξ is well defined up to quasi-isometry. If $\Sigma_i = \beta(i)$, for a Lipschitz path, β , as above, then Ξ is equivariantly quasi-isometric to $\tilde{P}(\beta)$.

We can generalise this to construct a stack out of any sequence, $(X_i)_i$ of geodesic spaces X_i together with uniform quasi-isometries between consecutive X_i . If each X_i is uniformly Gromov hyperbolic, then it follows from an argument of Bestvina and Feighn [BeF] that Ξ is Gromov hyperbolic if and only if it satisfies a certain “flaring” condition, which we express here in terms of quasiconvexity of the distance functions.

We refer to such a stack as a “hyperbolic stack”. Note the the boundaries, ∂X_i , are canonically homeomorphic to each other, so we get a “horizontal boundary”, $\partial^0 \Xi \cong \partial X_i$, of Ξ . We will show that there is a canonical “Cannon-Thurston map”, $\omega : \partial^0 \Xi \longrightarrow \partial \Xi$ (cf. [Mit1, Mit2]).

In this introduction, we have only discussed bi-infinite stacks, indexed by \mathbb{Z} , though we can similarly deal with stacks indexed by any set of consecutive integers. We will give a summary of the main result in Section 2.

Since the original draft of this paper was written, there has been considerable progress in several directions. As noted above, we now have a complete proof of the Ending Lamination Theorem, [Min6, BroCM1, BroCM2] (see also [Bow4] for some discussion). This has lead to constructions of Cannon-Thurston maps in a more general setting, see in particular, [Mj2, Mj3].

2. SUMMARY OF RESULTS

We give a summary of the main results. More precise statements will be given in later sections. First we recall some basic facts and definitions.

A *geodesic* in a metric space is a path whose length is equal to the distance between its endpoints. (We will usually assume this to be parameterised by arc-length.) A *geodesic space* is a metric space in which every pair of points are connected by a geodesic. A metric space is *proper* if it is complete and locally compact. It’s not hard to see that a geodesic space is proper if and only if every closed bounded subset is compact.

We recall the notion of a (Gromov) hyperbolic space [Gr1, GhH]. This is a geodesic space in which every geodesic triangle is “thin”, that is, every side lies a bounded distance from the union of the other two sides. Such a space, X , has associated to it an ideal “boundary”, ∂X , which a metrisable topological space, well defined up to homeomorphism. If X is proper, then ∂X is compact.

We also recall the standard notion of a quasi-isometry between two geodesic spaces, X, Y . This is a map $\phi : X \longrightarrow Y$ (not necessarily continuous) which respects distances up to linear bounds, and such that Y is a bounded neighbourhood of $\phi(X)$. (We will describe a variation on this definition, in terms of relations, in Section 3.) If X is hyperbolic, then so is Y , and ϕ induces a homeomorphism from ∂X to ∂Y .

We now introduce the notion of a “stack” of spaces (cf. [BeF, Mit1]). For the purposes of this summary, we deal with this informally. It will be discussed further in Section 6.

A “stack” will consist of a geodesic space, (Ξ, ρ) , with a collection, $(X_i)_{i \in \mathcal{I}}$, of closed connected subsets indexed by a set $\mathcal{I} \subseteq \mathbb{Z}$ of consecutive integers. We assume (in this section) that X_i separates X_{i-1} from X_{i+1} , that $\rho(X_i, X_{i+1})$ is bounded below by a positive constant, and that Ξ is a bounded neighbourhood of $\bigcup_{i \in \mathcal{I}} X_i$. Moreover, we assume that there is a uniform quasi-isometry, ϕ_i , between X_i and X_{i+1} in their induced path metrics and such that $\rho(x, \phi_i(x))$ is uniformly bounded above for all $x \in X_i$. Here “uniform” means that the relevant bounds are independent of i . (In fact, we can use slightly weaker hypotheses than those listed above, which are more natural in the coarse geometry setting, see Section 6.) We refer to X_i as a “sheet” of Ξ , which we imagine as being “horizontal”. By a *chain* in Ξ , we mean a sequence, $(x_i)_i$, of points indexed by $i \in \mathcal{I}$, such that $x_i \in X_i$ for all i and such that $\rho(x_i, x_{i+1})$ is bounded above (for example, if $x_{i+1} = \phi_i(x_i)$). We can interpolate between the consecutive x_i to give us a “vertical” path which is transverse to the X_i and which proceeds at a linear rate in the vertical direction — in particular, it is a quasigeodesic in Ξ .

If, conversely, we start with a sequence, (X_i, ρ_i) , of geodesic spaces, and with uniform quasi-isometries, ϕ_i , between consecutive X_i , we can construct a stack of this sort. For example, we can take the disjoint union, $\bigsqcup_i X_i$, and, for all i , connect each $x \in X_i$ to $\phi_i(x) \in X_{i+1}$ by a real interval of fixed bounded length. Provided this length is large enough, ρ_i will be the induced path-metric on X_i . There are many possible variations on this construction, but any sensible procedure will yield a stack, Ξ , that is equivalent up to quasi-isometry. Moreover, if each X_i is proper, we can arrange that (Ξ, ρ) be proper.

Suppose now that each of the sheets X_i is uniformly Gromov hyperbolic. Then the boundaries, ∂X_i , are all canonically homeomorphic. We can therefore define a “horizontal” boundary, $\partial^0 \Xi$, of Ξ , which is canonically homeomorphic to each ∂X_i .

The paper of Bestvina and Feighn [BeF], now gives us a criterion under which the stack Ξ is itself hyperbolic. We express this as follows:

Theorem 2.1. *Let Ξ be a stack, where all sheets are uniformly hyperbolic. Then Ξ is hyperbolic if and only if whenever $(x_i)_i$ and $(y_i)_i$ are chains, the function $[i \mapsto \rho_i(x_i, y_i)]$ is uniformly quasiconvex.*

The term “quasiconvex” refers to a coarse version of convexity defined in Section 5. “Uniform” means that the relevant constants are independent of the chains $(x_i)_i$ and $(y_i)_i$. Intuitively, this can be thought

of as saying that whenever two chains diverge, then they do so at a uniform rate. It is essentially equivalent to the “flaring” condition of [BeF]. Theorem 2.1 is a consequence of Theorems 6.2 and 6.4.

We will refer to a stack arising as in Theorem 2.1 as a *hyperbolic stack*. Note that it has its own boundary, $\partial\Xi$, as a hyperbolic space. Following the argument of [Mit1, Mit2], we see:

Theorem 2.2. *There is a natural continuous map $\omega : \partial^0\Xi \longrightarrow \partial\Xi$.*

This is referred to as the “Cannon-Thurston map”. Such a map was introduced in a slightly different context in [CannT]. Theorem 2.2 is an immediate consequence of Lemma 7.4 of this paper.

There are essentially three cases. If \mathcal{I} is finite, then the inclusion of any sheet, X_i , into Ξ is a quasi-isometry. This induces a homeomorphism, $\omega : \partial X_i \longrightarrow \partial\Xi$, so this case is elementary. The interesting cases are where $\mathcal{I} = \mathbb{N}$ or where $\mathcal{I} = \mathbb{Z}$. We refer to these as “semi-infinite” or “bi-infinite”.

For the purposes of this section, we assume henceforth that $\mathcal{I} = \mathbb{Z}$. In this case, we write $\Xi = \Xi^- \cup \Xi^+$, where $\Xi^- \cap \Xi^+ = X_0$, and where Ξ^- and Ξ^+ are semi-infinite stacks with sheets X_i , indexed by $-\mathbb{N}$ and by \mathbb{N} respectively. (Again, we only really need a coarse version of these statements.) Note that we can identify $\partial^0\Xi = \partial^0\Xi^- = \partial^0\Xi^+$.

Of course, the path-metric induced on Ξ^\pm may be quite different from ρ . Nevertheless, we have:

Theorem 2.3. *If Ξ is hyperbolic, then Ξ^- and Ξ^+ are hyperbolic.*

The proof we give here passes via Theorem 2.2 (see Proposition 8.4).

For the purposes of this summary, we will view the case of a semi-infinite stack as subsumed within the case of a bi-infinite stack (though all the statements we make hold intrinsically). Note that we have three Cannon-Thurston maps, namely $\omega : \partial^0\Xi \longrightarrow \partial\Xi$, $\omega^- : \partial^0\Xi \longrightarrow \partial\Xi^-$ and $\omega^+ : \partial^0\Xi \longrightarrow \partial\Xi^+$.

We now restrict to the case where each X_i is proper, so that Ξ can be assumed proper. In this case, the boundaries, $\partial^0\Xi$, $\partial\Xi^\pm$, $\partial\Xi$, are all compact.

Theorem 2.4. *The are natural continuous maps, $\tau^\pm : \partial\Xi^\pm \longrightarrow \partial\Xi$ such that $\omega = \tau^\pm \circ \omega^\pm$. Moreover, we have $\tau^-(\partial\Xi^-) \cup \tau^+(\partial\Xi^+) = \partial\Xi$ and $\tau^-(\partial\Xi^-) \cap \tau^+(\partial\Xi^+) = \omega(\partial^0\Xi)$.*

This is a combination of results of Section 8. See in particular, Lemma 8.11.

Recall that a *Peano continuum* is a connected locally connected compact topological space. A *dendrite* is a metrisable Peano continuum in which every pair of points are separated by a cut point.

Theorem 2.5. *$\partial\Xi$ is a metrisable Peano continuum, and $\partial\Xi^\pm$ is a dentrite.*

(See Proposition 10.2.)

We now restrict further to the case where each sheet, X_i , is isometric to the hyperbolic plane, \mathbb{H}^2 . Thus, $\partial^0\Xi \equiv \partial X_i \equiv \partial\mathbb{H}^2$ is homeomorphic to the circle, S^1 . In this case, the Cannon-Thurston map is surjective:

Theorem 2.6. *If Ξ is a hyperbolic stack of hyperbolic planes, then $\omega^\pm(\partial^0\Xi) = \partial\Xi^\pm$ (and so $\omega(\partial^0\Xi) = \partial\Xi$).*

For a proof, see Proposition 11.1. We can elaborate on this as follows.

A “lamination” on \mathbb{H}^2 is a disjoint collection bi-infinite geodesics called “leaves” whose union is closed in \mathbb{H}^2 . One can define an equivalence relation on $\partial\mathbb{H}^2$ by deeming two points to be equivalent if they are not separated by the closure any leaf in $\mathbb{H}^2 \cup \partial\mathbb{H}^2$. (In particular, the two endpoints of any leaf are equivalent.) One can verify that the quotient of $\partial\mathbb{H}^2$ is hausdorff. In fact (see Proposition 12.3):

Theorem 2.7. *If Ξ is a hyperbolic stack of hyperbolic planes, then there are unique laminations, Λ^- and Λ^+ , such that each Cannon-Thurston map, $\omega^\pm : \partial^0\Xi \rightarrow \partial\Xi^\pm$ is the quotient map under the associated equivalence relation on $\partial\mathbb{H}^2$.*

Of course, this implies a particular identification of $\partial^0\Xi$ with $\partial\mathbb{H}^2$, which we cannot expect to be canonical. We note that Λ^+ and Λ^- are transverse in sense that a leaf of Λ^+ can intersect a leaf of Λ^- in at most one point.

We now suppose that we have a group, Γ , acting by isometry on Ξ , and fixing setwise each sheet, X_i , of Ξ . We suppose that the action of Γ on each X_i is properly discontinuous and cocompact and preserves orientation. Thus, $\Sigma_i = X_i/\Gamma$ is a hyperbolic surface. We can write $\Gamma = \pi_1(\Sigma)$, where Σ is a closed topological surface which has a natural homotopy equivalence to each Σ_i . We can therefore think of each Σ_i as determining a point of the Teichmüller space, \mathcal{T} , of Σ . We will suppose, in addition, that the injectivity radius of each Σ_i is bounded below. In other words, each Σ_i lies in the η -thick part of \mathcal{T} for some fixed $\eta > 0$. In this case, one can show that consecutive Σ_i lie a bounded distance apart in \mathcal{T} , and so we can interpolate between them by paths of bounded length to give us a stable path in \mathcal{T} , that is, a lipschitz path lying in the thick part of \mathcal{T} , as defined in the introduction. Conversely, if we have a stable path, β , in \mathcal{T} , we can construct a stack $\Xi(\beta)$ of this sort, so that $\Sigma(i) = \beta(i)$ for all integers, i . Indeed we can take $\Xi(\beta)/\Gamma$ to be a riemannian manifold diffeomorphic to $\Sigma \times \mathbb{R}$. By any

sensible construction, $\Xi(\beta)$ will be well defined up to Γ -equivariant quasi-isometry. We will say more about this in Section 22.

If Ξ is also hyperbolic, then we have two laminations, Λ^- and Λ^+ , given by Theorem 2.7. These will now be Γ -equivariant, and will descend to a geodesic lamination on each Σ_i . This will be the realisation in Σ_i of a fixed abstract lamination, $\Lambda_\Sigma^\pm = \Lambda^\pm/\Gamma$ in Σ (see [CasB]). We refer to Λ_Σ^+ and Λ_Σ^- as the negative and positive “ending laminations” of the stack. If $\Xi = \Xi(\beta)$, then they depend only on β .

Only certain laminations can arise in this way. Let Λ_Σ be a lamination on Σ . Fix any hyperbolic structure on Σ and realise Λ_Σ as a geodesic lamination in Σ .

Definition. We say that Λ_Σ is *regular* if there is a linear function, f , such that if $E \subseteq \Sigma$ is any compact interval lying in any leaf of Λ_Σ and γ is any homotopically non-trivial curve in $\Sigma \setminus E$, then $\text{length}(E) \leq f(\text{length}(\gamma))$.

One can check that this condition is independent of the choice of hyperbolic structure (though, of course, the function, f , may change), see Propositions 13.9 and 13.11. Intuitively, this condition says that the leaves of Λ_Σ fill up Σ at a uniform rate.

We show (Proposition 13.6):

Theorem 2.8. *A regular lamination is uniquely ergodic.*

That is, any two transverse measures on the lamination agree up to scale. Moreover, we will show (Proposition 17.4).

Theorem 2.9. *Any ending lamination of a hyperbolic surface stack is regular.*

Moreover, any such lamination admits a transverse measure, and so, by Theorem 2.8 has a unique structure as a projective (measured) lamination.

Suppose that $\alpha : \mathbb{R} \rightarrow \mathcal{T}$ is a stable Teichmüller geodesic. It has associated to it two laminations. These arise from the singular foliations corresponding the real and imaginary parts of the associated quadratic differential (see, for example, [Ga] or [IT]).

Recall that a stable path is a Lipschitz path lying in the η -thick part of \mathcal{T} .

Theorem 2.10. *If α is a stable Teichmüller geodesic, then $\Xi(\alpha)$ is a hyperbolic surface stack. Moreover, the ending laminations are precisely the foliations associated to α .*

This statement is also proven in [Mo]. It has an interpretation in terms of singular Sol geometry (cf. [CannT]), as discussed in Section 18, see Proposition 18.8 and Lemma 18.9.

Note that, putting together Theorems 2.8, 2.9 and 2.10, we conclude that each of the two laminations associated to a stable Teichmüller geodesic is uniquely ergodic. We therefore recover the result of Masur [Mas].

In fact, there is a converse to Theorem 2.10 as follows. This has been obtained independently by Mosher [Mo].

Theorem 2.11. *Suppose that $\beta : \mathbb{R} \rightarrow \mathcal{T}$ is a stable path. Then β is a bounded distance from a Teichmüller geodesic if and only if $\Xi(\beta)$ is Gromov hyperbolic.*

The statement also holds if \mathbb{R} is replaced by $[0, \infty)$. The result will be proven in Section 20 (see Theorem 20.1).

Combining Theorems 2.1 and 2.11, we can give a simple criterion for recognising such paths among stable paths. We can formulate this as follows.

Given some $\sigma \in \mathcal{T}$, and a free homotopy class, γ , of closed curve in Σ , we write $l(\sigma, \gamma)$ for the length of γ when realised as a closed geodesic in the hyperbolic structure determined by σ . We will show:

Theorem 2.12. *Suppose that $\beta : \mathbb{R} \rightarrow \mathcal{T}$ is a stable path. Then β is a bounded distance from a Teichmüller geodesic if and only if the maps $[t \mapsto l(\beta(t), \gamma)] : \mathbb{R} \rightarrow [0, \infty)$ are uniformly quasiconvex for all simple closed curves, γ .*

By this we mean that the constants of quasiconvexity are independent of γ . Again, the same result holds if the domain, \mathbb{R} , is replaced by $[0, \infty)$. This will be proven in Section 21.

One key application of Theorem 2.11 is already discussed in the introduction, and will be elaborated on in Section 22. Let M be a hyperbolic 3-manifold homeomorphic to $\Sigma \times \mathbb{R}$, and without cusps. It has two ending laminations, which, in general are geometric laminations. In the case where M has positive injectivity radius, one can see, via the result of Masur [Mas], that these are uniquely ergodic. (This is also a consequence of Theorems 2.8 and 2.9 here.) In other words, they have a unique structure as projective laminations, and so can be thought of points of the Thurston boundary, $\partial\mathcal{T}$, of \mathcal{T} . As we describe in Section 22, we recover the Ending Lamination Theorem in this case:

Theorem 2.13. *If $M \cong \Sigma \times \mathbb{R}$ is a complete doubly degenerate hyperbolic 3-manifold with positive injectivity radius, then M is determined, up to isometry, by its end invariants.*

This is also valid where M is singly degenerate, or quasifuchsian.

3. COARSE GEOMETRY

In this section, we explain some of the conventions and principles of coarse geometry that we shall be using throughout this paper. In particular, we will give definitions of the terms “path-metric”, “geodesic space”, “net”, “quasi-isometry”, “equivalence” of quasi-isometries, “quasi-isometric map”, “straight” subset, and the “stable length” of an isometry.

The term “coarse geometry” might be loosely interpreted as the geometry of path-metric spaces defined up to quasi-isometry. Most of the spaces with which we shall be dealing have various constants or functions (of hyperbolicity, quasiconvexity etc.) associated to them. Such constants or functions will be termed “parameters”. We shall speak of other constants or functions as being “uniform” to mean that they depend only these parameters. Precisely which parameters are involved should be clear from the context.

At times, it will be convenient to adopt the following convention (cf. [Bow2]). Given $x, y \in \mathbb{R}$ and $k \geq 0$, we shall write $x \simeq_k y$, $x \preceq_k y$ and $x \ll_k y$ to mean respectively $|x - y| \leq k$, $x \leq y + k$ and $x \leq y - k$. We shall frequently drop the subscript, k , and behave as though the relations \simeq and \preceq were transitive. Unless it is otherwise stated, it is to be assumed that the additive constants involved depend (in some explicit way) only on the parameters of the spaces with which we are dealing.

Let (X, ρ) be a metric space. If $x, y \in X$, we write $x \simeq y$ to mean that $\rho(x, y) \simeq 0$. Given $Q \subseteq X$ and $r \geq 0$, write $N(Q, r) = \{y \in X \mid \rho(y, Q) \leq r\}$. Given $P, Q \subseteq X$, write $\text{hd}(P, Q) = \text{hd}_\rho(P, Q) = \inf\{r \mid P \subseteq N(Q, r), Q \subseteq N(P, r)\} \in [0, \infty]$ for the *Hausdorff distance* between P and Q . We write $P \simeq Q$ if $\text{hd}(P, Q) \simeq 0$.

A path, α , connecting x to y in X is *geodesic* if $\text{length}(\alpha) = \rho(x, y)$. We say that (X, ρ) is a *geodesic space* if every pair of points are connected by a geodesic. We have the slightly weaker notion of a *path-metric* space, where given any $x, y \in X$ and any $\epsilon > 0$, x and y are connected by a path of length at most $\rho(x, y) + \epsilon$. Any subset of X has a (possibly) infinite induced path-metric. Any proper (i.e. complete locally compact) path-metric space is geodesic. For simplicity, we shall only explicitly deal with geodesic spaces. However, all results (appropriately phrased) are easily seen to remain valid for path-metric spaces, on the insertion of ϵ 's in the appropriate places.

A subset $Q \subseteq X$ is r -*quasidense* if $X = N(Q, r)$. It is s -*separated* if $\rho(x, y) \geq s$ for all distinct $x, y \in Q$. An r -*net* is an $(r/2)$ -separated r -quasidense subset. Every metric space admits an r -net for every $r > 0$.

Let (X, ρ) and (Y, σ) be path-metric spaces. Here it will be convenient to define a “quasi-isometry”, ϕ , between X and Y formally as a particular kind of relation, denoted $\sim_\phi \subseteq X \times Y$. We write ϕ^{-1} for the *inverse* relation, i.e. $y \sim_{\phi^{-1}} x$ if and only if $x \sim_\phi y$. We write $\phi(X) = \{y \in Y \mid (\exists x \in X)(x \sim_\phi y)\}$, and similarly define $\phi^{-1}(Y) \subseteq X$. We say that ϕ is a *quasi-isometry* if $\phi(X)$ is quasidense in Y , $\phi^{-1}(Y)$ is quasidense in X , and if there is a linear function F such that $\sigma(y, y') \leq F(\rho(x, x'))$ and $\rho(x, x') \leq F(\sigma(y, y'))$ for all $x, x' \in X$ and $y, y' \in Y$ satisfying $x \sim_\phi y$ and $x' \sim_\phi y'$.

We shall say that two quasi-isometries ϕ and ψ are *equivalent* if $\sim_\phi \cup \sim_\psi \subseteq X \times Y$ is a quasi-isometry. Note that, if this is the case, $\phi(\{x\}) \cup \psi(\{x\})$ has bounded diameter for all $x \in X$. We shall say that ψ is *weakening* of ϕ if it is equivalent to ϕ and if $\sim_\phi \subseteq \sim_\psi$. In particular, if ϕ is any quasi-isometry, and $r \geq 0$, we can define another quasi-isometry, ψ , which we refer to as the r -*weakening* of ϕ by $x \sim_\psi y$ if and only if there exist $x' \in X$ and $y' \in Y$ satisfying $\rho(x, x') \leq r$, $\sigma(y, y') \leq r$ and $x' \sim_\phi y'$. By choosing r large enough we obtain $\psi(X) = Y$ and $\psi^{-1}(Y) = X$. In this way we can view ψ as a “multivalued function”. We shall sometimes abuse notation by writing $\psi(x)$ or indeed $\phi(x)$ for some choice of $y \in Y$ with $x \sim_\psi y$.

Note that we can compose quasi-isometries in the obvious way, namely, if $\phi_1 : X \rightarrow Y$ and $\phi_2 : Y \rightarrow Z$ are quasi-isometries, we define the quasi-isometry $\psi : X \rightarrow Z$ by $x \sim_\psi z$ if $(\exists y \in Y)(x \sim_{\phi_1} y$ and $y \sim_{\phi_2} z)$. Such composition is consistent with composition of functions, at least up to uniformly bounded distance. Note also that $\phi^{-1}\phi x \simeq x$ for all $x \in X$.

Suppose ϕ is a quasi-isometry between path-metric spaces, X and Y , and $r > 0$. We construct a space Z by taking disjoint copies of X and Y and connecting every pair of points $x \in X$ and $y \in Y$ with $x \sim_\phi y$ by an arc of length r , and equipping Z with the induced path-metric. Usually r will be appropriately chosen with respect to the parameters (of the quasi-isometry etc.). Note that we recover the original metrics on X and Y as the induced path-metrics. If X and Y are both proper, we can modify the construction so that Z is proper. To do this, choose nets in X and Y , take an appropriate weakening, ψ of ϕ , in the sense described above, and connect only points in the respective nets that are related under ψ . Clearly this depends on various choices. In any case, the resulting space, Z , is quasi-isometric to both X and Y . We shall denote it by $X \vee_\phi Y$. In practice, we shall

be applying this construction to infinite sequences of spaces. The result is always well defined up to uniform quasi-isometry.

Suppose, again, that (X, ρ) and (Y, σ) are path-metric spaces. A map $f : X \rightarrow Y$ (not necessarily continuous) is said to be *straight* if there are functions, $F_1, F_2 : [0, \infty) \rightarrow [0, \infty)$ such that for all $x, y \in X$ we have $F_1(\rho(x, y)) \leq \sigma(f(x), f(y)) \leq F_2(\rho(x, y))$, where $F_1(t) \rightarrow \infty$ as $t \rightarrow \infty$. Clearly, F_2 can always be taken to be linear. If F_1 can also be taken to be linear, we refer to ϕ as a *quasi-isometric* map. Clearly, a quasi-isometric map with quasidense image is (or gives rise to) a quasi-isometry. Indeed this is true of any straight map with quasidense image. If X is a subinterval of the real line, we refer to a quasi-isometric map (or its image) as a *quasigeodesic*. Usually it is convenient to assume that quasigeodesics are continuous.

If $X \subseteq Y$, we refer to X as a *straight* subset if its inclusion is a straight map with respect to the induced path-metric on X . We can replace straight maps by straight subspaces by a construction similar to that of $X \vee_\phi Y$ described earlier. Suppose that $f : X \rightarrow Y$ is a straight map. This time we connect x to $f(x)$ by an interval of a fixed length, to obtain a path-metric space Z . We recover the original metrics on X and Y as the induced path metrics. The inclusion of X in Z is straight, and the inclusion of Y in Z is a quasi-isometry. Again, we can assume this construction preserves properness, by restricting the construction to a net in X .

Suppose that $f, g : X \rightarrow Y$ are straight maps of path-metric spaces with $\text{hd}(f(X), g(X)) \leq r$. The relation, \sim_ϕ on $X \times Y$ defined by $x \sim_\phi y$ if $\rho(f(x), g(y)) \leq r$ gives us a quasi-isometry, ϕ , from X to Y . Note that we can simultaneously extend f and g to a straight map of $X \vee_\phi Y$ into Z .

Finally, given an isometry of a metric space, $\psi : X \rightarrow X$, we define its *stable length* (or stable translation distance) as $\text{sl}(\psi) = \lim_{n \rightarrow \infty} \rho(x, \psi^n(x))$ where $x \in X$. This is well defined and independent of the choice of x . Note that $\text{sl}(\psi^n) = n \text{sl}(\psi)$ for all $n \in \mathbb{N}$.

4. HYPERBOLIC SPACES

In this section, we give a brief summary of relevant facts relating to hyperbolic spaces as defined by Gromov [Gr1]. Further expositions can be found in [Bow2, CoDP, GhH, Sh]

Let (X, ρ) be a hyperbolic geodesic space, and let ∂X be its boundary. Thus, $X \cup \partial X$ carries a natural topology. This is compact if X is proper (though we shall not assume that in this section). We shall use the notations \simeq , \preceq and \ll as defined in the last section. The additive

constants involved will depend on the constants of hyperbolicity (and on quasiconvexity etc.).

Given $x, y \in X \cup \partial X$, we shall write $[x, y]$ for some choice of geodesic connecting x to y (with the convention that $[x, x] = \{x\}$). This is well defined up to bounded Hausdorff distance. If $z, w \in [x, y]$, we shall assume that $[z, w] \subseteq [x, y]$. A *centre* of three points $x, y, z \in X \cup \partial X$ is a point, $w \in X$ such that $\rho(w, [x, y]) \simeq 0$, $\rho(w, [y, z]) \simeq 0$ and $\rho(w, [z, x]) \simeq 0$. Such a point exists provided no two of x, y, z are equal to the same boundary point. It is well defined up to a bounded distance. Given subsets $P, Q \subseteq X$, we shall write $P \simeq Q$ to mean that $P \cap \partial X = Q \cap \partial X$ and $\text{hd}(P \cap X, Q \cap X) \simeq 0$.

We note that the property of hyperbolicity is quasi-isometry invariant among path-metric spaces. Clearly a quasi-isometry will map geodesics onto quasigeodesics. An important property of hyperbolicity is that quasigeodesics remain close to geodesics. More precisely if α and β are two quasigeodesics with the same endpoints, then $\alpha \simeq \beta$, where the additive constant depends only on the parameters. Note that it follows that a quasi-isometry respects centres up to bounded distance.

Suppose $Q \subseteq X \cup \partial X$. We write $\partial_X Q = \bar{Q} \cap \partial X$, where \bar{Q} is the closure of Q in $X \cup \partial X$. (In many cases, $\partial_X Q$ can be thought of as a Gromov boundary defined intrinsically to Q , and we will abbreviate $\partial_X Q$ to $\partial_X Q$.) Given $r \geq 0$ we write $N(Q, r) = Q \cup N(Q \cap X, r)$. Let $\text{join}(Q) = \bigcup \{[x, y] \mid x, y \in Q\}$. Thus, $\partial \text{join}(Q) = \partial_X Q$. We say that Q is *r-quasiconvex* if $\text{join}(Q) \subseteq N(Q, r)$. For any Q , the set $\text{join}(Q)$ is uniformly quasiconvex. If Q is *r-quasiconvex*, then the induced path metric on $N(Q \cap X, r)$ agrees with ρ up to an additive constant. Thus, $N(Q, r)$ is intrinsically hyperbolic. Moreover, $\partial_X Q = \partial N(Q, r)$ can be identified with the boundary of $N(Q, r)$. If $P \subseteq X \cup \partial X$ is another subset with $P \simeq Q$, then P is also quasiconvex, and $\partial_X P = \partial_X Q$. If $R \subseteq X \cup \partial X$ is a quasiconvex subset with $\rho(Q \cap X, R \cap X) \simeq 0$, then $Q \cup R$ is quasiconvex, and $\partial_X(Q \cup R) = \partial_X Q \cup \partial_X R$. Note that this applies if $Q \cap R \neq \emptyset$. Note also that if $Q \subseteq X \cup \partial X$, with $Q \cap X$ dense in Q and $Q \cap X$ is quasiconvex, the Q is quasiconvex (possibly after increasing the constant).

Suppose that $Q \subseteq X$ is quasiconvex and $x \in X \cup (\partial X \setminus \partial_X Q)$. We write $\pi_Q(x)$ for a point of Q with the property that $\rho(\pi_Q(x), [x, y]) \simeq 0$ for all $y \in Q$. Such a point exists and is well defined up to bounded distance. If $x \in X$, then it can be characterised by the property that $\rho(x, \pi_Q(x)) \simeq \rho(x, Q)$. After moving the image a bounded distance, it can always be assumed that $\pi_Q(x) = x$ for all $x \in Q$, and that $\pi_Q(y) = \pi_Q(x)$ whenever $y \in [x, \pi_Q(x)]$. We shall write $[x, Q] = [x, \pi_Q(x)]$. We shall refer to the map π_Q as a *projection* of $X \cup (\partial X \setminus \partial_X Q)$ onto

Q . Note that it need not be continuous. If $x \in \partial_X Q$, we shall take $\pi_Q(x) = x$, thereby giving us a map $\pi_Q : X \cup \partial X \longrightarrow Q \cup \partial_X Q$.

We note:

Lemma 4.1. *For all $x, y \in X$, $\rho(\pi_Q(x), \pi_Q(y)) \preceq \rho(x, y)$.*

Here, as usual, the additive constant is “uniform”, i.e. depends only on the parameters (of hyperbolicity and quasiconvexity).

We have the following converse:

Lemma 4.2. *Suppose $Q \subseteq X$ and $\psi : X \longrightarrow Q$ is a map which restricts to inclusion on Q and with the property that $\rho(\psi(x), \psi(y))$ is bounded above by a fixed linear function of $\rho(x, y)$. Then Q is quasiconvex.*

Proof. Let $r \geq 0$ and suppose $x, y \in N(Q, r)$. We choose points $x = x_0, x_1, \dots, x_n = y$ in $x_i \in [x, y]$ so that $\rho(x_i, x_{i+1}) \simeq 0$ for all i and where n is bounded above by a linear function of $\rho(x, y)$. Let $\alpha[x, \pi_Q(x)] \cup [\pi_Q(x_0), \pi_Q(x_1)] \cup \dots \cup [\pi_Q(y), y]$. If r is sufficiently large in relation to the parameters, then $\alpha \subseteq N(Q, r)$. We deduce that any pair of points, $x, y \in N(Q, r)$ are connected by a path in $N(Q, r)$ whose length is bounded by a linear function of $\rho(x, y)$.

Now, let β be a path in $N(Q, r)$ connecting x to y of minimal length (up to an additive constant). Again, $\text{length}(\beta)$ is bounded above by the a linear function of the distance between its endpoints. Since the same argument applies to any subpath of β , we deduce that β is quasi-geodesic. Thus, $[x, y]$ remains a bounded distance from β and hence from Q . Since, in particular, this applies to $x, y \in Q$, we deduce that Q is quasiconvex. \square

(Note that this is stated in [Mit2], though the argument presented there is incomplete.)

Suppose that $Q \subseteq X \cup \partial X$ is quasiconvex and $x \in X \cup \partial X$. Recall that $\pi_Q(x)$ is characterised (up to bounded distance) by the property that $\rho(\pi_Q(x), [x, y]) \simeq 0$ for all $y \in Q$. If $P \subseteq X \cup \partial X$ with $P \simeq Q$, then we see that $\pi_P(x) \simeq \pi_Q(x)$ for all $x \in X \cup (\partial X \setminus Q)$. Since quasigeodesics remain a bounded distance from geodesics, we similarly deduce:

Lemma 4.3. *Suppose $\phi : X \longrightarrow Y$ is a quasi-isometry of hyperbolic spaces and $Q \subseteq X$ is quasiconvex. Then, $\phi(Q)$ is quasiconvex. Moreover, $\phi(\pi_Q(x)) \simeq \pi_{\phi(Q)}(\phi(x))$ for all $x \in X \cup (\partial X \setminus Q)$.*

(In writing $\phi : X \longrightarrow Y$ as though it were a function, we are implicitly making a choice of $\phi(x) \sim_\phi x$, as discussed in Section 3.) Here we

are using the fact that any quasi-isometry extends to a homeomorphism of boundaries.

Note that if $y, z \in X \cup \partial X$, then $[y, z]$ is quasiconvex, and $\pi_{[y,z]}$ is a centre of the triple x, y, z . If $w \in \pi_{[y,z]}^{-1}[y, x]$, then $[w, [y, z]] = [w, \pi_{[y,z]}(w)] \subseteq \pi_{[y,z]}^{-1}[y, x]$. We note:

Lemma 4.4. $\pi_{[y,z]}^{-1}[y, x]$ is quasiconvex.

Proof. This follows easily from the fact that any pair of points, $a, b \in \pi_{[y,z]}^{-1}[y, x]$ are connected by a broken geodesic with a bounded number of segments, namely $[a, \pi_{[y,z]}(a)] \cup [\pi_{[y,z]}(a), \pi_{[y,z]}(b)] \cup [\pi_{[y,z]}(b), b]$. \square

Note also that if $y \in \partial X$ (or if $\rho(x, y) \gg 0$), then y lies in the interior of $\pi_{[y,z]}^{-1}[y, x]$. We deduce:

Lemma 4.5. If $x \in \partial X$, then x has a countable decreasing base of uniformly quasiconvex neighbourhoods.

Proof. Let $(\alpha_n)_{n \in \mathbb{N}}$ be a decreasing sequence of rays tending to x . The sets $(\pi_{\alpha_0}^{-1}(\alpha_n))_n$ form a neighbourhood base. \square

Let $(P_n)_{n \in \mathbb{N}}$ be a sequence of non-empty subsets of $X \cup \partial X$. We say that $(P_n)_n$ is *escaping* if $P_m \subseteq P_n$ for $m \leq n$ and if $\rho(x, P_n \cap X) \rightarrow \infty$ for some (hence all) $x \in X$. (Note that if X is proper, and each P_n is closed, it is enough that $\bigcap_n P_n \subseteq \partial X$.)

Lemma 4.6. Suppose that $(P_n)_n$ is an escaping sequence of closed uniformly quasiconvex subsets of $X \cup \partial X$. Then $\bigcap_n P_n = \{x\}$ for some $x \in \partial X$. Moreover, if U is any neighbourhood of x in $X \cup \partial X$, then $P_n \subseteq U$ for all sufficiently large n .

Proof. Fix any $a \in X$. If $m \leq n$, then $[a, P_m]$ is uniformly close to $[a, P_n]$. Thus, there is a geodesic ray, $\alpha = [a, x]$ such that $[a, P_n]$ is uniformly close to α for all n . Note that $\pi_{P_n}(a) \in P_n$ converges to x , and so $x \in \bigcap_n P_n$. Let U be a neighbourhood of x in $X \cup \partial X$. By (the construction of) Lemma 4.5, there is some $b \in [x, a]$ such that $\pi_{\alpha}^{-1}[b, x] \subseteq U$. Suppose for contradiction that there is some $x_n \in P_n \setminus \pi_{\alpha}^{-1}[b, x]$ for infinitely many n . Then b is a bounded distance from $[x_n, x]$ and hence P_n , contradicting the assumption that $(P_n)_n$ is escaping. \square

We shall need the following observation regarding divergence of (quasi)geodesics.

Suppose $x, y \in X$. Let $\rho(x, y) = 2r$, and let m be the midpoint of $[x, y]$. Suppose $z, w \in X$ are such that $\max\{\rho(x, y), \rho(z, w)\} \leq r$ then $\rho(m, [z, w]) \simeq 0$. In particular, if $\rho(m, [z, w]) \gg 0$, then $\rho(x, z) + \rho(y, w) \geq r$. (This follows easily by considering an approximating

tree for the four points x, y, z, w .) From the fact that quasigeodesics remain a bounded distance from geodesics, we have a similar statement for quasigeodesic segments. In this case, $2r$ is the length of the domain interval, m is the “midpoint” as defined by the parameterisation, and we deduce that $\rho(x, z) + \rho(y, w)$ is bounded below by an increasing linear function of r . More precisely:

Lemma 4.7. *Suppose $r \geq 0$ and $\alpha, \beta : [-r, r] \rightarrow X$ are quasigeodesic segments. Suppose $\rho(\alpha(0), \beta(0)) \geq l$, where l is a constant sufficiently large depending on the parameters (of hyperbolicity and quasigeodesicity). Then, $\rho(\alpha(t), \beta(t)) + \rho(\alpha(-t), \beta(-t))$ is bounded below by an increasing linear function of t which depends only on the parameters.*

In particular, provided the domain, $[-r, r]$ is sufficiently large, we can arrange that $\rho(\alpha(t), \beta(t)) + \rho(\alpha(-t), \beta(-t))$ is arbitrarily large.

Definition. We shall say that a hyperbolic space is *taut* if every point is (or is a bounded distance from) the centre of three ideal points.

The following is an immediated consequence of the fact that quasigeodesics remain a bounded distance from geodesics, so that quasi-isometries map centres a bounded distance from centres.

Lemma 4.8. *Suppose that $\phi, \psi : X \rightarrow Y$ are quasi-isometries whose extensions to ∂X are identical. Suppose that X (and hence also Y) is taut. Then ϕ and ψ are equivalent (as defined in Section 3).*

In particular, we see that $\rho(\phi(x), \psi(x))$ is bounded for all $x \in X$.

5. QUASICONVEX FUNCTIONS

In this section, we describe a notion of “quasiconvexity” for positive functions. This gives a convenient way of expressing the Bestvina-Feighn flaring condition (Section 6). Let $I \subseteq \mathbb{R}$ be an interval.

Let $f : \mathbb{R} \rightarrow [0, \infty)$.

Definition. We say that f has property (Q1) if there exist constants $t > 0$, $k > 0$ and $0 < \lambda < 1/2$, such that for all $x \in \mathbb{R}$, we have

$$f(x) \leq \lambda(f(x-t) + f(x+t)) + k.$$

We note that $f(x) \leq \lambda'(f(x-2t) + f(x+2t)) + k'$, where $\lambda' = \frac{\lambda^2}{1-2\lambda^2} < \lambda$ and $k' = \left(\frac{1+2\lambda}{1-2\lambda^2}\right)k$. By iterating this process, we see that, at the cost of increasing t and k , we can assume that λ is arbitrarily small.

Let μ be any constant greater than $\frac{1}{1-2\lambda}$. If $f(x) \geq \mu k$, then $f(x-t) + f(x+t) > \frac{1}{\lambda}(f(x) - k) \geq \frac{1}{\lambda} \left(1 - \frac{1}{\mu}\right) f(x) \geq 2f(x)$. In particular, $\max\{f(x-t), f(x+t)\} > sf(x)$, where $s = \frac{1}{2\lambda} \left(1 - \frac{1}{\mu}\right) > 1$.

Moreover, if $\mu k \leq f(x) \leq f(x+t)$, then $f(x+2t) \geq \frac{1}{\lambda}(f(x+t) - k) - f(x) \geq \left(\frac{1}{\lambda} - 1\right) f(x+t) - \frac{k}{\lambda} \geq \left(\frac{1}{\lambda} - 1 - \frac{1}{\lambda\mu}\right) f(x+t) \geq rf(x+t)$, where $r = \frac{\mu - \lambda\mu - 1}{\lambda} > 1$. We see, inductively, that $f(x+nt) \geq r^{n-1}f(x)$ for all $n \in \mathbb{N}$. Similarly, if $\mu k \leq f(x) \leq f(x-t)$, then $f(x-nt) \geq r^{n-1}f(x)$.

Putting the above two observations together, we see that for any $x \in \mathbb{R}$ and any $h \geq \mu k$, the set of $n \in \mathbb{Z}$ for which $x+nt \in I$ and $f(x+nt) \leq h$ form a (possibly empty) sequence of consecutive integers.

To be able to interpolate between these values, we shall need to add another hypothesis.

Definition. We shall say that f has property (Q2) if there are constants $t > 0$, $a, b \geq 0$, such that for all $x, y \in I$ with $|x-y| \leq t$, then $f(y) \leq af(x) + b$.

We can assume (at the cost of increasing a and b) that t is the same constant as that appearing in (Q1). Clearly, (Q2) implies that the growth of f is at most exponential.

Definition. We shall say that a function $f : \mathbb{R} \rightarrow [0, \infty)$ is *quasiconvex* if it satisfies (Q1) and (Q2).

For the remainder of this section, we suppose that $I = \mathbb{R}$ (though the statements hold for $I = \mathbb{R}$).

Moreover, as $n \rightarrow \pm\infty$, the sequence $f(x+nt)$ is either bounded, or grows exponentially, with an exponent that depends only on λ and k .

From this and the earlier observations, we see:

Lemma 5.1. *Suppose f is quasiconvex. There are constants $\zeta > 0$, $B \geq 0$ and $l \geq 0$ depending only on the parameters such that the following hold. Either there is some $k_1 > 0$ such that $f(x) \geq k_1 e^{\zeta x}$ for all x , or else $f(x)$ is bounded as $x \rightarrow \infty$. Similarly, either there is some $k_2 > 0$ such that $f(-x) \geq k_2 e^{\zeta x}$ for all x , or else $f(-x)$ is bounded as $x \rightarrow \infty$. Moreover, there is a (possibly empty) subinterval $I \subseteq \mathbb{R}$ such that $f(x) \leq B$ for all $x \in I$, and $f(x) > B$ for all $x \in \mathbb{R}$ a distance at least l from I .*

Note that it follows that if $f(x)$ is bounded as $x \rightarrow \infty$ then $f(x) \leq B$ for all sufficiently large x . The same statement holds for $x \rightarrow -\infty$.

Moreover if f is bounded (in both directions), then $f(x) \leq B$ for all $x \in \mathbb{R}$.

The following is an elementary observation:

Lemma 5.2. *Let \mathcal{F} be a family of uniformly quasiconvex functions. Define $f : \mathbb{R} \rightarrow [0, \infty]$ by $f(x) = \sup\{g(x) \mid g \in \mathcal{F}\}$. Then either $f(x) = \infty$ for all x , or else f is a uniformly quasiconvex function to $[0, \infty)$.*

We also note that any positive linear combination of uniformly quasiconvex functions is uniformly quasiconvex.

The property of quasiconvexity is quite robust. In particular:

Lemma 5.3. *Suppose that $h : [0, \infty) \rightarrow [0, \infty)$ is any non-decreasing function. Suppose that $f : \mathbb{R} \rightarrow [0, \infty)$ has the property that for any $x \in \mathbb{R}$, there is a uniformly quasiconvex function $g : \mathbb{R} \rightarrow [0, \infty)$ such that $|g(x) - f(x)| \leq h(0)$ and $g(y) \leq f(y) + h(|x - y|)$ for all $y \in \mathbb{R}$. Then f satisfies (Q1).*

Proof. To simplify notation, we suppose that $x = 0$. For large enough t , we have $g(0) \leq \frac{1}{9}(g(t) + g(-t)) + k$, where k is some constant. Let $c = h(t)$. We have $|f(y) - g(y)| \leq c$ for all $y \in [-t, t]$. Suppose that $f(0) \geq l = \max\{2c, c + k\}$. Since $|f(0) - g(0)| \leq h(0) \leq c$ and $f(0) \geq c + k$, we have $g(0) \geq k$. Since $|f(0) - g(0)| \leq f(0)/2$, we have $f(0) \leq 2g(0)$. We deduce that $g(t) + g(-t) \geq 8g(0)$ and so $f(t) + f(-t) \geq 8g(0) - 2c \geq 4f(0) - 2c \geq 3f(0)$. We see that, whatever the value of $f(0)$, we have $f(0) \leq \frac{1}{3}(f(t) + f(-t)) + l$, where l is a fixed constant. This proves (Q1). \square

Lemma 5.4. *Suppose that $h : [0, \infty) \rightarrow [0, \infty)$ is any non-decreasing function. Suppose that $f : \mathbb{R} \rightarrow [0, \infty)$ has the property that for any $x \in \mathbb{R}$, there is a uniformly quasiconvex function $g : \mathbb{R} \rightarrow [0, \infty)$ such that $|g(y) - f(y)| \leq h(|x - y|)$ for all $y \in \mathbb{R}$. Then f is quasiconvex.*

Proof. By Lemma 5.3, f satisfies (Q1). To verify (Q2), note that if $|y| \leq t$, then $f(y) \leq g(y) + c \leq (ag(0) + b) + c \leq af(0) + ac + b + c$. \square

Finally, suppose that $f : \mathbb{R} \rightarrow [0, \infty)$ has the property that multiples $nf : \mathbb{R} \rightarrow [0, \infty)$ are uniformly quasiconvex for $n \in \mathbb{N}$. Then we see that f is quasiconvex with the same multiplicative constants and additive constants equal to 0.

In this case, we note that if $f(x + t) \geq f(x)$, then $f(x + 2t) \geq \frac{1-\lambda}{\lambda}f(x + t)$. Thus, as similar analysis as above shows that if f is bounded as $x \rightarrow \infty$, then in fact, it decays exponentially. Moreover, we note:

Lemma 5.5. *If f is quasiconvex with additive constants 0, then there is some $\zeta > 0$ and $k > 0$ such that if $x, u \in \mathbb{R}$, then $f(x) \leq ke^{-\zeta u}(f(x+u) + f(x-u))$.*

6. STACKS

In this section, we shall describe more precisely the notion of a “stack”, referred to in Sections 1 and 2. This can be thought of as a geodesic space sliced up by a sequence of subspaces termed “sheets”. If we assume that the sheets are uniformly hyperbolic in the induced path-metrics, then hyperbolicity of the stack turns out to be equivalent to a quasiconvexity condition on the distance functions — see Theorems 6.2 and 6.4. This is essentially the same as the “flaring condition” of [BeF]. Eventually, we will want to restrict attention to “hyperbolic stacks”, i.e. those satisfying all of the conditions (S1)–(S6) below, though we shall only introduce these hypotheses as we need them.

For simplicity, we shall deal for the moment only with “bi-infinite stacks”, i.e. those indexed by the integers, \mathbb{Z} . However, everything we say in this section applies equally well to stacks indexed by an arbitrary set of consecutive integers. Of particular interest, of course, will be “semi-infinite stacks”, i.e. those indexed by the natural numbers \mathbb{N} . We shall begin by giving a fairly general formulation of the notion of a stack, but we shall see how, without loss of generality, we can make a number of simplifying assumptions.

Let (Ξ, ρ) be a geodesic space, and let $((X_i, \rho_i))_{i \in \mathbb{Z}}$ be a sequence of geodesic spaces. Let $f_i : X_i \rightarrow \Xi$ be a sequence of maps (not necessarily continuous). Recall the definition of “straight” from Section 3. We assume

(S1₀) The maps f_i are uniformly straight, and

(S2₀) $\rho(f_i(X_i), f_j(X_j))$ is bounded below by an increasing linear function of $|i - j|$.

Now, the constructions described at the end of Section 3 allow us to allow us to replace the maps f_i with locally isometric embeddings (except that we are now dealing with a sequence of maps rather than just two). The new images will be uniformly separated. Moreover, the second condition ensures that the original images were locally finite, and so (by using nets) the construction can be assumed, if necessary, to preserve properness. Thus, (at the cost of modifying Ξ by a quasi-isometry) we can replace each $f_i : X_i \rightarrow \Xi$ by a subspace $X_i \subseteq \Xi$,

and assume:

(S1) With respect to the induced path-metric, ρ_i , each of the spaces X_i is uniformly straight. Moreover, $\rho(X_i, X_j)$ is bounded away from 0 for $i \neq j$.

(S2) $\rho(X_i, X_j)$ is bounded below by an increasing linear function of $|i - j|$.

(The remaining axioms (S3)–(S6) we introduce can be readily reinterpreted in the more general set-up above.)

We remark that (S2) is implied by the assumption that each X_i “coarsely separates” Ξ in the following sense:

(S2') If $i < k < j$, then any path connecting X_i to X_j in Ξ passes within a distance r of X_k , where r is strictly less than half the separation constant appearing in (S1).

To see this, suppose that α is such a path. Let $x_k \in \alpha$ be a point with $\rho(x_k, X_k) \leq r$. The conditions imply that $\rho(x_k, x_l)$ is bounded below for distinct $k, l \in \{i, i + 1, \dots, j\}$. It follows that $\text{length}(\alpha)$ is bounded below by an increasing linear function of $|i - j|$.

Now, in addition to (S1) and (S2), we assume:

(S3) $\text{hd}(X_i, X_{i+1})$ is bounded above for $i \in \mathbb{Z}$.

Given $r \geq 0$ and $i \in \mathbb{Z}$, we define a relation between X_i and X_{i+1} by writing $x \sim_{r,i} y$ if $\rho(x, y) \leq r$. Provided that r is sufficiently large in relation to the parameters, this defines a quasi-isometry between (X_i, ρ_i) and (X_{i+1}, ρ_{i+1}) . We refer to a quasi-isometry arising in this way, or which agrees with such a quasi-isometry up to a bounded distance, as a *natural quasi-isometry*. We can consider the union of all the relations, $\sim_{r,i}$ for $i \in \mathbb{Z}$ as a relation on $\bigcup_{i \in \mathbb{Z}} X_i$. We can construct a space $\bigvee_i X_i = \dots \vee_\phi X_{-1} \vee_\phi X_0 \vee_\phi X_1 \vee_\phi \dots$, exactly as we did for a pair of spaces in Section 3. Moreover, we get a uniformly proper map $\bigvee_i X_i \rightarrow \Xi$. If necessary, we can assume that the construction preserves properness. We also note, more generally that if $i, j \in I$ we can construct a quasi-isometry between X_i and X_j , for an example by taking an iterate of ϕ as already defined, or by redefining ϕ by choosing r sufficiently large in relation to $|i - j|$. Such a quasi-isometry will be “natural” in the following sense:

Definition. We say that a quasi-isometry, ϕ , between the sheets (X_i, ρ_i) and (X_j, ρ_j) is *natural* if there is some $r \geq 0$ such that $\rho(x, y) \leq r$ for all $x \in X_i$ and $y \in X_j$ with $x \sim_\phi y$.

Note that r can always be chosen less than some (linear) function of $|i - j|$, so that the parameters of such a quasi-isometry are bounded in terms of $|i - j|$.

Definition. An r -chain, $(x_i)_{i \in \mathbb{Z}}$ is a sequence of points, $x_i \in X_i$ such that $\rho(x_i, x_{i+1}) \leq r$ for all i .

We shall sometimes write $\underline{x} = (x_i)_i$. If r is sufficiently large (in relation to the parameters), we can assume that every point of $\bigcup_i X_i$ lies in an r -chain. Frequently, we shall make this assumption, and refer simply to “chains”.

Note that chains are uniformly quasigeodesic, or, more precisely, can be extended to uniform quasigeodesics on Ξ , by interpolating with geodesic segments of bounded length.

Let us now assume:

(S4) The spaces (X_i, ρ_i) are uniformly hyperbolic geodesic spaces.

Given $x_i, y_i \in X_i$, write $[x_i, y_i]_i$ for some choice of geodesic from x_i to y_i in (X_i, ρ_i) . We note:

Lemma 6.1. *Suppose $(x_i)_i$ and $(y_i)_i$ are chains. Then, for all i , $\text{hd}([x_i, y_i]_i, [x_{i+1}, y_{i+1}]_{i+1})$ is uniformly bounded.*

Proof. This follows from the fact that the image of $[x_i, y_i]_i$ under the quasi-isometry, ϕ , is uniformly quasigeodesic, and hence a bounded distance from $[x_{i+1}, y_{i+1}]_{i+1}$. \square

Thus, if $z_j \in [x_j, y_j]_j$ for some j , then we can extend z_j to a chain, $(z_i)_i$, with $z_i \in [x_i, y_i]_i$ for all i . (This may entail increasing the chain constant by a controlled amount.) Indeed, if $z_j^0, z_j^1, \dots, z_j^n \in [x_j, y_j]_j$ are a sequence of points appearing in this order, then we can extend them to chains $(z_i^0)_i, (z_i^1)_i, \dots, (z_i^n)_i$ such that $z_i^0, z_i^1, \dots, z_i^n$ appear in the same order on $[x_i, y_i]_i$ for all i .

Let us now assume:

(S5) (Ξ, ρ) is hyperbolic.

Theorem 6.2. *Suppose that $(x_i)_i$ and $(y_i)_i$ are chains. The map $[i \mapsto \rho_i(x_i, y_i)]$ is uniformly quasiconvex.*

Here “uniformity means that the constants involved in quasiconvexity depend on the parameters featuring in conditions (S1)–(S5) and on the chain constant, but not on the chains themselves.

Proof. First note that property (Q2) is an immediate consequence of the fact that ϕ is a quasi-isometry. To simplify notation, we shall verify property (Q1) for $i = 0$. In particular, we shall find a constant $L \geq 0$ and some fixed $i \geq 0$, such that $\rho_i(x_i, y_i) + \rho_{-i}(x_{-i}, y_{-i}) \geq 3\rho_0(x_0, y_0) - L$. To this end, we fix a constant $h \geq 0$ to be chosen as described below, and take points $x_0 = z_0^0, z_0^1, \dots, z_0^n, z_0^{n+1} = y_0$ on $[x_0, y_0]$ with $\rho_0(z_0^0, z_0^m) = mh$ for $m \leq n$, so that $\rho_0(z_0^n, z_0^{n+1}) \leq h$. As described above, we can extend z_0^m to chains $(z_i^m)_i$, where $z_i^m \in [x_i, y_i]_i$, $z_i^0 = x_i, z_i^{n+1} = y_i$, and $z_i^0, z_i^1, \dots, z_i^{n+1}$ appear in the correct order along $[x_i, y_i]_i$. The chain constant involved is independent of x_i, y_i or h . Now the chain $(z_i^m)_i$ lies in a uniform quasigeodesic path β^m . Let l be the associated constant arising from Lemma 4.7, given the hyperbolicity of Ξ . Choosing h sufficiently large in relation to l , we can assume that $\rho(z_0^m, \beta^m) \geq l$ for all $m \in \{0, \dots, n-1\}$. Thus, by Lemma 4.7, the map $[i \mapsto \rho(z_i^m, z_i^{m+1}) + \rho(z_{-i}^m, z_{-i}^{m+1})] : \mathbb{Z} \rightarrow [0, \infty)$ is bounded below by a fixed linear function of i . Now, $\rho_i \geq \rho$ for all i . Thus, for i sufficiently large, we have $\rho_i(z_i^m, z_i^{m+1}) + \rho_{-i}(z_{-i}^m, z_{-i}^{m+1}) \geq L = 3h$. Summing over $m \in \{0, \dots, n-1\}$, we see that $\rho_i(x_i, y_i) + \rho_{-i}(x_{-i}, y_{-i}) \geq nL \geq \left(\frac{\rho_0(x_0, y_0)}{n} - 1\right)L \geq (L/h)\rho_0(x_0, y_0) - L \geq 3\rho_0(x_0, y_0) - L$. This proves (Q1) as required. \square

Let us now assume, in addition that:

(S6) $\bigcup_{i \in \mathbb{Z}} X_i$ is quasidense in Ξ .

In this case, the image of the proper map $f : \bigvee_i X_i \rightarrow \Xi$ is quasidense, and so f must be a quasi-isometry. Thus, up to uniform quasi-isometry, we can always assume that Ξ is obtained in this manner. Though not essential, this will be convenient for certain constructions.

Indeed, we can construct a stack abstractly from a sequence of geodesic spaces, $(X_i)_i$, together with a sequence of uniform quasismetries, between successive spaces. Again, if each of the X_i is proper, we can assume that the resulting stack, $\Xi = \bigvee_i X_i$, is proper.

In this construction, properties (S1), (S2), (S3) and (S6) necessarily hold. Only the straightness of each of the spaces X_i (as required by property (S1)) calls for comment. In fact, we can formulate the argument as follows. Let $Y_i = X_i \vee_\phi X_{i+1} \subseteq \Xi$ be the “strip” lying between X_i and X_{i+1} , and let σ_i be the induced path-metric on Y_i . We note that

$\sigma_i(X_i, X_i)$ is bounded below, and that the inclusions of (X_i, ρ_i) and of (X_{i+1}, ρ_{i+1}) in (Y_i, σ_i) are both uniform quasi-isometries. Moreover, we note that $Y_{i-1} \cap Y_i = X_i$ and that $\Xi = \bigcup_{i \in \mathbb{Z}} Y_i$. With this set-up, we can deduce:

Lemma 6.3. *If $x, y \in X_0$, then $\rho_0(x, y)$ is bounded above by an exponential function of $\rho(x, y)$, whose exponent depends only on the parameters.*

Proof. It is easily seen that any path connecting x to y in Ξ can be approximated by a broken geodesic with break points $x = x^0, x^1, \dots, x^n = y$, where n is bounded above by a linear function of $\text{length}(\alpha)$ and where for each m , there is some $i = i(m)$ such that $x^m, x^{m+1} \in X_i \cup X_{i+1}$ and $\sigma_i(x^m, x^{m+1}) \simeq 0$. Note that the maximal value of $|i|$ is bounded by n . Let $y^m \in X_0$ be the basepoint of a chain that includes x^m . We can assume that $x^0 = y^0$ and $x^n = y^n$. We see that $\rho_0(y^m, y^{m+1})$ is bounded above by an exponential function of $i(m)$ and hence of n . Summing over all $m \in \{0, \dots, n\}$, we see that $\rho_0(x, y)$ is also bounded by an exponential function of n and hence, in turn, of $\text{length}(\alpha)$. By taking α to be a geodesic from x to y with respect to the metric ρ we get the desired result. \square

To conclude this section, we note that we have the following converse to Lemma 6.2. Suppose that Ξ is a stack satisfying axioms (S1)–(S4) and (S6). We choose r large enough in relation to the parameters so that the union of all r -chains is quasidense in Ξ .

Theorem 6.4. *Suppose for any pair of r -chains, $(x_i)_i$ and $(y_i)_i$, the function $[i \mapsto \rho_i(x_i, y_i)] : \mathbb{Z} \rightarrow [0, \infty)$ is uniformly quasiconvex. Then Ξ is uniformly hyperbolic.*

Proof. This follows from (a slight reinterpretation of) the main result of [BeF]. \square

In fact, we may weaken the hypotheses of Theorem 6.4 as follows:

Corollary 6.5. *Suppose we have a family $\{(x_i^p)_i \mid p \in P\}$ of r -chains in Ξ indexed by some set P . Suppose that the sets $\{x_i^p \mid p \in P\}$ are uniformly quasidense in X_i for all $i \in \mathbb{Z}$. Suppose also that for all $p, q \in P$, the function $[i \mapsto \rho_i(x_i^p, x_i^q)] : \mathbb{Z} \rightarrow [0, \infty)$ is uniformly quasiconvex. Then Ξ is hyperbolic.*

Proof. Suppose that $(a_i)_i$ is any chain in Ξ , and that $j \in \mathbb{Z}$. There is some $p \in P$ so that $\rho_j(a_j, x_j^p)$ is uniformly bounded. It follows that $\rho_i(a_i, x_i^p) \leq h(|i - j|)$ for some fixed non-decreasing function $f : \mathbb{N} \rightarrow \mathbb{R}$.

Now, if $(b_i)_i$ is another chain, we can similarly find q with $\rho_j(b_j, x_j^q)$ uniformly bounded. It follows that $|\rho_i(a_i, b_i) - \rho_i(x_i^p, x_i^q)| \leq 2h(|i - j|)$. Since the function $[i \mapsto \rho_i(x_i^p, x_i^q)]$ is by hypothesis uniformly quasiconvex, it follows by Lemma 5.3 that $[i \mapsto \rho_i(a_i, b_i)]$ is also. Since $(a_i)_i$ and $(b_i)_i$ are arbitrary r -chains, we can apply Theorem 6.4 to deduce that Ξ is hyperbolic. \square

We shall be applying Corollary 6.5 to the case where we have a sheet-preserving action on Ξ so the the quotients of the sheets have bounded diameter, and where we take the orbit of a given r -chain.

We finish off with a corollary of Theorem 6.2 that will be used in Section 15. Suppose $\psi : \Xi \rightarrow \Xi$ is sheet preserving isometry of the stack Ξ . Thus, $\psi|_{X_i}$ is an isometry of the sheet (X_i, ρ_i) . We write $\text{sl}_i(\psi)$ for its stable length (as defined at the end of Section 3).

Lemma 6.6. *If $\psi : \Xi \rightarrow \Xi$ is a sheet-preserving isometry. Then the map $[i \mapsto \text{sl}_i(\psi)] : \mathbb{Z} \rightarrow [0, \infty)$ is uniformly quasiconvex with additive constant 0.*

Proof. Choose any $x \in X$. By Theorem 6.2 the functions $[i \mapsto \rho_i(x, \psi^n(x))]$ are uniformly quasiconvex for $n \in \mathbb{N}$, and so therefore are $[i \mapsto \frac{1}{n}\rho_i(x, \psi^n(x))]$. The same therefore applies to the limit $\text{sl}_i(\psi)$. Since $\text{sl}_i(\psi^n) = n \text{sl}_i(\psi)$, and the functions $[i \mapsto \text{sl}_i(\psi^n)]$ are uniformly quasiconvex, it follows that the additive constant must be 0. \square

The above discussion applies equally well where we replace the indexing set, \mathbb{Z} by any set, $\mathcal{I} \subseteq \mathbb{Z}$, of consecutive integers, in particular for \mathbb{N} . Note that, starting from a stack Ξ indexed by \mathbb{Z} , and $\mathcal{I} \subseteq \mathbb{Z}$, we can view the metric r -neighbourhood, $\Xi(\mathcal{I})$, of $\bigcup_{i \in \mathcal{I}} X_i$ as a stack, with sheets $(X_i)_{i \in \mathcal{I}}$, provided we choose r large enough in relation to the parameters. In particular, setting $\Xi^\pm = \Xi(\pm\mathbb{N})$, we have $\Xi^+ \cup \Xi^- = \Xi$ and $\Xi^+ \cap \Xi^-$ is a uniform neighbourhood of X_0 . (In fact, in the construction of Ξ described earlier, we would have $\Xi^+ \cap \Xi^- = X_0$.)

Proposition 6.7. *Suppose that Ξ is a stack with uniformly hyperbolic sheets $(X_i)_{\mathbb{Z}}$. If Ξ is hyperbolic, then so is $\Xi(\mathcal{I})$, where $\mathcal{I} \subseteq \mathbb{Z}$ is any set of consecutive integers. Moreover, the hyperbolicity constant of $\Xi(\mathcal{I})$ depends only on the parameters and hyperbolicity constant of Ξ .*

Proof. We first apply Theorem 6.2 to see that chains in Ξ are uniformly quasiconvex. Restricting to \mathcal{I} to see that that the same holds for $\Xi(\mathcal{I})$. Therefore Theorem 6.4 tells us that Ξ is hyperbolic, as required. \square

7. THE CANNON-THURSTON MAP

In this section, we give a construction of quasiconvex sets in a hyperbolic stack, and explain how this may be used to construct the “Cannon-Thurston” map from the boundary of a sheet to the boundary of a stack. This was defined originally in the context of cyclic covers of 3-manifolds fibring over the circle [CannT], and generalised in [Min2]. The argument follows a similar strategy to that in [Mit1, Mit2]. For further discussion and generalisations, see [AIDP], and the references in [Bow3]. In this section, we shall not need to assume the properness of the stack.

Definition. We say that a stack is *hyperbolic* if it satisfies axioms (S1)–(S6) of Section 6.

Let Ξ be a hyperbolic stack, with sheets, $(X_i)_i$ indexed by $i \in \mathcal{I} \subseteq \mathbb{Z}$ as before. Let $\partial\Xi$ be its Gromov boundary. If $i, j \in \mathbb{Z}$, then there is a natural quasi-isometry between the sheets X_i and X_j . We thus get a homeomorphism between ∂X_i and ∂X_j . Since such a quasi-isometry is well defined up to bounded distance, (linear in $|i - j|$) this homeomorphism is canonical. We thus get canonical homeomorphisms between the spaces ∂X_i for all $i \in \mathbb{Z}$. We write $\partial^0\Xi$ for the topological space arising in this way. In other words, $\partial^0\Xi$ comes equipped with a canonical homeomorphism to ∂X_i for all i . One of the main objectives of this section will be to define the “Cannon-Thurston” map $\omega : \partial^0\Xi \rightarrow \partial\Xi$. We begin with a construction of quasiconvex sets in Ξ .

Suppose that $Q \subseteq \Xi \cup \partial\Xi$ is non-empty. Let $Q_i = Q \cap X_i$. We say that Q is *stratified* if:

- (T1) $Q \cap \Xi$ is (topologically) dense in Q ,
- (T2) $\text{hd}(Q_i, Q_{i+1}) \simeq 0$ for all $i \in \mathbb{Z}$, and
- (T3) $\bigcup_{i \in \mathbb{Z}} Q_i$ is quasidense in $Q \cap \Xi$.

Note that (T2) implies that $Q_i \neq \emptyset$ for all $i \in \mathbb{Z}$. If ϕ is a natural quasi-isometry from X_i to X_{i+1} , then $\phi(Q_i) \simeq Q_{i+1}$.

We shall say that a stratified set, Q , is *quasiconvex-stratified* if, in addition to (T1)–(T3), we have:

- (T4) The sets Q_i are uniformly quasiconvex in X_i .

Lemma 7.1. *Let Q be a quasiconvex-stratified set. Then Q is quasiconvex (as a subset of $\Xi \cup \partial\Xi$).*

Proof. By (1), it is enough to show that $Q \cap \Xi$ is quasiconvex. To this end, we define a map $\psi : \Xi \rightarrow Q \cap \Xi$ satisfying the hypotheses of Lemma 4.2. Since $\bigcup_i X_i$ and $\bigcup_i Q_i$ are quasidense in Ξ and Q respectively, it is enough to define $\psi : \bigcup_i X_i \rightarrow \bigcup_i Q_i$. Given $x \in X_i$, we set $\psi(x) = \pi_i(x)$, where $\pi_i = \pi_{Q_i} : X_i \rightarrow Q_i$ is the projection to Q_i in X_i . If $x, y \in \bigcup_i X_i$, we want to put an upper bound on $\rho(\psi(x), \psi(y))$ as a linear function of $\rho(x, y)$.

In fact, it is enough to show that $\rho(\psi(x), \psi(y))$ is uniformly bounded for $x \in X_i$ and $y \in X_j$ with $\rho(x, y) \simeq 0$ and $|i - j| \leq 1$. If $i = j$, then since X_i is straight, we also have $\rho_i(x, y) \simeq 0$. Moreover, since $\rho \leq \rho_i$, this case follows from Lemma 4.1. Otherwise, we can assume that $j = i + 1$. Let ϕ be a natural quasi-isometry between X_i and X_{i+1} . By Lemma 4.3 and the preceding discussion, we see that $\rho_{i+1}(\phi(\psi(x)), \psi(\phi(x))) \simeq 0$. Thus $\psi(x) \simeq \phi(\psi(x)) \simeq \psi(\phi(x)) \simeq \psi(y)$ as required. \square

Lemma 7.2. *Suppose that $Q \subseteq \Xi$ is a quasiconvex-stratified set, and that $x \in X_0$. Then $\rho_0(x, Q_0) \leq f(\rho(x, Q))$, where $f : [0, \infty) \rightarrow [0, \infty)$ is an increasing function depending only on the various parameters.*

Proof. Given $r \geq 0$, suppose that $\rho(x, Q) \leq r$. Then $\rho(x, y) \leq r$ for some $y \in Q_i$, where $|i|$ is bounded in terms of r . By property (T2), there is some $z \in Q_0$ with $\rho(y, z)$ bounded in terms of $|i|$ and hence r . Since X_0 is straight in Ξ , we see that $\rho(x, z)$ is bounded in terms of $\rho(x, y) \leq \rho(x, y) + \rho(y, z)$. \square

Lemma 7.3. *Suppose $R \subseteq X_0$ is quasiconvex in X_0 . Then there is a uniformly quasiconvex-stratified set $Q \subseteq \Xi$ such that $Q \cap X_0 = R$.*

Proof. By induction on i , we shall construct sets Q_j for $j \in \mathcal{I}$, with $|j| \leq i$, so that Q_j is uniformly quasiconvex, and with $\text{hd}(Q_j, Q_{j+1}) \simeq 0$. We start with $Q_0 = R$. For the inductive step, let $\phi : X_i \rightarrow X_{i+1}$ be a natural quasi-isometry. Now, $\phi(Q_i)$ is quasiconvex in X_i . If we were simply to iterate this process, we would lose control of the quasiconvexity constant. Thus, instead we set $Q_{i+1} = \text{join}(\phi(X_i))$. The quasiconvexity constant thus remains uniformly bounded. Moreover, since the quasiconvexity constant of $\phi(X_i)$ is bounded in terms of the parameters, we see that $\text{hd}_{\rho_{i+1}}(Q_{i+1}, \phi(Q_i))$ is uniformly bounded. Thus, $Q_i \simeq \phi(Q_i) \simeq Q_{i+1}$ with respect to the metric ρ . We similarly define $Q_{-(i+1)}$. Continuing inductively, we obtain a sequence $(Q_i)_{i \in \mathbb{Z}}$ of uniformly quasiconvex sets. We set $Q = \bigcup_i Q_i$. \square

Given a (uniformly) quasiconvex subset, $P \subseteq X_0 \cup \partial X_0$, let $S(P)$ be the closure, in $\Xi \cup \partial \Xi$, of the set Q constructed in Lemma 7.3, starting

with $R = P \cap X_0$. Thus, $S(P)$ is a closed uniformly quasiconvex-stratified set with $S(P) \cap X_0 = P \cap X_0$. Moreover, if $R \subseteq P$, then $S(R) \subseteq S(P)$.

We are now ready to define the Cannon-Thurston map. We identify $\partial^0 \Xi$ with ∂X_0 . Suppose $x \in \partial X_0$. Let $(P^n)_{n \in \mathbb{N}}$ be a decreasing base of uniformly quasiconvex neighbourhoods of x in $X_0 \cup \partial X_0$, as given by Lemma 4.5. Let $Q^n = S(P^n)$ be the quasiconvex-stratified set defined above. Thus, $(Q^n)_n$ is a decreasing sequence of closed uniformly quasiconvex subsets of $\Xi \cup \partial \Xi$. Since $(P^n)_n$ is escaping (in the sense of Lemma 4.6), we see by Lemma 7.2 that $(Q^n)_n$ is also escaping. By Lemma 4.6, there is some $y \in \partial \Xi$ such that $\bigcap_n Q_n = \{y\}$.

We set $\omega(x) = y$. A-priori, this might depend on the choice of the sets $(P^n)_n$. The fact that it is well defined is a consequence of the following lemma. We set $\omega(x) = x$ for all $x \in X_0$, so as to give us a map $\omega : X_0 \cup \partial X_0 \longrightarrow \Xi \cup \partial \Xi$.

Lemma 7.4. *The map $\omega : X_0 \cup \partial X_0 \longrightarrow \Xi \cup \partial \Xi$ is continuous.*

Proof. Suppose $x \in X_0$. Let P^n and $Q^n = S(P^n)$ be as in the construction that defines $\omega(x)$. Let U be a neighbourhood of $\omega(x)$ in $\Xi \cup \partial \Xi$. By Lemma 4.6, there is some m such that $Q^m \subseteq U$. Suppose that z lies in the interior of P^m in $X_0 \cup \partial X_0$. If $z \in X_0$, then $z \in Q^m \subseteq U$, so $\omega(z) = z \in U$. If $z \in \partial X_0$, let $(R^n)_n$ be a base of quasiconvex neighbourhoods of z in $X_0 \cup \partial X_0$ which defines $\omega(z)$. Now there is some n such that $R^n \subseteq Q^m$, and so $S(R^n) \subseteq S(P^m) = Q^m \subseteq U$. Thus again $\omega(z) \in U$. \square

In retrospect, we see that in the definition of $\omega(x)$, we could have taken any base $(P^n)_n$, and indeed any set of closed uniform quasiconvex-stratified sets $(Q^n)_n$ with $Q^n \cap X_0 = P^n \cap X_0$.

Note that we could have carried out the same construction for any $i \in \mathbb{Z}$ to give us a continuous map, $\omega' : X_i \cup \partial X_i \longrightarrow \Xi \cup \partial \Xi$. There is a natural quasi-isometry, $\phi^i : X_0 \longrightarrow X_i$ which extends to a homeomorphism $\phi^i : \partial X_0 \longrightarrow \partial X_i$, which is the identity on $\partial^0 \Xi \equiv \partial X_0 \equiv \partial X_i$.

If $x \in \partial X_0$, and $x^n \in X_0$ is a sequence converging to x , then $\phi^i(x^n)$ converges to $\phi^i(x)$. Thus, $\omega'(\phi^i(x)) = \phi^i(\omega(x))$. Thus, under the identification with $\partial^0 \Xi$, we see that ω agrees with ω' . In this way, we have a canonical map $\omega : \partial^0 \Xi \longrightarrow \partial \Xi$.

Definition. We refer to ω as the *Cannon-Thurston map*.

For future reference, we note that if Ξ has the form $\Xi = \bigvee_{i \in \mathbb{Z}} X_i$, and $Q \subseteq \Xi$ is a quasiconvex-stratified set, then Q can (up to bounded Hausdorff distance) be assumed to have the form $Q = \bigvee_{i \in \mathcal{I}} Q_i$.

8. PROPER SPACES

In this section, we restrict attention to proper stacks, and consider the relation between bi-infinite (or “two-sided”) and semi-infinite (or “one-sided”) stacks.

Suppose that Ξ is a bi-infinite proper hyperbolic stack. We fix some $r \geq 0$ sufficiently large so that every point of $\bigcup_i X_i$ is contained in an r -chain. (Later, we will want to place additional conditions on r .) Normally we shall refer to an r -chain simply as a “chain”. We shall distinguish between *bi-infinite*, *positive* and *negative* chains, indexed respectively by \mathbb{Z} , \mathbb{N} and $-\mathbb{N} = \{-n \mid n \in \mathbb{N}\}$. Each positive or negative chain is a (quasidense subset of a) uniform quasigeodesic, and hence determines an ideal point of $\partial\Xi$. Clearly, a positive chain cannot be asymptotic to a negative chain, so they cannot determine the same point. We define $\partial^+\Xi = \partial_r^+\Xi$ (respectively $\partial^-\Xi = \partial_r^-\Xi$) to be the subsets of $\partial\Xi$ determined by positive (respectively negative) chains. Thus $\partial^+\Xi \cap \partial^-\Xi = \emptyset$.

By a *partial* positive chain, we mean a chain, $(x_i)_{i \in \mathcal{I}}$, indexed by an initial segment, \mathcal{I} , of \mathbb{N} , i.e. either $\mathcal{I} = \{0, \dots, n\}$ for some $n \in \mathbb{N}$ or $\mathcal{I} = \mathbb{N}$. We define the *length* of $(x_i)_i$ to be n or ∞ respectively. We define the *terminal point* of $(x_i)_i$ to be either x_i^n or the limit of the sequence $(x_i)_i$ in $\partial^+\Xi$. We shall frequently use the notation \underline{x} for a partial chain, and write x for its terminal point. We adopt similar terminology and conventions for partial negative chains.

Lemma 8.1. *Suppose that $(\underline{x}^n)_n$ is a sequence of partial positive chains in Ξ , with terminal points x^n . Suppose that x_0^n converges to $a \in X_0 \cup \partial X_0$. Then x^n converges to $\omega(a) \in \Xi \cup \partial\Xi$.*

Proof. By continuity ω (Lemma 7.4), $\omega(x_0^n)$ converges to $\omega(a)$ in $\Xi \cup \partial\Xi$. Now the chains \underline{x}^n are uniformly quasigeodesic, and their distance from any fixed point tends to infinity (cf. Lemma 7.2). It follows easily that their terminal points must converge to same ideal point. \square

Now suppose that $(\underline{x}^n)_n$ is any sequence of partial positive chains whose lengths tend to ∞ . Passing to a subsequence, we can suppose that x_0^n converges to a point $x_0 \in X_0 \cup \partial X_0$. If $x_0 \in X_0$, then x_1^n remains in a compact subset of X_1 , and so subconverges on a point $x_1 \in X_1$. Continuing by a diagonal sequence argument, we can pass to a subsequence of our partial chains, so that for each i , x_i^n converges to a point x_i in X_i . Clearly $\underline{x} = (x_i)_i$ is an (infinite) positive chain. Now, for large n , x_i^n remains uniformly close to x_i for arbitrarily many i . It follows that the terminal points of \underline{x}^n converge on the terminal point, x , of \underline{x} .

Proposition 8.2. *If Ξ is a bi-infinite hyperbolic stack, then $\partial\Xi = \partial^+\Xi \cup \partial^-\Xi \cup \omega(\partial^0\Xi)$.*

Proof. Let $x \in \partial\Xi$. Let $(x^n)_n$ be a sequence of points of $\bigcup_{i \in \mathbb{Z}} X_i$ tending to x . Passing to a subsequence, we can suppose (without loss of generality) that $x_n \in \bigcup_{i \in \mathbb{N}} X_i$. We can represent x^n as the terminal point of a partial positive chain, \underline{x}^n . Again, passing to a subsequence, we can suppose that \underline{x}_0^n converges to a point, x_0 , in $X_0 \cup \partial X_0$. If $x_0 \in \partial X_0$, then by Lemma 8.1, x^n converges to $\omega(x_0)$ and so $x = \omega(x_0) \in \omega(\partial^0\Xi)$. If $x_0 \in X_0$, then the lengths of the partial chains necessarily tend to ∞ . Thus, as discussed above, some subsequence of x^n converges to a point of $\partial^+\Xi$. In this case, $x \in \partial^+\Xi$. \square

Note that exactly the same reasoning applies to a semi-infinite stack, Ξ^+ . We obtain a Cannon-Thurston map $\omega^+ : \partial^0\Xi^+ \longrightarrow \partial\Xi^+$. In this case, we can obtain:

Proposition 8.3. *If Ξ^+ is a semi-infinite proper hyperbolic stack, then $\partial\Xi^+ = \partial^+\Xi^+ \cup \omega^+(\partial^0\Xi^+)$.*

Now, again let Ξ be a bi-infinite proper hyperbolic stack. Now X_0 coarsely separates Ξ into semi-infinite substacks, Ξ^+ and Ξ^- , with sheets $(X_i)_{i \in \mathbb{N}}$ and $(X_i)_{i \in -\mathbb{N}}$ respectively. These substacks are well defined up to bounded Hausdorff distance.

To make this more explicit, we can assume (up to quasi-isometry) that Ξ has the form $\Xi = \bigvee_{i \in \mathbb{Z}} X_i$. We set $\Xi^+ = \bigvee_{i \in \mathbb{N}} X_i$ and $\Xi^- = \bigvee_{i \in -\mathbb{N}} X_i$, which we can take to be substacks of Ξ . Let ρ^\pm be the induced path-metric on Ξ^\pm . Note that, with this construction, we have $\Xi^+ \cup \Xi^- = \Xi$ and $\Xi^+ \cap \Xi^- = X_0$.

We note the following particular case of Proposition 6.7.

Proposition 8.4. *If Ξ is a bi-infinite hyperbolic stack, then the semi-infinite substacks Ξ^+ and Ξ^- are also uniformly hyperbolic.*

Now, Ξ^+ and Ξ^- have intrinsic hyperbolic boundaries, $\partial\Xi^+$ and $\partial\Xi^-$, respectively. We can identify the spaces $\partial^0\Xi^+ \equiv \partial^0\Xi^- \equiv \partial^0\Xi \equiv \partial X_0$. Clearly, every positive chain in Ξ^+ is a positive chain in Ξ and vice versa. Moreover two such chains are asymptotic in Ξ^+ if and only if they are asymptotic in Ξ . The same applies to negative chains. We thus have set-theoretical identifications $\partial^+\Xi^+ \equiv \partial^+\Xi$ and $\partial^-\Xi^- \equiv \partial^-\Xi$.

In summary, we have Cannon-Thurston maps:

$$\omega^\pm : \partial^0\Xi \longrightarrow \partial\Xi^\pm$$

and

$$\omega : \partial^0\Xi \longrightarrow \partial\Xi,$$

which extend to continuous maps:

$$\omega^\pm : X_0 \cup \partial X_0 \longrightarrow \Xi^\pm \cup \partial \Xi^\pm$$

and

$$\omega : X_0 \cup \partial X_0 \longrightarrow \Xi \cup \partial \Xi$$

via inclusion on X_0 .

The next step will be to define maps:

$$\tau^\pm : \partial \Xi^\pm \longrightarrow \partial \Xi,$$

which extend to continuous maps:

$$\tau^\pm : \Xi^\pm \cup \partial \Xi^\pm \longrightarrow \Xi \cup \partial \Xi$$

via inclusion on Ξ^\pm , and such that $\omega = \tau^\pm \circ \omega^\pm$.

We shall need:

Lemma 8.5. *Suppose $a \in \partial^0 \Xi$ and $y \in \partial^+ \Xi$. Then $\omega(a) = y$ if and only if there is a sequence, $(\underline{x}^n)_n$, of positive chains, each converging to y , and with x_0^n converging to a in $X_0 \cup \partial X_0$.*

Proof. The “if” part is an immediate consequence of Lemma 8.1. The “only if” part will be proven in Section 9 (see Lemma 9.2). \square

(In fact, we can choose the chains so that x_0^n converges to a along a geodesic ray in X_0 .)

Note that this lemma applies (by the same argument) to semi-infinite stacks. Thus, the existence of such a chain is also equivalent to the statement $\omega^+(a) = y$. We deduce:

Lemma 8.6. *Given $a \in \partial^0 \Xi$ and $y \in \partial^+ \Xi$, we have $\omega^+(a) = y$ if and only if $\omega(a) = y$.*

We also note;

Lemma 8.7. *If $x, y \in \partial^0 \Xi$ are distinct and $\omega^+(x) = \omega^+(y)$, then $\omega^+(x) \in \partial^+ \Xi$.*

We postpone the proof until Section 9 (see Lemma 9.3).

Putting this together with Lemma 8.6, we deduce:

Lemma 8.8. *If $x, y \in \partial^0 \Xi$ and $\omega^+(x) = \omega^+(y)$, then $\omega(x) = \omega(y)$.*

Although not needed for the construction of the maps τ^\pm , we note for future reference that:

Lemma 8.9. *If $x, y \in \partial^0 \Xi$ and $\omega(x) = \omega(y)$, then either $\omega^+(x) = \omega^+(y)$ or $\omega^-(x) = \omega^-(y)$.*

Again we postpone the proof until Section 9. Note that by Lemma 8.7, and the fact that $\partial^+\Xi \cap \partial^-\Xi = \emptyset$, if $x \neq y$, then the two latter possibilities are mutually exclusive. Indeed, we note:

Lemma 8.10. *If $x, y, z \in \partial^0\Xi$ and $\omega^+(x) = \omega^+(y)$ and $\omega^-(x) = \omega^-(z)$, then either $x = y$ or $x = z$.*

Note that, in particular, the map $(\omega^+, \omega^-) : \partial^0\Xi \longrightarrow \partial\Xi^+ \times \partial\Xi^-$ is injective. This places a restriction in the topology of $\partial^0\Xi$, as we shall see in Section 10.

We now define the map $\tau^+ : \partial^+\Xi \longrightarrow \partial\Xi$ by setting $\tau^+(y) = y$ for $y \in \partial^+\Xi^+ \equiv \partial^+\Xi$, and $\tau^+(\omega^+(x)) = \omega(x)$ for $x \in \partial^0\Xi$. By Proposition 8.3 and Lemmas 8.6 and 8.8, this is well defined. Clearly $\tau^+ \circ \omega^+ = \omega$. We set $\tau^+|_{\Xi^+}$ to be the inclusion of Ξ^+ into Ξ .

Lemma 8.11. *The map $\tau^+ : \Xi^+ \cup \partial\Xi^+ \longrightarrow \Xi \cup \partial\Xi$ is continuous.*

Proof. Let $F = \omega^+(\partial^0\Xi) \subseteq \partial\Xi$. Now $\omega = (\tau^+|_F) \circ \omega^+ : \partial^0\Xi \longrightarrow \partial\Xi$ is continuous, and $\tau^+|_F$ is surjective. Since all spaces are compact, it follows that $\tau^+|_F$ is continuous.

Clearly $\tau^+|_{\Xi^+}$ is continuous. It therefore remains to show that if $y^n \in \Xi^+ \cup \partial\Xi^+$ is a sequence of points converging on $y \in \partial\Xi^+$, then $\tau^+(y^n)$ converges to $\tau(y)$. In fact, it's enough to show that some subsequence converges to $\tau(y)$. We can assume that $y^n \in \bigcup_{i \in \mathbb{N}} X_i \cup \partial^+\Xi$. Let \underline{x}^n be a sequence of partial positive chains with terminal points y^n . As in the proof of Proposition 8.2, we can reduce to two cases. Either \underline{x}^n converges pointwise to an infinite positive chain, \underline{x} , or x_0^n converges to a point $a \in \partial X_0 \equiv \partial^0\Xi$. In the former case, y^n converges to the terminal point of \underline{x} , both in $\Xi^+ \cup \partial\Xi^+$ and $\Xi \cup \partial\Xi$. Thus, the terminal point of \underline{x} is y , and so $\tau^+(y^n) = y^n \rightarrow \tau^+(y) = y$. In the latter case, we see by Lemma 8.1 (or by its variant for semi-infinite stacks) that y^n converges to $\omega^+(a)$ in $\Xi^+ \cup \partial\Xi^+$ and to $\omega(a)$ in $\Xi \cup \partial\Xi$. Thus $y = \omega^+(a)$ and $y^n = \tau^+(y^n) \rightarrow \omega(x) = \tau^+(\omega(a)) = \tau^+(\omega^+(a))$ in $\Xi \cup \partial\Xi$ as required. \square

Clearly the same reasoning gives us a continuous map $\tau^- : \Xi^- \cup \partial\Xi^- \longrightarrow \Xi \cup \partial\Xi$.

Finally, we note that it follows from Proposition 8.2 that:

Lemma 8.12. *If Ξ is a bi-infinite proper hyperbolic stack, then $\partial\Xi = \tau^+(\partial\Xi^+) \cup \tau^-(\partial\Xi^-)$.*

Clearly, $\tau^+(\partial^+\Xi) \cap \tau^-(\partial^-\Xi) = \omega(\partial^0\Xi)$.

9. INTERVAL STACKS

In this section we shall consider “interval stacks”, i.e. hyperbolic stacks whose sheets are all isometric to subintervals of the real line. We shall see that the boundary of a semi-infinite interval stack is a (possibly degenerate) interval. The boundary of a bi-infinite interval stack may be a circle, an interval or the disjoint union of two intervals. With future applications in mind, we shall phrase our arguments in terms of interval substacks of a given stack, Ξ . (A similar construction was described in [Mit1, Mit2], where it was termed a “ladder”.)

Let Ξ be, for the moment, a semi-infinite proper hyperbolic stack. We use the notation $\underline{x} = (x_i)_{i \in \mathbb{N}}$ for a chain in Ξ , and write Ξ for the terminal point of \underline{x} .

First, we consider the case where the sheets are compact intervals. Suppose \underline{x} and \underline{y} are chains. Let $\Upsilon = \Upsilon(\underline{x}, \underline{y}) = \bigcup_{i \in \mathbb{N}} [x_i, y_i]_i$. This is well defined up to uniformly bounded Hausdorff distance. By Lemma 6.1, Υ is a quasiconvex-stratified set, and hence quasiconvex, by Lemma 7.1. As usual, we write $\partial\Upsilon = \partial_{\Xi}\Upsilon$ for the closure of Υ in $\Xi \cup \partial\Xi$ intersected with $\partial\Xi$. Clearly, $x, y \in \partial\Upsilon$. Note that if $\Xi = \bigvee_{i \in \mathbb{N}} X_i$, then (at least with appropriate quasi-isometries) we can identify the stack $\hat{\Upsilon} = \bigvee_{i \in \mathbb{N}} [x_i, y_i]_i$ as a subset of Ξ , containing Υ as a quasidense subset. Now $\hat{\Upsilon}$ is intrinsically hyperbolic, and we may identify $\partial\Upsilon$ with the Gromov boundary, $\partial\hat{\Upsilon}$ of $\hat{\Upsilon}$. Note that $\hat{\Upsilon}$ is one-ended, since a base of neighbourhoods of the end is given by the collection $\left(\bigvee_{j \geq i} [x_j, y_j]_j \right)_{i \in \mathbb{N}}$. We conclude that $\partial\Upsilon = \partial\hat{\Upsilon}$ is connected. Moreover, $\partial^0\hat{\Upsilon} = \emptyset$ and so, by Proposition 8.5, we have $\partial^+\hat{\Upsilon} = \partial\hat{\Upsilon}$. Thus, every point of $\partial\Upsilon$ is the terminal point of a positive chain in Υ . Moreover, we can assume that every point of Υ is contained in a chain in Υ . (We, of course, need to choose the chain constant appropriately taking account of the fact that the parameters of $\hat{\Upsilon}$ depend on the parameters of Ξ . We can assume that such chains do not cross, i.e. if $z_j, w_j \in [x_j, y_j]_j$, with $z_j \in [x_j, w_j]_j$, then we can extend to chains $(z_i)_i$ and $(w_i)_i$ such that $z_i \in [x_i, z_i]_i$ for all i .)

Suppose \underline{z} is a chain in Υ , so that $\Upsilon(\underline{x}, \underline{z})$ and $\Upsilon(\underline{z}, \underline{y})$ are quasiconvex subsets of Υ which intersect precisely in the chain \underline{z} . It follows that $\partial\Upsilon(\underline{x}, \underline{z}) \cup \partial\Upsilon(\underline{z}, \underline{y}) = \partial\Upsilon$ and $\partial\Upsilon(\underline{x}, \underline{z}) \cap \partial\Upsilon(\underline{z}, \underline{y}) = \{z\}$. In particular, if $z \neq x, y$, then z separates x from y in $\partial\Upsilon$. Since every point of $\partial\Upsilon$ is represented by a positive chain in Υ , we see that every point of $\partial\Upsilon \setminus \{x, y\}$ separates x from y . Since $\partial\Upsilon$ is a metrisable continuum, it follows that it is an interval.

In summary, we have shown:

Lemma 9.1. *If \underline{x} and \underline{y} are chains in Ξ , then $\partial\Upsilon(\underline{x}, \underline{y})$ is a (possibly degenerate) topological interval with endpoints x and y . Moreover, $\partial\Upsilon(\underline{x}, \underline{y}) \subseteq \partial^+\Xi$.*

Next we move on to consider the case where the sheets are geodesic rays. Suppose \underline{x} is a chain in Ξ and that $a \in \partial^0\Xi$. Let $\Upsilon = \Upsilon(\underline{x}, a) = \bigcup_{i \in \mathbb{N}} [x_i, a]$. Again this is (uniformly) quasiconvex, and $x, \omega(a) \in \partial\Upsilon$. By similar arguments as above, we see that $\partial\Upsilon$ is connected. Moreover, $\partial\Upsilon \setminus \{\omega(a)\} \subseteq \partial^+\Xi$. In fact, every element, z , of $\partial\Upsilon \setminus \{\omega(a)\}$ is the terminal point of a chain, \underline{z} , in Υ . Note that $\Upsilon(\underline{x}, \underline{z}) \cup \Upsilon(\underline{z}, a) = \Upsilon$, and that $\Upsilon(\underline{x}, \underline{z}) \cap \Upsilon(\underline{z}, a) = \{z\}$. We therefore conclude, as before, that:

Lemma 9.2. *$\partial\Upsilon(\underline{x}, a)$ is a (possibly degenerate) interval with endpoints x and $\omega(a)$. Moreover, $\partial\Upsilon(\underline{x}, a) \subseteq \partial^+\Xi \cup \{\omega(a)\}$.*

Note that if $x = \omega(a)$, then $\partial\Upsilon(\underline{x}, a) = \{x\}$. In particular, every chain in Υ converges on x . Since every point of $[x_0, a]_0$ is contained in a chain, this proves Lemma 8.5.

On the other hand, if $x \neq \omega(a)$, we can find a chain \underline{z} in Υ such that $x \notin \Upsilon(\underline{z}, a)$.

We now move on to the case where all sheets are bi-infinite geodesics. Suppose $a, b \in \partial^0\Xi$ are distinct. Let $\Upsilon = \Upsilon(a, b) = \bigcup_{i \in \mathbb{N}} [a, b]_i$. By the same reasoning as before, we have:

Lemma 9.3. *$\partial\Upsilon(a, b)$ is a (possibly degenerate) interval with endpoints $\omega(a)$ and $\omega(b)$. Moreover, $\partial\Upsilon(a, b) \subseteq \partial^+\Xi \cup \{\omega(a), \omega(b)\}$.*

Note that if $\omega(a) = \omega(b)$, then every point of $[a, b]_0$ lies on a positive chain. In particular, $\omega(a) \in \partial^+\Xi$. This proves Lemma 8.7.

On the other hand, if $\omega(a) \neq \omega(b)$, then we can find chains \underline{x} and \underline{y} in Υ such that $\partial\Upsilon(\underline{x}, a) \cap \partial\Upsilon(\underline{y}, b) = \emptyset$. Indeed, given that every point of Υ lies on a chain, we can arrange that $\rho(\Upsilon(\underline{x}, a), \Upsilon(\underline{y}, b))$ is arbitrarily large.

We now move on to consider the case where Ξ is a bi-infinite stack. Here, we shall only deal with the case where $a, b \in \partial^0\Xi$. The other cases can be dealt with similarly.

Let $\Upsilon = \bigcup_{i \in \mathbb{Z}} [a, b]_i$, $\Upsilon^+ = \bigcup_{i \in \mathbb{N}} [a, b]_i$ and $\Upsilon^- = \bigcup_{i \in -\mathbb{N}} [a, b]_i$.

We begin with the proof of Lemma 8.9:

Proof. Suppose that $\omega^+(a) \neq \omega^+(b)$ and $\omega^-(a) \neq \omega^-(b)$. We want to prove that $\omega(a) \neq \omega(b)$. From an earlier observation, we can find positive chains \underline{x}^+ and \underline{y}^+ in Υ^+ such that $\rho^+(\Upsilon^+(\underline{x}^+, a), \Upsilon^+(\underline{y}^+, b))$ is arbitrarily large. Similarly, we can find negative chains, \underline{x}^- and \underline{y}^- in Υ^- such that $\rho^-(\Upsilon^-(\underline{x}^-, a), \Upsilon^-(\underline{y}^-, b))$ is arbitrarily large. In fact, since every point of $[a, b]_0$ is the basepoint of both a positive and

a negative chain and that such chains can be taken not to cross, we can assume, without loss of generality, that $x_0^+ = x_0^-$ and $y_0^+ = y_0^-$. We can therefore combine the positive and negative chains to give bi-infinite chains \underline{x} and \underline{y} . Moreover, $\rho(\Upsilon(\underline{x}, a), \Upsilon(\underline{y}, b))$ is arbitrarily large. In particular, we can arrange that $\partial\Upsilon(\underline{x}, a) \cap \partial\Upsilon(\underline{y}, b) = \emptyset$. Since $\omega(a) \in \partial\Upsilon(\underline{x}, a)$ and $\omega(b) \in \partial\Upsilon(\underline{y}, a)$, we have $\omega(a) \neq \omega(b)$. \square

Now we have already proven all the essential ingredients of Lemmas 8.11 and 8.12. Applying these to Υ (or more precisely to $\hat{\Upsilon} = \bigvee_{i \in \mathbb{Z}} [a, b]_i$), we see that $\partial\Upsilon = F^+ \cup F^-$ and $F^+ \cap F^- = \{\omega(a), \omega(b)\}$, where $F^\pm = \tau^\pm(\partial\Upsilon^\pm)$.

Now $\partial\Upsilon^+$ and $\partial\Upsilon^-$ are intervals with endpoints $\{\omega^+(a), \omega^+(b)\}$ and $\{\omega^-(a), \omega^-(b)\}$. Moreover, τ^\pm restricted to the interior of $\partial\Upsilon^\pm$ is injective, and the image of the interiors of $\partial\Upsilon^+$ and $\partial\Upsilon^-$ are disjoint. By Lemma 8.7 we see that we cannot have simultaneously $\omega^+(a) = \omega^+(b)$ and $\omega^-(a) = \omega^-(b)$. Thus at least one of the intervals $\partial\Upsilon^\pm$, and hence one of the intervals F^\pm , must be non-degenerate. Finally note that (by Lemma 8.9) if $\omega(a) = \omega(b)$ then either $\omega^+(a) = \omega^+(b)$ or $\omega^-(a) = \omega^-(b)$, and so in this case that (precisely) one of the intervals F^\pm is degenerate. In summary, we see that $\partial\Upsilon$ is either the union of two non-degenerate intervals connected at their endpoints or else a non-degenerate interval with its endpoints identified. Either way, we obtain:

Proposition 9.4. *If Ξ is a bi-infinite proper hyperbolic stack, and $a, b \in \partial^0\Xi$ are distinct, then $\partial\Upsilon(a, b)$ is a topological circle.*

The remaining cases are readily dealt with similarly. We see that the boundary of a bi-infinite stack of rays is a non-degenerate interval, and the boundary of a bi-infinite stack of compact intervals is the disjoint union of two (possibly degenerate) intervals.

10. DENDRITES

In this section, we show that the boundary of a semi-infinite proper hyperbolic stack is a dendrite (Proposition 10.2). We deduce that the boundary of proper bi-infinite hyperbolic stack is locally connected (Proposition 10.4). We note that Proposition 10.2 greatly restricts the geometry of the sheets in an infinite stack (Proposition 10.5).

A *dendrite* can be defined as a metrisable continuum (i.e. compact and connected) in which every pair of distinct points are separated by a third. A dendrite is necessary locally connected, and every pair of points are connected by a unique arc, i.e. a subset homeomorphic to

a (possibly degenerate) interval. There are many equivalent ways of characterising dendrites. We shall use the following description.

Suppose that M is a metrisable continuum, and that to every pair of points, $x, y \in M$, we associate a closed subset, $[x, y] \subseteq M$, satisfying the following for all $x, y, z \in M$:

$$(D1) \quad [x, y] = [y, x],$$

$$(D2) \quad [x, x] = \{x\},$$

$$(D3) \quad [x, y] \subseteq [x, z] \cup [z, y],$$

$$(D4) \quad \text{If } x \neq y, \text{ then } [x, y] \neq \{x, y\}, \text{ and}$$

$$(D5) \quad \text{If } x \neq y, \text{ then there is a neighbourhood, } U, \text{ of } x \text{ in } M \text{ such that for all } w \in U, y \notin [x, w].$$

Note that it follows that $x, y \in [x, y]$. In fact, we shall see directly from our construction that $[x, y]$ is an arc connecting x to y .

Lemma 10.1. *If M is a metrisable continuum satisfying properties (D1)–(D5) above, then M is a dendrite.*

Proof. Given $z \in M$, define a relation on $M \setminus \{z\}$ by deeming $x, y \in M \setminus \{z\}$ to be related if $z \notin [x, y]$. By properties (D1), (D2) and (D3), this is an equivalence relation, and by (D5), each equivalence class is open.

Suppose now that $x, y \in M$ are distinct. By (D4) there is some $z \in [x, y] \setminus \{x, y\}$. Now x and y lie in different equivalence classes of $M \setminus \{z\}$, and so z separates x from y . \square

Suppose now that Ξ is a semi-infinite proper hyperbolic stack. Recall, from Section 9, the definition of the quasiconvex subsets $\Upsilon(\underline{x}, \underline{y})$, $\Upsilon(\underline{x}, a)$ and $\Upsilon(a, b)$ where $\underline{x}, \underline{y}$ are chains, and $a, b \in \partial^0 \Xi$. These sets are well defined up to finite (in fact uniformly bounded) Hausdorff distance. If \underline{x}' and \underline{y}' are chains asymptotic to \underline{x} and \underline{y} respectively, then $\text{hd}(\Upsilon(\underline{x}, \underline{y}), \Upsilon(\underline{x}', \underline{y}')) < \infty$ and $\text{hd}(\Upsilon(\underline{x}, a), \Upsilon(\underline{x}', a)) < \infty$ (though, of course, these need not be uniformly bounded). If $x, y \in \partial^+ \Xi$ are terminal points of \underline{x} and \underline{y} , we write $[[x, y]] = \partial \Upsilon(\underline{x}, \underline{y})$, $[[x, a]] = \partial \Upsilon(\underline{x}, a)$ and $[[a, b]] = \partial \Upsilon(a, b)$. These are all well defined by the above discussion. If $a = b \in \partial^0 \Xi$, we set $[[a, b]] = \{a\}$. We have thus defined $[[p, q]]$ for all p and q in the formal disjoint union $\partial^+ \Xi \sqcup \partial^0 \Xi$. By Lemmas 9.1, 9.2 and 9.3, we see that, in all cases, $[[p, q]]$ is an interval with endpoints $\omega(p)$ and $\omega(q)$, where (for the purposes of the present discussion) we set $\omega(x) = x$ for all $x \in \partial^+ \Xi$. Moreover, $[[p, q]]$ is degenerate (i.e. a singleton) if and only if $\omega(p) = \omega(q)$.

Suppose $\underline{x}, \underline{y}, \underline{z}$ are three chains in Ξ converging respectively to $x, y, z \in \partial^+\Xi$. For each i , $[x_i, y_i]_i$ lies in a uniformly bounded ρ_i -neighbourhood of $[x_i, z_i]_i \cup [z_i, y_i]_i$. Thus, $\Upsilon(\underline{x}, \underline{y})$ lies in a uniformly bounded neighbourhood of $\Upsilon(\underline{x}, \underline{z}) \cup \Upsilon(\underline{z}, \underline{y})$, from which it follows that $[[x, y]] \subseteq [[x, z]] \cup [[z, x]]$. Note that if $\omega(y) = \omega(z)$, then $[[y, z]]$ is degenerate, and it follows that $[[x, y]] = [[x, z]]$. The same argument applies when one or more of x, y or z belong to $\partial^0\Xi$. Now by Proposition 9.3, $\partial\Xi = \omega(\partial\Xi \sqcup \partial^0\Xi)$, and so if $x, y \in \partial\Xi$, we get a well defined interval, $[x, y]$, with endpoints x and y , by setting $[x, y] = [[p, q]]$ where $x = \omega(p)$ and $y = \omega(q)$. If $z \in \partial\Xi$, we see that $[x, y] \subseteq [x, z] \cup [z, y]$. Note also that $[x, y] = [y, x]$, $[x, x] = \{x\}$ and $[x, y] \subseteq \partial^0\Xi \cup \{x, y\}$.

We are now ready to prove:

Proposition 10.2. *If Ξ is a semi-infinite proper hyperbolic stack, then $\partial\Xi$ is a dendrite.*

Proof. Given any $x, y \in \partial\Xi$, there is an arc, $[x, y]$, defined as above, connecting x to y . It follows that $\partial\Xi$ is a continuum. We have already observed that hypotheses (D1)–(D4) of Lemma 10.1 hold. It remains to verify hypothesis (D5). We rephrase this in the form of Lemma 10.3 below. \square

Lemma 10.3. *Suppose that $x, y \in \partial\Xi$ and that the sequence $x^n \in \partial\Xi$ converges to x . Suppose that $y \in [x^n, x]$ for all n . Then $x = y$.*

Proof. Suppose, for contradiction, that $x \neq y$. We can suppose that $y \neq x^n$ for all n . Since $y \in [x^n, x] \setminus \{x^n, x\}$, it follows that $y \in \partial^+\Xi$. Moreover, passing to a subsequence, we can assume either (Case (1)) that $x^n \in \partial^+\Xi$ for all n , or (Case (2)) that $x^n \in \omega(\partial^0\Xi)$ for all n (see Proposition 8.3).

First, we consider Case (1). Let x^n be the terminal point of a chain, \underline{x}^n . Again after passing to a subsequence, we can assume (cf. Proposition 8.2) that either \underline{x}^n converges pointwise to a chain \underline{x} (Case (1a)) or that x_0^n converges in $X_0 \cup \partial X_0$ to a point $a \in \partial X_0 \equiv \partial^0\Xi$ (Case (1b)).

Let us consider Case (1a). It is necessarily the case that x is the terminal point of \underline{x} . By definition, $[x^n, x] = [[x^n, x]] = \partial\Upsilon(\underline{x}^n, \underline{x}) \subseteq \partial^+\Xi$. Since $y \in [x^n, x]$, for each n , we can represent y by a chain \underline{y}^n in $\Upsilon(\underline{x}^n, \underline{x})$. In other words, for all i and n , we have $y_i^n \in [x_i^n, x_i]_i$. But now, for any given i , $x_i^n \rightarrow x_i$ as $n \rightarrow \infty$. Thus, \underline{y} converges pointwise to \underline{x} . We arrive at the contradiction that $x = y$.

In Case (1b), since $x^n \rightarrow x$, we see by Lemma 8.1, that $x = \omega(a)$. Thus, by definition, $[x^n, x] = [x^n, \omega(a)] = [[x^n, a]] = \partial\Upsilon(\underline{x}^n, a)$. Again, since $y \in [x^n, x] \setminus \{x\}$, for each n we can represent y by a chain \underline{y}^n in $\Upsilon(\underline{x}^n, a)$. In other words, $y_i^n \in [x_i^n, a]_i$ for all i and n . In particular,

$y_0^n \in [x_0^n, a]_0$. Since $x_0^n \rightarrow a$, it follows that $y_0^n \rightarrow a$ in $X_0 \cup \partial X_0$. Again by Lemma 8.1, we see that $y = \omega(a)$, giving the contradiction that $x = y$.

Finally consider Case (2), where $x^n = \omega(a^n)$ for a^n in ∂X_0 . Again, without loss of generality, we can assume that $a^n \rightarrow a \in \partial X_0$. By continuity of ω (Lemma 7.4), we have $\omega(a^n) \rightarrow \omega(a)$, and so $x = \omega(a)$. By definition, $[x^n, x] = [\omega(a^n), \omega(a)] = [[a^n, a]] = \partial\Upsilon(a^n, a)$. Again, $y^n \in [x^n, x] \setminus \{x^n, x\}$ is represented by a chain \underline{y}^n in $\Upsilon(a^n, a)$. In particular, $y_0^n \in [a^n, a]_0$. Since $a^n \rightarrow a$ we have $y_0^n \rightarrow a$ in $X_0 \cup \partial X_0$. Thus by Lemma 8.1, we have the contradiction $y = \omega(a) = x$. \square

We have thus proven Proposition 10.2.

Now, in retrospect, we see that if $x, y \in \partial\Xi$, then $[x, y]$ is the unique arc connecting x to y in $\partial\Xi$. Given $x, y, z \in \partial\Xi$, we can define the *median*, $\text{med}(x, y, z)$, of the triple (x, y, z) as the unique intersection point of the arcs $[x, y]$, $[y, z]$ and $[z, x]$. This can be given a geometric interpretation as follows. Suppose, for example, that $x, y, z \in \partial^+\Xi$. We represent these three points by chains, $\underline{x}, \underline{y}, \underline{z}$. Let w_i be a centre of (x_i, y_i, z_i) in X_i . Thus $\underline{w} = (w_i)_i$ is a chain (possibly after increasing the chain constant by a controlled amount). Its terminal point, w , is equal to $\text{med}(x, y, z)$. Similarly, if $x = \omega(a)$, $y = \omega(b)$ and $z = \omega(c)$ with $a, b, c \in \partial^0\Xi$ distinct, we set w_i to be a centre of (a, b, c) in X_i . Again, $\underline{w} = (w_i)_i$ is a chain with terminal point $\text{med}(x, y, z)$. The remaining cases can all be described similarly.

We finish this section with a brief account of the bi-infinite case.

Let Ξ be a bi-infinite proper hyperbolic stack. Either all the sheets are compact (so that $\partial^0\Xi = \emptyset$), or all the sheets are non-compact. In the former case (by Lemma 8.12 and Proposition 10.2), we see that $\partial\Xi$ is the disjoint union of two dendrites. In the latter case, we obtain a Peano (i.e. locally connected) continuum:

Proposition 10.4. *If Ξ is a bi-infinite proper hyperbolic stack whose sheets are non-compact, then $\partial\Xi$ is a Peano continuum.*

Proof. By Lemma 8.12, $\partial\Xi = \tau^+(\partial\Xi^+) \cup \tau^-(\partial\Xi^-)$ and $\tau^+(\partial\Xi^+) \cap \tau^-(\partial\Xi^-) = \omega(\partial^0\Xi) \neq \emptyset$. By Proposition 10.2, $\tau^\pm(\partial\Xi^\pm)$ is a dendrite. Since τ^\pm is continuous (Lemma 8.11), $\tau^\pm(\partial\Xi^\pm)$ is a Peano continuum. The result follows, given that the union of two intersecting Peano continua is a Peano continuum. \square

It was pointed out to me by Bruce Kleiner that the results of this section place severe restriction on the geometry of sheets that can form an infinite hyperbolic stack. It is impossible, for example, to have an

infinite stack of hyperbolic n -spaces for $n \geq 3$. I am grateful to Bruce Kleiner for his permission to reproduce the argument below.

Suppose that Ξ is a bi-infinite proper hyperbolic stack. By Lemma 8.10, the map $(\omega^+, \omega^-) : \partial^0 \Xi \longrightarrow \partial \Xi^+ \times \partial \Xi^-$ is injective. The latter space is a product of two dendrites and hence contractible and 2-dimensional. It follows, for example, that $\partial^0 \Xi$ has topological dimension at most 2, and cannot contain a 2-sphere.

We can apply this to other situations. We say that Ξ is *uniformly proper* if for each $r \geq 0$, the space of r -balls in Ξ is precompact in the Gromov-Hausdorff topology [Gr2]. This equivalent to (uniform) uniform properness of the sheets.

Proposition 10.5. *If Ξ is an infinite uniformly proper hyperbolic stack, then $\partial \Xi$ embeds in a 2-dimensional contractible space.*

Proof. (Sketch) Take a sequence of basepoints, $p_i \in X_i$, and balls about p_i in Ξ , whose radii tend to ∞ . By a diagonal sequence argument, using compactness in the Gromov-Hausdorff topology [Gr2] (or using ultrafilters), one can construct a bi-infinite stack, Ξ' , so that $\partial^0 \Xi$ embeds in $\partial^0 \Xi'$. By the above observation, this in turn embeds in a 2-dimensional contractible space. \square

The above applies to hyperbolic n -spaces, proving the earlier remark. One can adapt this argument to other spaces. For example, an infinite hyperbolic stack of half-spaces in \mathbb{H}^n for $n \geq 3$, is ruled out, since by a similar diagonal sequence argument, one could construct from it a bi-infinite stack of \mathbb{H}^n 's. No doubt one can give a more systematic elaboration of these ideas, though we shall not pursue that here.

11. STACKS OF HYPERBOLIC PLANES

In this section we consider a semi-infinite proper hyperbolic stack, Ξ , of hyperbolic planes. In other words, X_i is isometric to \mathbb{H}^2 for all $i \in \mathbb{N}$. Together with the results of Section 12, we see that the Cannon-Thurston map arises from a lamination. We write d for the hyperbolic metric. Thus, $\partial^0 \Xi \equiv \partial \mathbb{H}^2$ is a topological circle. First we shall show:

Proposition 11.1. *The Cannon-Thurston map, $\omega : \partial^0 \Xi \longrightarrow \partial \Xi$ is surjective.*

Before setting about the proof, we make some preliminary observations. The following are standard constructions involved in proving the quasi-isometry invariance of isodiametric functions. Suppose that ϕ is a (possibly multivalued) quasi-isometry between geodesic spaces X and Y . Given a closed loop, γ , in X , we can construct a closed

loop γ' in Y as follows. We choose points $x_0, x_1, \dots, x_n = x_0$ along γ which cut γ into $n = O(\text{length}(\gamma))$ subpaths of bounded length, and set $\gamma' = \bigcup_{i=0}^{n-1} [\phi(x_i), \phi(x_{i+1})]$. We note that $\text{hd}(\gamma', \phi(\gamma))$ is bounded, and that $\text{length}(\gamma')$ is bounded above by a linear function of $\text{length}(\gamma)$. (We shall choose the various constants dependent on the parameters of ϕ .)

Suppose that Y is uniformly simply connected, i.e., every curve of a given length bounds a disc of bounded diameter. Suppose that γ bounds a disc, D , (not necessarily embedded). Then γ' bounds a disc, D' , contained a bounded neighbourhood of $\phi(D)$. This can be seen by triangulating D sufficiently finely, mapping vertices under ϕ , mapping edges to geodesic segments, and extending to 2-cells using uniform simple connectedness.

Now consider the case of the hyperbolic plane. We say that a loop β *links* a point x not in the image of β , if it has non-zero winding number about x . Note that β bounds a disc, D , whose image is the union of those compact complementary regions that are unlinked by β . In particular, if β does not link x , then $d(x, D) = d(x, \beta)$.

Suppose now that $\phi : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ is a quasi-isometry. Suppose that $x \in \mathbb{H}^2$ and Γ is a loop in \mathbb{H}^2 . Let γ' be constructed as above. If γ' does not link $\phi(x)$, then it bounds a disc, D , as above. From the earlier discussion, we see that γ bounds a disc, D' , within a bounded neighbourhood of $\phi^{-1}D$. If γ links x , then x must lie in the image of D' . This places an upper bound on $d(\phi(x), D) = d(\phi(x), \gamma')$ and hence on $d(x, \gamma)$. Put another way, if $d(x, \gamma)$ is sufficiently large in relation to the parameters of ϕ , and γ links x , then γ' links $\phi(x)$.

Lemma 11.2. *Suppose $\gamma \subseteq \mathbb{H}^2$ is an embedded circle enclosing a point x , and suppose that $\phi : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ is a quasi-isometry. Then, there is an embedded circle, $\beta \subseteq \mathbb{H}^2$, contained in a uniformly bounded neighbourhood of $\phi(\beta)$ and enclosing the point $\phi(x)$. Moreover, $\text{length}(\beta)$ is bounded above by a linear function of $\text{length}(\gamma)$.*

Proof. Let γ' be the loop arising from the above construction. If $d(x, \gamma)$ is greater than some uniform constant (in fact, the only case that really interests us), then γ' necessarily links $\phi(x)$. If not, then we can always modify γ' by a bounded amount so that this is the case. Now let β' be the boundary of the complementary region of γ' which contains $\phi(x)$. \square

We now return to our stack, Ξ . Suppose that $\underline{x} = (x_i)_i$ is a chain in Ξ , and that $\gamma_0 \subseteq X_0$ is an embedded circle enclosing x_0 . By Lemma 11.2, we can find embedded circles γ_i in X_i , enclosing x_i for each i ,

such that γ_{i+1} is contained in a uniformly bounded ρ -neighbourhood of γ_i . Moreover, $\text{length}(\gamma_i)$ is bounded above by an exponential function in i , depending on $\text{length}(\gamma_0)$ and the parameters of the stack. Note that (for an appropriate chain constant) every point of $\bigcup_{i \in \mathbb{N}} \gamma_i$ is the terminal point of a partial chain contained in $\bigcup_{i \in \mathbb{N}} \gamma_i$.

We now prove Proposition 11.2:

Proof. By Proposition 9.3, $\partial\Xi = \partial^+\Xi \cup \omega(\partial^0\Xi)$, so it's enough to show that the terminal point, x , of any infinite chain, \underline{x} , lies in $\omega(\partial^0\Xi)$. By Lemma 8.1, it suffices to find a sequence, $(\underline{y}^n)_n$, of partial chains, with $\rho_0(x_0, y_0^n) \rightarrow \infty$, whose terminal points converge to x in $\Xi \cup \partial\Xi$. (Since, on passing to a subsequence, we can suppose that $y_0^n \rightarrow a \in \partial X_0$, and then $x = \omega(a)$.)

Let γ_0^n be the circle of radius n in $X_0 \cong \mathbb{H}^2$, centred on x_0 . Let $(\gamma_i^n)_n$ be the sequence of topological circles constructed as above. Suppose that there is some r such that for infinitely many $n \in \mathbb{N}$, there is some $i(n)$ such that $\rho_i(x_i, \gamma_{i(n)}^n) \leq r$. Choose $y^n \in \gamma_{i(n)}^n$ with $\rho(x_i, y^n) \leq r$. Now, y^n is the terminal point of a partial chain, \underline{y}^n lying in $\bigcup_i \gamma_i^n$. In particular, $y_0^n \in \gamma_0^n$ and so $\rho_0(x_0, y_0^n) = n \rightarrow \infty$. Moreover, $i(n) \rightarrow \infty$, and so y^n converges to x . This gives us our desired sequence of partial chains.

Thus, we suppose for contradiction, that for all r , we have $\rho(x_i, \gamma_i^n) \geq r$ for all sufficiently large n and for all i . Fix any such n . Suppose $z_i \in \gamma_i^n$. Thus, z_i is the terminal point of a partial chain, $(z_j)_j$ in $\bigcup_i \gamma_i$. Now the function $[j \mapsto \rho_j(x_j, z_j)] : \{0, \dots, i\} \rightarrow [0, \infty)$ is uniformly quasiconvex (Theorem 6.2), and bounded below by r . Moreover, $\rho_0(x_0, z_0) = n$ is fixed. Thus, provided r is sufficiently large in relation to the parameters of the stack, we see, by Lemma 5.1, that $\rho_j(x_j, z_j)$ is also bounded below by a fixed exponential function of j . Since $z_i \in \gamma_i$ was chosen arbitrarily, we see that $\rho_i(x_i, \gamma_i)$ is bounded below by a fixed exponential function.

Now, $\text{length}(\gamma_i)$ is at least the circumference of a circle of radius $\rho_i(x_i, \gamma_i)$ in $X_i \cong \mathbb{H}^2$. Thus, $\text{length}(\gamma_i)$ is bounded below by a superexponential function of i . However, by construction, it is also bounded above by an exponential function. We therefore arrive at a contradiction. \square

Given four distinct point, x, y, z, w in the circle, S^1 , we say that $\{x, y\}$ *links* $\{z, w\}$, if x and y lie in different components of $S^1 \setminus \{z, w\}$. Identifying $S^1 \cong \partial^0\Xi$, this means that $[x, y]_i$ intersects $[z, w]_i$ for all i .

Recall that $\partial\Xi$ is a dendrite, so that every pair of points, $p, q \in \partial\Xi$ are connected by a unique arc $[p, q]$. We note:

Lemma 11.3. *Suppose that $x, y, z, w \in \partial\Xi$ are distinct, and that $\{x, y\}$ links $\{z, w\}$. Then $[\omega(x), \omega(y)] \cap [\omega(z), \omega(w)] \neq \emptyset$.*

Proof. Since $[x_i, y_i]_i \cap [z_i, w_i]_i \neq \emptyset$ for all i , we see that $\Upsilon(x, y)$ and $\Upsilon(z, w)$ intersect in an unbounded set. Since they are quasiconvex and Ξ is proper, it follows that their boundaries must intersect. \square

In fact, suppose the four points appear in the cyclic order x, z, y, w around $\partial^0\Xi$. Let u_i be a centre for x, y, z in X_i . Then u_i is a bounded distance from $[z, w]_i$. Thus, the chain $(u_i)_i$ determines a point in the intersection of $[\omega(x), \omega(y)]$ and $[\omega(z), \omega(w)]$.

Note, in particular, that if $x, y, z, w \in \partial^0\Xi$ are distinct and $\omega(x) = \omega(y) \neq \omega(z) = \omega(w)$, then $\{x, y\}$ and $\{z, w\}$ are unlinked.

Definition. We say that two disjoint subsets, A and B , of the circle, S^1 , are *linked* if there are distinct $x, y \in A$ and distinct $z, w \in B$ such that $\{x, y\}$ links $\{z, w\}$.

Definition. We say that an equivalence relation on the circle is *unlinked* if the equivalence classes are pairwise unlinked.

Recall that an equivalence relation is *closed* if it is closed as a subset of $S^1 \times S^1$. In this case all the equivalence classes are closed.

We now define an equivalence relation on $\partial^0\Xi \cong S^1$ by writing $x \sim y$ if $\omega(x) = \omega(y)$. By Lemma 10.3 and Proposition 10.1, we see that this closed unlinked equivalence relation with quotient $\partial\Xi \cong \partial^0\Xi/\sim$. Moreover, by Proposition 10.2, this quotient is a dendrite.

12. LAMINATIONS

In this section we consider equivalence relations arising from (geodesic) laminations in the hyperbolic plane, \mathbb{H}^2 . To avoid any confusion with ideal boundaries, we shall write $\text{fr}(Q)$ for the topological boundary or “frontier” of a subset, Q , of \mathbb{H}^2 or of $\mathbb{H}^2 \cup \partial\mathbb{H}^2$. We identify $S^1 \cong \partial\mathbb{H}^2$. Given any closed subset, $A \subseteq \partial\mathbb{H}^2$, we write $\text{hull}(A)$ for its hyperbolic convex hull in $\mathbb{H}^2 \cup \partial\mathbb{H}^2$. This is closed, and $\partial\mathbb{H}^2 \cap \text{hull}(A) = A$. If $A, B \subseteq \partial\mathbb{H}^2$ are disjoint, then A and B are unlinked (as defined in Section 11) if and only if $\text{hull}(A) \cap \text{hull}(B) = \emptyset$. A closed connected subset, Q , of $\mathbb{H}^2 \cup \partial\mathbb{H}^2$ arises as such a convex hull if and only if $\text{fr}(Q \cap \mathbb{H}^2)$ is a union of bi-infinite geodesics. If it is not a single ideal point, then we refer to such a set as a *polygon*. Note that Q has finite area if and only if $\partial_X Q = Q \cap \mathbb{H}^2$ is finite.

Let $\Theta_2(S^1)$ be the space of unordered pairs of S^1 in the Hausdorff topology (or equivalently, $S^1 \times S^1$ minus the diagonal quotiented by the involution that swaps the coordinates). Thus, $\Theta_2(S^1)$ is a Möbius band.

Definition. An (*abstract*) *lamination* on \mathbb{H}^2 is a closed subset of $\Theta_2(S^1)$ such that any pair of non-intersecting elements are unlinked.

We may “realise” each element, $\{x, y\}$, of such a lamination as a bi-infinite geodesic, $[x, y] = \text{hull}\{x, y\} \subseteq \mathbb{H}^2 \cup \partial\mathbb{H}^2$. We refer to such a set as a *leaf* of the lamination. Let Λ be the set of all such leaves. We refer to Λ as the *realisation* of the abstract lamination, and to a set arising in this way as a (*geodesic*) *lamination* in \mathbb{H}^2 . Note that distinct leaves can only intersect in ideal points. Moreover, $\bigcup \Lambda$ is a closed subset of $\mathbb{H}^2 \cup \partial\mathbb{H}^2$ which we refer to as the *support* of Λ . A leaf, $\lambda \in \Lambda$, is *isolated* if $\lambda \cap \partial\mathbb{H}^2$ is an isolated point of the abstract lamination.

Suppose that Λ is a lamination with no isolated leaves. We define a relation \sim on $\mathbb{H}^2 \cup \partial\mathbb{H}^2$ by writing $x \sim y$ if the geodesic segment $[x, y]$ does not cross any leaf of Λ , or more precisely, for all $\lambda \in \Lambda$, either $[x, y] \subseteq \lambda$ or $\lambda \cap [x, y] \subseteq \{x, y\}$. One can readily verify that this is a closed equivalence relation. Note that its restriction to $\partial\mathbb{H}^2$ can be defined in terms of the abstract lamination. Moreover, we see that:

Lemma 12.1. *The relation \sim thus defined on $\partial\mathbb{H}^2$ is a closed unlinked equivalence relation.*

A *stratum* of Λ is an equivalence class of the relation \sim in $\mathbb{H}^2 \cup \partial\mathbb{H}^2$. Equivalently a stratum is the convex hull of an equivalence class in $\partial\mathbb{H}^2$. Let $\mathcal{H}(\Lambda)$ be the set of strata of Λ . We see immediately that $\bigcup \mathcal{H}(\Lambda) = \mathbb{H}^2 \cup \partial\mathbb{H}^2$. Note that since every stratum meets $\partial\mathbb{H}^2$, we can identify the quotients $(\mathbb{H}^2 \cup \partial\mathbb{H}^2)/\sim$ and $\partial\mathbb{H}^2/\sim$. It is well known that such a quotient is a dendrite. (We omit the proof here, since we shall not be requiring it directly.)

We note that we can allow for a lamination with isolated leaves. In this case, we take \sim to be the transitive closure of the relation defined above. A *stratum* of Λ is then either a leaf of Λ or the closure of a complementary component of $\mathbb{H}^2 \setminus \bigcup \Lambda$. Each \sim -class is a union of strata.

Conversely one can define laminations in terms of equivalence relations on $\partial\mathbb{H}^2$. Let \sim be a closed unlinked equivalence relation on $\partial\mathbb{H}^2$. Let \mathcal{H} be the set of convex hulls of equivalence classes. Thus each element of \mathcal{H} is a polygon, and we write Λ for the set of geodesics that arise in the frontiers of such polygons. It is easily verified that Λ is a lamination without isolated leaves, and that each element of \mathcal{H} is a stratum of Λ , in other words, $\mathcal{H} \subseteq \mathcal{H}(\Lambda)$. From this it follows that $\bigcup \mathcal{H}$ is a closed subset of $\mathbb{H}^2 \cup \partial\mathbb{H}^2$.

Suppose that $Q \in \mathcal{H}(\Lambda) \setminus \mathcal{H}$. Let $A = \partial_x Q$, so that $Q = \text{hull}(A)$. Now each of the frontier geodesics in Q correspond to distinct \sim -classes.

In particular, A has not isolated points. Thus, A is a perfect subset of the circle $S^1 \equiv \partial\mathbb{H}^2$, and each \sim -class restricted to A is either a singleton or the pair of endpoints of a complementary interval of $S^1 \setminus A$. Now, A/\sim is obtained by collapsing each complementary interval to a point, and is thus a topological circle. In other words, A projects to an embedded circle in $\partial\mathbb{H}^2/\sim$. This is impossible if \mathbb{H}^2/\sim is a dendrite. We conclude that if this quotient is a dendrite, then $\mathcal{H} = \mathcal{H}(\Lambda)$, and so $\bigcup \mathcal{H} = \mathbb{H}^2 \cup \partial\mathbb{H}^2$. In other words, \mathcal{H} is precisely the set of strata of Λ , and \sim is the equivalence class arising from Λ as described earlier.

In summary, we have shown:

Lemma 12.2. *Suppose that \sim is a closed unlinked equivalence relation on $\partial\mathbb{H}^2$ such that $\partial\mathbb{H}^2/\sim$ is a dendrite. Then there is a unique lamination Λ with no isolated leaves such that the strata of Λ are precisely the convex hulls of equivalence classes of \sim .*

We note that if we start with a lamination Λ with isolated leaves and take the equivalence relation, \sim , as defined earlier, then the lamination we recover by the above construction is equal to Λ with all its isolated leaves removed.

Finally, we note that the equivalence relation arising from a stack of hyperbolic planes satisfy the the hypotheses of Lemma 11.3, as discussed in the last section. Thus:

Proposition 12.3. *Suppose that Ξ is a semi-infinite proper hyperbolic stack of hyperbolic planes. Let $\omega : \partial^0\Xi \rightarrow \partial\Xi$ be the (surjective) Cannon-Thurston map. Then there is a unique abstract lamination, Λ , on $\partial^0\Xi$, without isolated leaves such that if $x, y \in \partial^0\Xi$, then $\omega(x) \neq \omega(y)$ if and only if there is an element of the lamination, disjoint from $\{x, y\}$, which links $\{x, y\}$.*

We shall say that two leaves of Λ are *separated* if they map to different points in the quotient dendrite.

13. GEOMETRIC AND ERGODIC PROPERTIES OF LAMINATIONS ON SURFACES

In this section, we describe a property of a lamination on a compact surface which we term ‘‘regularity’’. We show that this implies unique ergodicity of the lamination. We shall see later that the ending lamination of a stable hyperbolic stack of surfaces is regular (Proposition 15.4). Putting this together with results of [Min5], one can deduce that the end of a hyperbolic 3-manifold has bounded geometry if and only if its ending lamination is regular. Further discussion of unique ergodicity in relation to Teichmüller rays can be found in [K, Mas]. The

notion of regularity plays a central role in the proof of Theorem 20.1. In the course of our analysis, we introduce the notions of “projected distance” and “decompositions” which will also feature in subsequent sections.

We begin by recalling some basic facts about laminations on hyperbolic surfaces. Some standard references are [CanaEG, CasB].

Let Σ be a closed surface with $\text{genus}(\Sigma) \geq 2$, and let $\Gamma = \pi_1(\Sigma)$. Given a hyperbolic structure on Σ , we can represent it as a quotient $\Sigma = \mathbb{H}^2/\Gamma$ of \mathbb{H}^2 under a free properly discontinuous action of Γ . In particular we get an induced action of Γ by homeomorphism on $S^1 = \partial\mathbb{H}^2$. Any two such actions (corresponding to different hyperbolic structures) will be topologically conjugate. We can define an *abstract lamination* on Σ as an equivariant lamination in S^1 . Thus, given a hyperbolic structure, we get a realisation as a geometric lamination, Λ . This projects to a lamination, Λ_Σ , on Σ , which we can view as a set of disjoint geodesics, whose *support*, $\bigcup \Lambda_\Sigma$, is closed in Σ . It is always the case that the interior of $\bigcup \Lambda_\Sigma$ is empty. (Indeed $\bigcup \Lambda_\Sigma$ has Hausdorff dimension 1.) We say that $\bigcup \Lambda_\Sigma$ *fills* Σ if every closed curve on Σ intersects $\bigcup \Lambda_\Sigma$. In this case, each complementary component is a finite-sided polygon. We say that Λ_Σ (or Λ) is *maximal* if it is not contained in any strictly larger lamination, or equivalently, if every complementary region is an ideal triangle. Every lamination, Λ_Σ , lies inside a maximal lamination obtained by adding a finite number of isolated leaves.

A *transverse measure* to a lamination, Λ_Σ , is a Borel measure on each transversal geodesic segment to the lamination, which is invariant under pushing along the lamination. (See [CanaEG] for more precision.) A *measured lamination* is a lamination equipped with a transversal measure. (We won't necessarily assume that its support is all of $\bigcup \Lambda_\Sigma$.) Any geodesic lamination admits a non-zero transverse measure. We say that the lamination is *uniquely ergodic* if this measure is unique up to scale.

Let us fix, for the moment, a hyperbolic structure, $\Sigma = \mathbb{H}^2/\Gamma$. We denote by d the metric on either \mathbb{H}^2 and Σ . Recall that the *injectivity radius*, $\text{inj}(\Sigma)$, is half the length of the shortest geodesic, or equivalently,

$$\text{inj}(\Sigma) = \frac{1}{2} \min\{d(x, gx) \mid x \in \mathbb{H}^2, g \in \Gamma \setminus \{1\}\}.$$

Let Λ and Λ_Σ be respectively the realisations of an abstract lamination in \mathbb{H}^2 and Σ . We shall define the notion of “regularity” for Λ_Σ referred to above with reference to this particular hyperbolic structure. Eventually we will need to verify that it only depends on the underlying abstract lamination (Propositions 13.9 and 13.11).

Definition. A *horizontal interval* (in \mathbb{H}^2 or Σ) is a compact subinterval of a leaf of the lamination.

We write \mathcal{H} (respectively \mathcal{H}_Σ) for the set of horizontal intervals in \mathbb{H}^2 (respectively Σ). Given $E \in \mathcal{H}_\Sigma$, a *surrounding curve* of E is a closed curve in $\Sigma \setminus E$ which is essential in $\Sigma \setminus E$.

Definition. The *exterior length*, $\text{ext}(E)$, of E is the infimum of the lengths of surrounding curves of E .

(Clearly we get the same definition if we restrict to simple surrounding curves.)

The exterior length can be thought of a measure of the extent to which the interval E “fills up” the surface Σ .

Definition. A function $f : [0, \infty) \rightarrow [0, \infty)$ is a *regulating function* if for all $E \in \mathcal{H}_\Sigma$ we have $\text{length}(E) \leq f(\text{ext}(E))$.

Clearly if such a function exists, then Λ_Σ fills Σ . Conversely, if Λ_Σ fills Σ and there are no isolated leaves then it is not hard to see that Λ admits some regulating function. (We omit the proof since we shall not be requiring this fact.) However, by placing some restriction on the regulating function (for example by supposing that it only depends on certain parameters) we get a genuine geometric restriction on Λ_Σ . In particular we have the following:

Definition. We say that a lamination is *regular* if it admits a linear regulating function.

Suppose that Λ_Σ is regular with linear regulating function f . We present a series of lemmas aimed at showing that Λ_Σ is uniquely ergodic.

For the purposes of this section, we set $\epsilon = \text{inj}(\Sigma)$ and define a *transversal* to be a non-horizontal geodesic arc in Σ of length at most ϵ . If $E \in \mathcal{H}_\Sigma$, then a *diverting arc* for E is a transversal, α , which meets E precisely in its endpoints $\partial\alpha = \alpha \cap E$. We write E_α for the subinterval of E lying between these endpoints. Thus $\alpha \cup E_\alpha$ is an essential closed curve.

Lemma 13.1. *Let Λ_Σ be a regular lamination with regulating function, f . There is an integer $p \in \mathbb{N}$, depending only on f and $\text{inj}(\Sigma)$ such that if $E, F \in \mathcal{H}_\Sigma$ and α is a transversal with $\alpha \cap E = \emptyset$, then $\text{length}(E) \geq [n/p] \text{length}(F)$, where $n = |\alpha \cap F|$ and $[.]$ denotes integer part.*

Proof. Note that α cuts F into (at most) $n+1$ segments, F_0, F_1, \dots, F_n . For $1 \leq i \leq n-1$, let α_i be the subarc of α connecting the endpoints of F_i . Thus $\gamma_i = F_i \cup \alpha_i$ is a surrounding curve of E of length at most

$\text{length}(F_i) + \epsilon$ (though not necessarily embedded). By regularity, it follows that $\text{length}(F_i)$ is bounded below by an increasing linear function of $\text{length}(E)$. In particular, we can find constants $a, b > 0$ such that if $\text{length}(E) \geq a$, then $\text{length}(F_i) \geq b \text{length}(E)$.

Case (1), $\text{length}(E) \leq a$:

Since $\gamma_i = \alpha_i \cup F_i$ is essential, it must have length at least 2ϵ , and so $\text{length}(F_i) \geq \epsilon$. Thus $\text{length}(F) \geq (n-1)\epsilon \geq (\epsilon/a)(n-1) \text{length}(E)$, and so $\text{length}(F) \geq [n/p] \text{length}(E)$ for large enough p depending on ϵ and a and hence on $\text{inj}(\Sigma)$ and f as required.

Case (2), $\text{length}(E) \geq a$:

In this case, $\text{length}(F) \geq \sum_{i=1}^{n-1} \text{length}(F_i) \geq (n-1)b \text{length}(E)$ and so $\text{length}(F) \geq [n/p] \text{length}(E)$, where p depends on b and hence on f . \square

Lemma 13.2. *There is a constant $q \in \mathbb{N}$ depending on $\text{inj}(\Sigma)$ and f such that if $E, F \in \mathcal{H}_\Sigma$, α is a transversal with $\alpha \cap E = \emptyset$ and β is a diverting arc of E , then $|\beta \cap F| \geq [n/q]$, where $n = |\alpha \cap F|$.*

Proof. Let $m = |\beta \cap F|$. We can cut F into (at most) $m+1$ segments, F_0, \dots, F_m , not meeting β . Now $\beta \cup E_\beta$ is a surrounding arc for each F_i . Thus $\text{length}(F_i) \leq f(\text{length}(E) + \epsilon)$ so that $\text{length}(F) \leq (m+1)f(\text{length}(E) + \epsilon)$. But by Lemma 13.1 we have $\text{length}(F) \geq [n/p] \text{length}(E)$. Since f is linear, the result follows. \square

As an immediate corollary, we note:

Lemma 13.3. *If $E, F \in \mathcal{H}_\Sigma$ and α, β are diverting arcs of E and F respectively, then $\min\{|\alpha \cap F|, |\beta \cap E|\}$ is bounded above in terms of $\text{inj}(\Sigma)$ and f .*

Indeed, we can take the bound to be equal to q of Lemma 13.2.

Lemma 13.4. *There is a constant $N \geq 0$ such that if $E, F \in \mathcal{H}_\Sigma$ and α, β are transversals with $\min\{|\alpha \cap E|, |\beta \cap E|, |\alpha \cap F|, |\beta \cap F|\} \geq 2$, then*

$$\frac{|\alpha \cap E| |\beta \cap F|}{|\alpha \cap F| |\beta \cap E|} \leq N.$$

Proof. This is an easy consequence of Lemmas 13.3 and 13.4 considering the order of the points of intersection of E and F along each of the transversals α and β . \square

Now let Λ^0 be Λ minus the (finite) set of isolated leaves. Thus Λ^0 is a sublamination of Λ with no isolated leaves. Clearly Λ^0 is also regular

and thus fills Λ . In particular, its support is connected, and so there is no proper sublamination. Thus, Λ^0 is the support of any non-zero transverse measure to Λ . Moreover, it follows that every half-leaf of Λ^0 is dense. This allows us to define the “first return map” as follows. Let α be transversal with $\alpha \cap \Lambda^0 \neq \emptyset$ and with $\partial\alpha \cap \Lambda^0 = \emptyset$. Thus, K^0 is a Cantor set. We choose positive and negative sides of the transversal α . Thus, if $x^0 \in K^0$ is the intersection point of α and $\lambda \in \Lambda^0$, then we can define the positive and negative unit tangent vectors, x^+ and x^- to λ at x^0 . Let $K = K^+ \cup K^-$. Thus (as a subset of the unit tangent bundle to Σ) K is also a Cantor set. Let $\pi : K \rightarrow K^0$ be the projection map sending x^\pm to x^0 .

Let $f : K \rightarrow K$ be the first return map following Λ^0 . In other words, if $x^\pm \in K$, then $\pi \circ f(x^\pm)$ lies in the same leaf, $\lambda \in \Lambda^0$, as x^\pm . If $E(x^\pm, 1)$ is the subinterval of λ between $\pi(x^\pm)$ and $\pi \circ f(x^\pm)$, then $E(x^\pm, 1) \cap \alpha = \{\pi(x^\pm), \pi \circ f(x^\pm)\}$.

Given $x \in \mathbb{N}$, define $E(x^\pm, n)$ inductively by $E(x^\pm, n+1) = E(x^\pm, n) \cup E(f^n(x^\pm), 1)$. Thus $E(x^\pm, n)$ is a horizontal interval meeting α in the $n+1$ points $f^i(x^\pm)$ for $i = 0, \dots, n$. More generally, if $\beta \subseteq \alpha$ is a subarc, then $|\beta \cap E(x^\pm, n)| = |\{i \mid i \leq n, \pi \circ f^i(x^\pm) \in \beta\}|$. We write $A^\pm(x, \beta, n)$ for this quantity, where $x = \pi(x^\pm)$. If $\beta \cap \Lambda^0$ is infinite, then by the recurrence property we have $A^\pm(x, \beta, n) \rightarrow \infty$ as $n \rightarrow \infty$ for all $x \in \beta \cap \Lambda^0$. Reinterpreting Lemma 13.4 in these terms, we deduce:

Lemma 13.5. *Suppose that β_1, β_2 are subintervals of α , that $x_1, x_2 \in K^0$ and that $n \in \mathbb{N}$. If $A^+(x_i, \beta_j, n) \geq 2$ for all $i, j \in \{1, 2\}$, then*

$$\frac{A^+(x_1, \beta_1, n)A^+(x_2, \beta_2, n)}{A^+(x_1, \beta_2, n)A^+(x_2, \beta_1, n)} \leq N.$$

Note that any transverse measure on Λ^0 gives us a Borel measure, μ , on K^0 , which pulls back to an f -invariant measure μ' on K . Normalising we can suppose that μ has unit mass. If we have two distinct such measures, then applying the Hahn-Jordan decomposition theorem to their difference, we can suppose that they are mutually singular. Since they are assumed to be Borel measures, it follows that for any $M \geq 0$, we can find intervals $\beta_1, \beta_2 \subseteq \alpha$, each intersecting K^0 in infinite subsets, such that $\mu_1(\beta_1)\mu_2(\beta_2) \geq M\mu_1(\beta_2)\mu_2(\beta_1)$.

Proposition 13.6. *A regular lamination is uniquely ergodic.*

Proof. Suppose for contradiction that Λ is not uniquely ergodic. Then in the above notation, we can find measures μ_1, μ_2 on α and intervals $\beta_1, \beta_2 \subseteq \alpha$ such that $\mu_1(\beta_1)\mu_2(\beta_2) > N\mu_1(\beta_2)\mu_2(\beta_1)$. The pull backs of μ'_1 and μ'_2 to K are f -invariant. Thus, applying the Birkhoff ergodic

theorem, we can find $x_1, x_2 \in K^0$ such that $\frac{1}{n}A^+(x_i, \beta_j, n) \rightarrow \mu_i(\beta_j)$ as $n \rightarrow \infty$ for $i, j \in \{1, 2\}$.

Moreover, by the recurrence property, we know that $A^+(x_i, \beta_j, n) \geq 2$ for all sufficiently large n . This gives a contradiction to Lemma 13.5. \square

Next we give a slightly different formulation of regularity which is useful for applications such as its invariance under change of hyperbolic structure (Proposition 13.9). First we introduce the notion of “projected distances” which play a central role in the next few sections.

Given $x \in \mathbb{H}^2$ and $\lambda \in \Lambda$, we write $\pi_\lambda(x) \in \lambda$ for the projection of x to λ , i.e. the point of λ which minimises $d(x, \pi_\lambda(x))$. Given $x, y \in \mathbb{H}^2$, we write

$$\sigma_\lambda(x, y) = d(\pi_\lambda(x), \pi_\lambda(y))$$

and

$$\sigma(x, y) = \sigma_\Lambda(x, y) = \sup\{\sigma_\lambda(x, y) \mid \lambda \in \Lambda\}.$$

We note that σ is a metric on \mathbb{H}^2 with $\sigma \leq d$. It restricts to the standard metric on each leaf of Λ . Moreover, it is “monotone” in the sense that if $x, y, z \in \mathbb{H}^2$ with $z \in [x, y]$, then $\sigma(x, z) \leq \sigma(x, y)$.

Definition. We refer to $\sigma(x, y)$ as the *projected distance* between x and y with respect to Λ .

Suppose $Q \subseteq \mathbb{H}^2$ and $\lambda \in \Lambda$. We write $\text{pdiam}_\lambda(Q)$ for the diameter of $\pi_\lambda(Q)$, and $\text{pdiam}(Q) = \sup\{\text{pdiam}_\lambda(Q) \mid \lambda \in \Lambda\}$.

Definition. We refer to $\text{pdiam}(Q)$ as the *projected diameter* of Q with respect to Λ .

If $P, Q \subseteq \mathbb{H}^2$ and $P \cap Q \neq \emptyset$, then $\text{pdiam}(P \cup Q) \leq \text{pdiam}(P) + \text{pdiam}(Q)$. If $P \subseteq Q$, then $\text{pdiam}(P) \leq \text{pdiam}(Q)$.

Suppose that α is a path in Σ . We define $\text{pdiam}(\alpha) = \text{pdiam}(\tilde{\alpha})$, where $\tilde{\alpha}$ is any lift of α to \mathbb{H}^2 . Clearly $\text{pdiam}(\alpha) \leq \text{length}(\alpha)$. These measurements coincide for horizontal intervals. We also note that the projected diameter of a path is at least that of the geodesic connecting its endpoints in the same homotopy class.

Suppose that γ is a closed curve, and x is any point on γ . This determines a path, γ_x , which starts and ends at x following γ . Suppose y is another point on γ . By considering the projected lengths of the two paths into which x and y cut γ , we see easily that $\text{pdiam}(\gamma_y) \leq 2 \text{pdiam}(\gamma_x)$.

Definition. We define the *projected length*, $\text{plength}(\gamma)$, of γ as the maximum value of $\text{pdiam}(\gamma_x)$ as x runs around γ .

Clearly, $\text{plength}(\gamma) \leq \frac{1}{2} \text{length}(\gamma)$. We also note that $\text{plength}(\gamma) \leq 2 \text{pdiam}(\gamma_x)$ for any x in γ . As a consequence, we note:

Lemma 13.7. *If $\gamma = \alpha_1 \cup \dots \cup \alpha_n$, then $\text{plength}(\gamma) \leq 2 \sum_{i=1}^n \text{pdiam}(\alpha_i)$.*

Given a horizontal interval, $E \in \mathcal{H}_\Sigma$, we can define its *projected exterior length*, $\text{pext}(E)$, as the infimum of the projected lengths of (simple) surrounding curves of E . We can similarly define regulating functions with respect to projected length. In particular, we have:

Definition. We say that Λ_Σ is *strongly regular* if there is a linear function, $f : [0, \infty) \rightarrow [0, \infty)$ such that for all $E \in \mathcal{H}_\Sigma$, $\text{length}(E) \leq f(\text{pext}(E))$.

It is clear that strong regularity implies regularity. In fact, we shall see that they are equivalent (Proposition 13.11). First we show that strong regularity is independent of hyperbolic structure.

Suppose that $\Sigma' = \mathbb{H}^2/\Gamma'$ where $\Gamma' \cong \Gamma$ represents another marked hyperbolic structure. Let $\Lambda'_{\Sigma'}$ be the realisation of the abstract lamination underlying Λ_Σ . Suppose that $\phi : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ is an equivariant quasi-isometry. (This induces an equivariant homeomorphism $\phi : \partial\mathbb{H}^2 \rightarrow \partial\mathbb{H}^2$ sending the abstract lamination underlying Λ to that of Λ' .)

Lemma 13.8. *In the set-up described above, there is another equivariant quasi-isometry, $\psi : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ which maps \mathbb{H}^2 homeomorphically to \mathbb{H}^2 , and which maps each leaf of Λ homeomorphically to the corresponding leaf of Λ' . Moreover, ψ sends the centre of each complementary triangle of Λ to the centre of the corresponding complementary triangle of Λ' . Finally the parameters of ψ depend only on those of ϕ .*

Proof. We can forget about the group actions. Since the construction we give is canonical, it is automatically equivariant.

We first define $\psi|_{\bigcup \Lambda}$. Suppose $x \in \lambda \in \Lambda$ with $\lambda \cap \partial\mathbb{H}^2 = \{a, b\}$. Let $y_1, y_2 \in \partial\mathbb{H}^2$ be the points such that $[y_1, y_2]$ meets λ orthogonally in x . Thus, $\pi_\lambda(y_1) = \pi_\lambda(y_2) = x$, and x is a centre of both $\{a, b, y_1\}$ and $\{a, b, y_2\}$.

Let $\lambda' \in \Lambda'$ be the geodesic connecting $\phi(a)$ to $\phi(b)$. Let $\psi_i(x) = \pi_{\lambda'}(\phi(y_i))$, and let $\psi(x)$ be the midpoint of $[\psi_1(x), \psi_2(x)] \subseteq \lambda'$. Now, $\phi(x)$ and $\psi_1(x)$ are both centres of $\{\phi(a), \phi(b), \phi(y_1)\}$ and so $d(\phi(x), \psi_1(x))$ is uniformly bounded. The same goes for $d(\phi(x), \psi_2(x))$. It follows that $d(\phi(x), \psi(x))$ is uniformly bounded. Moreover, since ψ_1 and ψ_2 both map λ homeomorphically to λ' , we see that ψ does so also. It is easily seen that $\psi|_{\bigcup \Lambda}$ is a homeomorphism onto $\bigcup \Lambda'$.

We need to extend this map over each complementary triangle, T , of Λ . Let T' be the corresponding complementary triangle to Λ' . Let c and c' be the centres of T and T' respectively. Let ϵ be a sufficiently small uniform constant. If $\lambda, \mu \in \Lambda$ are edges of T , and $y \in \lambda, z \in \mu$, with $d(y, z) \leq \epsilon$ and with y, z lying in the same horocircle about $\lambda \cap \mu$, then we map the geodesic segment $[y, z]$ linearly to the segment $[\psi(y), \psi(z)]$. We do this for all such pairs, $\{y, z\}$. This defines ψ on all but a hexagonal subset, H , of T , whose boundary consists of three horizontal intervals of bounded length (depending on ϵ) alternating with three segments of length exactly ϵ . Note that c lies in the interior of H . We have already defined ψ on the boundary, $\text{fr}(H)$, of H . Again by choosing ϵ sufficiently small in relation to the parameters of ϕ , we can assume that $\psi(\text{fr}(H))$ encircles c and that there is a positive lower bound on $d(\psi(\text{fr}(H)), c')$. We now extend ψ over H by sending c to c' and coning linearly. \square

Note that for each $\lambda \in \Lambda$, the homeomorphism $\psi|_{\lambda} \rightarrow \psi(\lambda) \in \Lambda'$ is a uniform quasi-isometry of from $\psi|_{\lambda} \rightarrow \psi(\lambda) \in \Lambda'$. Thus ψ distorts projected distances by at most a linearly bounded amount, depending only on the parameters of ψ and hence of ϕ .

Proposition 13.9. *Suppose that Λ_{Σ} and Λ'_{Σ} are realisations of the same abstract lamination with respect to different hyperbolic structures. Then Λ_{Σ} is strongly regular if and only if Λ'_{Σ} is.*

Proof. Write Σ' for the second structure. By symmetry, we can suppose that Λ'_{Σ} is strongly regular. Let $\phi : \mathbb{H}^2 \rightarrow \mathbb{H}^2$ be any equivariant quasi-isometry (for example, the lift of any diffeomorphism from Σ to Σ' that respects markings) and let ψ be the quasi-isometry given by Lemma 13.8. Thus ψ descends to a homeomorphism from Σ to Σ' which we also denote by ψ .

Let $E \in \mathcal{H}_{\Sigma}$, and let γ be a surrounding curve whose projected length is arbitrarily close to $\text{pext}(E)$. Now $\psi(E)$ is a surrounding curve of $\psi(E)$ in Σ' , and so $\text{length}(\psi(E))$ is linearly bounded in terms of $\text{plength}(\psi(\gamma))$ and hence $\text{plength}(\gamma)$. Moreover, $\text{length}(E)$ is linearly bounded in terms of $\text{length}(\psi(E))$. It follows that Λ_{Σ} is strongly regular. \square

(Note that we cannot apply the above argument directly to regularity, since in general ψ will not be Lipschitz.)

We remark that the relation between the regulating functions of Λ_{Σ} and Λ'_{Σ} depend only on the parameters of ψ and hence of ϕ . Moreover, we can choose ϕ so that these parameters depend only on the Teichmüller distance between Σ and Σ' .

To show that regularity implies strong regularity (and for further applications later) we introduce the notion of a “decomposition” of Σ . We shall use a number of variants on this construction, but they all have in common certain features which we go on to describe.

By a *decomposition* of Σ , we shall mean a collection, \mathcal{P} , of open subsets of Σ which are precisely the complementary components of a finite embedded connected 1-complex, $\mathcal{G} = \mathcal{G}(\mathcal{P})$, in Σ , with respect to a fixed metric (here ρ_0). The 1-cells of (or “edges”) of \mathcal{G} partitioned into *horizontal* and *vertical* types. Each horizontal edge is a horizontal interval (i.e. contained in a leaf). Each vertical edge is an arc meeting Λ_Σ in its endpoints, not homotopic into a leaf relative to its endpoints. We shall also assume the vertical edges to be pairwise disjoint. Each vertex of \mathcal{G} is either of degree 3 (the endpoint of a vertical edge) or degree 1 (the endpoint of a horizontal edge). We refer to vertices of degree 1 as *vertical 0-cells*. We refer to \mathcal{G} as the *1-skeleton* of \mathcal{P} .

We refer to an element of \mathcal{P} as a *region*. A region is *rectangular* if it is simply connected, and its boundary contains precisely two vertical edges. We shall assume that no two rectangular regions meet along a vertical edge.

All the above conditions will be assumed to hold for any decomposition. We make a few observations.

The union of all horizontal edges is a disjoint union of horizontal intervals, which we refer to as the *base* of \mathcal{P} . We define the *complexity* of \mathcal{P} to be equal to the number of regions, $|\mathcal{P}|$. This is bounded in terms of $\text{genus}(\Sigma)$ and the number of components of the base (which is half the number of vertical 0-cells). Note that the complexity in turn gives a bound on the number of edges and vertices on the 1-skeleton.

The boundary of any region is a union of (not necessarily embedded) closed curves each consisting of an alternating sequence of vertical cells and horizontal sides — a *side* being a non-empty union of horizontal edges. (Note that every vertical 0-cell is terminal in the 1-skeleton, so the boundary curve doubles back along the same edge after passing through such a cell.)

If R is a rectangle, we shall refer to its length, $\text{length}(R)$, as the minimum of the lengths of its two horizontal sides.

Definition. We define the *girth* of \mathcal{P} as the maximal length of all edges other than those incident on rectangles on both sides (these might be the same rectangle).

We shall normally place some sort of uniform bound on $\text{girth}(\mathcal{P})$. Note that the total length of the 1-skeleton of $\mathcal{G}(\mathcal{P})$ is bounded in terms of $|\mathcal{P}|$, $\text{girth}(\mathcal{P})$ and the length of the longest rectangle.

By the *trivial* decomposition, we mean the decomposition consisting of a single horizontal edge, E , and hence just one complementary region $\Sigma \setminus E$. We can think of the boundary of this region as a closed path that encircles E .

Note that we can remove from a decomposition, \mathcal{P} , any set of rectangles and amalgamate the adjacent regions to form another decomposition, \mathcal{P}' . Note that \mathcal{P}' is “coarser” than \mathcal{P} in the sense that each element of \mathcal{P} is a subset of an element of \mathcal{P}' .

We next describe how to construct a decomposition, $\mathcal{P}(E)$, with base E with the property that each vertical edge is incident on a rectangle, and such that its girth is bounded above in terms of $\text{genus}(\Sigma)$. Note also that $|\mathcal{P}|$ is bounded above in terms of $\text{genus}(\Sigma)$. In all applications, E will be a horizontal interval of Λ .

To give the idea, imagine, for the moment, that E were a non-separating simple closed geodesic rather than a geodesic arc. We cut Σ along E so as to obtain a hyperbolic surface S with surface with totally geodesic boundary. Let $\epsilon > 0$ be sufficiently small, and let $S_0 \subseteq S$ be the set of points of S which lie on a geodesic segment of length at most ϵ connecting two points of the boundary, but which does not lie entirely in the boundary. Thus S_0 is a disjoint union of rectangular strips. We can take the set of all rectangles which arise in this way as the rectangular regions of a decomposition of Σ . The case where E is an interval requires a slight modification to take account of the endpoints of E but the argument is similar. (If an endpoint lies in a component of S_0 , then we remove a neighbourhood of this endpoint from S_0 so as to cut this component into three rectangles, separated by a small non-rectangular region containing a vertical 0-cell.) Note that all but a subset of E of bounded length lies in one of the horizontal edges of a rectangle. We shall denote the resulting decomposition by $\mathcal{P}(E)$.

For future reference, we note that we can perform a similar construction given a disjoint union of intervals, $F = F_1, \dots, F_n$. In this case, the resulting decomposition, $\mathcal{P}(F)$ will have complexity, $|\mathcal{P}(F)|$ bounded in terms of n and $\text{genus}(\Sigma)$.

Our first application of decompositions is the following:

Lemma 13.10. *Suppose that $E \in \mathcal{H}_\Sigma$. We can find a surrounding curve, γ , of E with $\text{length}(\gamma) - n \text{pext}(E)$ bounded in terms of $\text{genus}(\Sigma)$, and such that γ has the form $\gamma = F_1 \cup \alpha_1 \cup F_2 \cup \alpha_2 \cup \dots \cup F_n \cup \alpha_n$, where the F_i are disjoint geodesic segments, where each α_i is a path of length bounded in terms of $\text{genus}(\Sigma)$ and $\text{inj}(\Sigma)$, and where n is bounded in terms of $\text{genus}(\Sigma)$.*

If Λ_Σ is maximal, then we can take each of the F_i to be a horizontal interval.

We remark for future reference (see Lemma 14.7) that the same argument will show that we can also choose such a γ so that $\text{length}(\gamma) - \text{ext}(E)$ is bounded.

Proof. Let $\mathcal{P}(E)$ be the decomposition with base E constructed as above. We can assume that the length of each rectangle exceeds $\text{girth}(\mathcal{P}(E))$. Let β be a surrounding curve of projected length arbitrarily close to $\text{pext}(E)$. We perform a series of operations to obtain γ from β .

Suppose $R \in \mathcal{P}(E)$ is a rectangular region. We can suppose that each component of $\beta \cap R$ crosses R , for otherwise we could push it off R , decreasing the length of β . Moreover, we can suppose that there is at most one such component, for if there were two, we could connect them by a geodesic segment δ . Now δ cuts β into two closed curves, at least one of which must be essential. The component of its intersection with R that contains δ can be pushed off R so as to obtain a new closed curve of projected length no greater than the original. We perform these operations for each rectangular region. The resulting curve intersects the set of rectangles in arcs which can be adjusted to be geodesic, thus giving us our arcs F_1, F_2, \dots, F_n . If Λ is maximal, they can also be assumed to be horizontal. Note that $n \leq |\mathcal{P}(E)|$ is bounded in terms of $\text{genus}(\Sigma)$, and that we have increased the projected length of β by at most an amount bounded in terms of $\text{genus}(\Sigma)$.

Now the diameters of the rectangular regions are bounded in terms of $\text{genus}(\Sigma)$ and $\text{inj}(\Sigma)$. Thus, after shortcutting γ in similar fashion if necessary, we get a bound on the length of the connecting arcs $\alpha_1, \dots, \alpha_n$. Moreover, we can similarly arrange that γ be simple, as required.

Finally note that the lengths of each interval F_i approximately equals its projected length, and so $\text{length}(\gamma) - n \text{plength}(\gamma)$ is bounded. The result follows. \square

As an immediate consequence we deduce that $\text{ext}(E)$ is bounded by a linear function of $\text{pext}(E)$ depending only on $\text{genus}(\Sigma)$ and $\text{inj}(\Sigma)$. In particular:

Proposition 13.11. *A lamination is regular if and only if it is strongly regular.*

As usual we have a control over the respective regulating functions.

Finally we note that Propositions 13.9 and 13.11 together tell us that the regularity of a lamination does not depend on the hyperbolic structure.

14. STABLE SEQUENCES OF SURFACES

In this sections, we consider a sequence of hyperbolic metrics on a surface which have particular properties in relation to a given lamination (see Propositions 14.2 and 14.6). Under these hypotheses, we show that such a lamination is regular. We apply these results in later sections.

Suppose that $(\rho_i)_{i \in \mathbb{N}}$ is a sequence of hyperbolic metrics on a closed hyperbolic surface, Σ . We also write ρ_i for the induced metric on the universal cover, $X = \tilde{\Sigma}$, of Σ . Thus, $\Sigma = X/\Gamma$, and (X, ρ_i) is isometric to the \mathbb{H}^2 for all i . Suppose that Λ (respectively Λ_Σ) is simultaneously (the realisation of) a lamination with respect to each of the metric ρ_i . We make the following assumptions about the metric ρ_i .

(L1) The identity maps from (X, ρ_i) to (X, ρ_{i+1}) are uniform quasi-isometries

(L2) If $x, y \in \lambda \in \Lambda$, then the maps $[i \mapsto \rho_i(x, y)] : \mathbb{N} \longrightarrow [0, \infty)$ are bounded and uniformly quasiconvex.

(L3) The injectivity radii of the surfaces (Σ, ρ_i) are bounded below by a positive constant.

(L4) For all $l \geq 0$, there is some $p \in \mathbb{N}$ such that if $i, j \in \mathbb{N}$ and γ is a simple closed curve of Σ with $\max\{l_i(\gamma), l_j(\gamma)\} \leq l$, then $|i - j| \leq p$.

In (L4), $l_i(\gamma)$ denotes the length of the closed geodesic in (Σ, ρ_i) in the same homotopy class as γ . In the case of interest to us (where the distance function is quasiconvex) it is not hard to see that (L3) and (L4) are equivalent (see Proposition 17.3). Property (L3) is the “stability” condition on the surfaces — they project to a bounded set in moduli space. The intuitive interpretation of (L2) is that the sequence of metrics “collapses” in the direction of the lamination Λ — cf. the ending lamination of a 3-manifold.

We begin by deriving some consequences of (L1) and (L2), so for the moment, the group Γ is irrelevant.

Firstly, it will be convenient to replace (L2) by a stronger property, namely that the sequence $a_i = \rho_i(x, y)$ satisfies $a_{i+1} \leq ta_i + r$ for all i , where $t < 1$ and $r \geq 0$ are constants depending only on the parameters. This is justified by replacing ρ_i by ρ_{pi} for a sufficiently large integer $p \in \mathbb{N}$ (cf. Lemma 5.1). In what follows, we will say that a function $[i \mapsto a_i]$ is *linearly controlled* if $a_{i+1} \leq ta_i + r$ for all i , where $0 < t < 1$ and $r \geq 0$. We shall consider other functions satisfying this property, where we may need to modify r depending on the parameters. Note

that such a sequence decreases monotonically to at most some uniform constant, k , and then remains bounded by k .

Definition. A sequence of points $(x_i)_i$ in X form a *chain* if there is some $h \geq 0$ such that $\rho_i(x_i, x_{i+1}) \leq h$ for all i .

Such a chain thus has a “chain constant”, h , associated to it, which we can choose to be uniformly bounded and sufficiently large for what follows. (If the sequence of metrics arises from a stack, then this terminology is consistent with that previously introduced — see Section 17.) Note that, by (L1) this is equivalent to bounding $\rho_{i+1}(x_i, x_{i+1})$.

If $(x_i)_i$ and $(y_i)_i$ are both chain sequences in a leaf $\lambda \in \Lambda$, then the sequence $\rho_i(x_i, y_i)$ is also linearly controlled (at the cost of increasing r in relation to the chain constant). In particular (by setting $y_i = x_0$ for all i), we see that $\rho_i(x_0, x_i)$ is bounded above.

Given $x \in X$ and $\lambda \in \Lambda$, write $\pi_{\lambda,i}(x)$ for the projection to x to λ in (X, ρ_i) . By hypothesis (L1) we see that $(\pi_{\lambda,i}(x))_i$ is a uniform chain sequence. As a consequence, we see that the maps $[i \mapsto \sigma_{\lambda,i}(x, y)]$ and $[i \mapsto \sigma_i(x, y)]$ are linearly controlled where $\sigma_{\lambda,i}$ and σ_i denote projected distances with respect to λ and Λ in the metric ρ_i . Moreover, if $Q \subseteq X$, then $[i \mapsto \text{pdiam}_{\lambda,i}(Q)]$ and $[i \mapsto \text{pdiam}_i(Q)]$ are also linearly controlled.

We shall say that a sequence of horizontal intervals $E_i \subseteq \lambda$ form a *chain* if they have the form $E_i = [x_i, y_i]$ for chain sequences $(x_i)_i$ and $(y_i)_i$. Clearly $[i \mapsto \text{length}_i(E_i)]$ is linearly controlled, as is $[i \mapsto \rho_i(E_i, F_i)]$ where $(F_i)_i$ is any other chain of intervals.

Suppose that $\lambda, \mu \in \Lambda$ are distinct leaves. The intervals $E_i = \pi_{\lambda,i}(\mu)$ thus form a chain. In particular, if $x \in \lambda$, $\rho_i(x, E_i)$ is linearly controlled and is hence uniformly bounded in terms of $\rho_0(x, E_0)$. Moreover a positive lower bound on $\text{length}_i(E_i)$ gives an upper bound on $\rho_i(\lambda, \mu)$. Suppose that $z \in \mu$. Let $z_i = \pi_{\mu,i}(z)$. Then $(z_i)_i$ is a chain sequence, and so $\rho_i(z_i, z)$ is bounded above in terms of $\rho_0(z_0, z) \leq \rho_0(x, z)$. Since $\rho_i(x, z_i) \leq \rho_i(\lambda, \mu)$, we conclude that a lower bound on $\text{length}(E_i)$ places an upper bound on $\rho_i(x, z)$ in terms of $\rho_0(x, z)$. As a consequence, we have:

Lemma 14.1. *There are constants, r_1 and r_2 , depending on the parameters of (L1) and (L2) with the following property. Suppose $x, y \in \lambda$ and $z, w \in \mu$ where $\lambda, \mu \in \Lambda$ are distinct. If $\rho_0(x, z) \leq r_1$, $\rho_0(y, w) \leq r_1$ and $\rho_i(x, y) \geq 3r_1$, then $\rho_i(x, z) \leq r_2$ and $\rho_i(y, w) \leq r_2$.*

Proof. Let $E_i = \pi_{\lambda,i}(\mu)$. Note that $\rho_0(x, E_0) \leq \rho_0(x, \mu) \leq \rho_0(x, y) \leq r_1$. Thus $\rho_i(x, E_i) - r_1$ is bounded above in terms of the parameters of (L1) and (L2). The same applies to $\rho_i(y, E_i)$. Since $\rho_i(x, y) \geq 3r_1$, we

get a lower bound on $\text{length}_i(E_i) - r_1$. Thus by choosing r_1 sufficiently large in relation to the parameters, we get a uniform lower bound on $\text{length}_i(E_i)$. From the discussion preceding the lemma, we thus obtain an upper bound on $\rho_i(x, z)$ and $\rho_i(y, w)$ and in terms of r_1 . We denote this upper bound by r_2 . \square

This lemma has an interpretation in terms of the rectangular regions of a decomposition of Σ — while the horizontal edges remain sufficiently long, the lengths of the vertical edges remain bounded.

Before proving the main result of this section (Proposition 14.6), we start with a weaker version which will be sufficient for some application, in particular, those of Section 17. We remark that some analogous ideas can be found in [MasM].

Proposition 14.2. *Suppose that Λ and (Σ, ρ_i) satisfy hypotheses (L1)–(L4) above. Then, in any of the metrics ρ_i , Λ has a regulating function which depends only on the parameters of (L1)–(L4).*

Note that an immediate consequence is that Λ fills the surface Σ . To simplify notation, we prove this for the metric ρ_0 . In other words, we want to show that if $E \in \mathcal{H}_\Sigma$, and γ is a surrounding curve of E , then $\text{length}_0(E)$ is bounded in terms of $\text{length}_0(\gamma)$.

To this end, we introduce another measure on the size of an interval $E \in \mathcal{H}_\Sigma$. First, let r_0 be a constant so that $\text{length}_i(E)$ decreases monotonically to at most r_0 , and then remains bounded by r_0 . (Indeed for future applications, it will be convenient to assume that r_0 is sufficiently large so that the other linearly controlled functions discussed earlier also have this property.) We write:

$$h(E) = \min\{i \in \mathbb{N} \mid \text{length}_i(E) \leq r_0\}.$$

Thus if $i \geq h(E)$ then $\text{length}_i(E) \leq r_0$. By property (L1), $\text{length}_0(E)$ is bounded by an (exponential) function of $h(E)$. It therefore suffices to find an upper bound for $h(E)$.

First, we construct a decomposition, \mathcal{P}_0 , with base E , with the property that γ is entirely contained in some region, P_0 , of \mathcal{P}_0 , and such that $\text{girth}_0(\mathcal{P}_0)$ is bounded above by a (linear) function of $\text{length}_0(\gamma)$. To this end, we start with the decomposition, $\mathcal{P}(E)$ described in Section 13. We delete any rectangle which is completely crossed by a subarc of γ , and push γ off any other rectangles. In this process we may increase the girth of the decomposition, but only by an amount bounded in terms of $\text{length}(\gamma)$. Finally we throw away any rectangles whose ρ_0 -length is at most r_0 . (Recall that the length of a rectangle is defined as the minimum of the lengths of the two horizontal sides.) We also

assume that $r_0 \geq 3r_1$, where r_1 is the constant of Lemma 14.1. We denote the resulting decomposition by \mathcal{P}_0 .

Now given $i \in \mathbb{N}$, let \mathcal{P}_i be the decomposition obtained by throwing away from \mathcal{P}_0 all those rectangles whose ρ_i -length is at most r_0 , and amalgamating the adjacent non-rectangular regions. By the choice of r_0 , it follows that if $i \leq j$, then \mathcal{P}_j is coarser than \mathcal{P}_i . Moreover, by Lemma 14.1, we see that if we straighten out any vertical edge, i.e. replace it by the homotopic geodesic with the same endpoints, then its ρ_i -length will be bounded by r_2 . Let P_i be the region of \mathcal{P}_i containing P_0 and hence γ .

Let $\{h_1, h_2, \dots, h_n\} = \{h(R) \mid R \text{ is a rectangle in } \mathcal{P}_0\}$ with $h_1 < h_2 < \dots < h_n$. Here, $h(R)$ denotes $\min\{h(F_1), h(F_2)\}$ where F_1 and F_2 are the horizontal sides of R . Thus, the h_m are precisely the levels at which the combinatorics of \mathcal{P}_i changes. Note that n is bounded by the number of rectangles, and hence by $\text{genus}(\Sigma)$.

At each critical level $i = h_m$, the ρ_i -lengths of the horizontal edges of the non-rectangular regions may increase, since we are amalgamating a number of horizontal edges, some coming from old non-rectangular regions and others coming from rectangles we are deleting. However, the ρ_i -lengths of edges the latter type are uniformly bounded. Since we have a bound on the total number of regions, we have an inductive control on the ρ_i -lengths of edges of the former type. In summary, we have a bound, depending on l and the various parameters on the ρ_i -length of each horizontal edge of each non-rectangular region of \mathcal{P}_i . It follows that we have a bound on the ρ_i -length of the homotopic closed geodesic to each boundary curve of such a region. (First straighten out the vertical edges, and then homotop to a closed geodesic.)

Now let β_i be any boundary curve of P_i . We note:

Lemma 14.3. *If $i < h_n$, then β_i is essential in Σ . Moreover, $l_i(\beta_i)$ is uniformly bounded.*

Proof. The second statement is a consequence of the discussion immediately preceding the lemma. For the first statement, we push β_i slightly into P_i so that it becomes embedded but remains disjoint from γ . Suppose that it bounds a disc, D . Since γ is essential, it follows that $D \cap \gamma = \emptyset$, and so $E \subseteq D$. Since no vertical arc can be homotoped into E , it follows that there can be no vertical arcs. In other words, the decomposition is trivial, and so $i \geq h_n$. \square

Now hypothesis (L4) puts a bound on the number of consecutive i for which β_i can remain homotopically constant. Since it is determined by the combinatorics of \mathcal{P}_i , we deduce:

Lemma 14.4. *There is some $p \in \mathbb{N}$ such that for all m , $h_m - h_{m-1} \leq p$.*

Finally, we note:

Lemma 14.5. *There is some bound $q \in \mathbb{N}$ such that $h(E) - h_n \leq q$.*

Proof. By definition, \mathcal{P}_{h_n} is trivial, from which it follows that the ρ_{h_n} -length of E is bounded. (For example it is encircled by β_{h_n} .) Since $\rho_i(E)$ grows uniformly exponentially in $h(E) - i$, this places an upper bound on $h(E) - h_n$ as required. \square

We can now prove Proposition 14.2:

Proof. By Lemmas 14.4 and 14.5, we have $h(E) \leq np + q$. Here, n depends on $\text{genus}(\Sigma)$ and p and q depend on l the parameters of (L1)–(L4). This puts a bound on $\text{length}_0(E)$. \square

We can now sharpen this to give the principal result of this section:

Proposition 14.6. *Let Λ and (Σ, ρ_i) satisfy hypotheses (L1)–(L4) above. Then Λ is regular. In fact, there is a linear regulating function which depends only on the parameters of (L1)–(L4).*

To begin with we define the *size* $\text{size}(\mathcal{P})$ of a decomposition, \mathcal{P} , as the total length of the 1-skeleton, $\mathcal{G}(\mathcal{P})$. Note that for a given complexity, $|\mathcal{P}|$, and girth, the size of \mathcal{P} is bounded above by a fixed linear function of the length of the longest (horizontal) interval (or of the longest rectangle).

In the following lemma, size and girth are measured in the metric ρ_0 . (In fact, the lemma is valid for any lamination on a hyperbolic surface.)

Lemma 14.7. *Suppose that Λ is a maximal lamination and that $E \in \mathcal{H}_\Sigma$. Then there is a decomposition, \mathcal{P} , of Σ such that $|\mathcal{P}|$ is bounded in terms of $\text{genus}(\Sigma)$ and $\text{girth}(\mathcal{P})$ and $\text{size}(\mathcal{P}) - \text{ext}(\mathcal{P})$ are each bounded in terms of $\text{genus}(\Sigma)$ and $\text{inj}(\Sigma)$. No subarc of any vertical edge is homotopic into E relative to its endpoints. Moreover, the union of all edges not meeting E is essential in Σ (i.e. contains an essential closed curve).*

Proof. Let $\gamma = F_1 \cup \alpha_1 \cup \cdots \cup F_n \cup \alpha_n$ be a surrounding curve of E with $F_i \in \mathcal{H}_\Sigma$ and $\text{length}(\gamma) - \text{ext}(\gamma)$ bounded in terms of $\text{genus}(\Sigma)$ as given by the remark following Lemma 13.10. Let $\mathcal{P}(F)$ be the decomposition with base $F = \bigcup_{i=1}^n F_i$ described in Section 13. Thus, $|\mathcal{P}(F)|$ and $\text{size}(\mathcal{P}(F)) - \text{ext}(E)$ are bounded as required. Moreover, since the vertical arcs of $\mathcal{P}(F)$ are geodesic, the second requirement is automatically satisfied. To meet the final requirement, however, we will need to modify the set of vertical arcs as follows.

Firstly, we delete any rectangular region entirely crossed by a subarc of any α_i . Now the endpoints of each α_i lie in non-rectangular regions. Thus, if α_i intersects any rectangular region, then it enters and leaves by the same vertical edge, β . We can thus homotop β so that it no longer intersects α_i by following close to a subarc of α_i . We continue this process so as to get rid of all intersections of each α_i with each vertical edge. This increases the length of each vertical edge by a bounded amount. We now push the endpoints of each α_i slightly so that they meet F_i and F_{i+1} at interior points. We can now adjoin the α_i as vertical edges. Finally we delete all vertical edges that are incident on two rectangular regions. This gives us our decomposition, \mathcal{P} , and γ gives us the required essential closed curve in the 1-skeleton.

There is a technical detail in that it is conceivable that one of the arcs, α_i might be homotopic into a leaf $\lambda \in \Lambda$, contrary to one of the requirements of a decomposition. In this case, let $K \subseteq \lambda$ be the interval between F_i and F_{i+1} . If $E \cap K = \emptyset$, then we can replace, at the outset, the arc $F_i \cup \alpha_i \cup F_{i+1}$ by a single horizontal interval $F_i \cup K \cup F_{i+1}$. If $E \cap K$, then the length of E is bounded, and one can find an alternative construction, though this case will be of no interest to us in applications. \square

To apply this, we note:

Lemma 14.8. *Any lamination satisfying (L1)–(L4) can be extended to a maximal lamination satisfying (L1)–(L4).*

Proof. We know by Proposition 14.2 that such a lamination fills Σ . Thus each complementary region is a finite-sided polygon that can be subdivided into ideal triangles by adding a finite number of isolated leaves. Each such leaf is the union of two geodesic rays each asymptotic to a leaf of the original lamination. From this property (L2) follows easily. \square

Thus, without loss of generality, we can assume that Λ is maximal. Let \mathcal{P} be the decomposition given by Lemma 14.7. If $R \in \mathcal{P}$ is a rectangular region, let $h(R) = \min\{i \mid \text{length}_i(R) \leq r_0\}$. We can suppose that r_0 is sufficiently large so that if $i = h(R)$ and α is any arc crossing R , then $\text{length}_i(R) \geq r_0/2$.

Lemma 14.9. *Let \mathcal{P} be as above. There is some $p \geq 0$ depending on $\text{genus}(\Sigma)$ and the parameters of (L1)–(L4) such that E intersects each edge of \mathcal{P} in at most p points.*

Proof. Let $\mathcal{Q} \subseteq \mathcal{P}$ be the set of regions that intersect E non-trivially. We prove the lemma by induction on $|\mathcal{Q}|$.

Let Q be the closure of $\bigcup \mathcal{Q}$, and let $K_0 = K \cap Q$. We claim that K_0 contains a closed curve, β , which is essential in $\Sigma \setminus E$. To see this, note that if all of the components of $\Sigma \setminus Q$ are discs, then we can homotop an essential closed in K (as given by the hypotheses) into K_0 . On the other hand, if some such component is not a disc, then one of its boundary components has the required property.

Now let $R \in \mathcal{Q}$ be a rectangular region for which $i = h(R)$ is maximal among rectangles in \mathcal{Q} . Let $q > 0$ be the number of components of $E \cap R$. All but at most two of the components of $E \cap R$ cross R , and so $\text{length}_i(E) \geq (q - 2)r_0/2$.

On the other hand, the ρ_i -lengths of all the horizontal intervals of K_0 are uniformly bounded (given the choice of R). Moreover, the ρ_0 -lengths of all the vertical intervals are bounded by $\text{girth}(\mathcal{P})$. This puts a bound on their projected diameters with respect to ρ_i in terms of the parameters of (L1)–(L4). By Lemma 13.7, we therefore have a bound on the total projected length of the curve β . Since $\beta \cap E = \emptyset$, Proposition 14.2 and Lemma 13.10 place an upper bound on the ρ_i -length of E . Putting this together with the inequality of the previous paragraph, we get an upper bound on q .

We now replace \mathcal{Q} by $\mathcal{Q} \setminus \{R\}$ and apply the induction hypothesis to the (at most $q + 1$) components of $E \setminus R$ (pushed slightly off R). The result follows. \square

Corollary 14.10. *With the same hypotheses as Lemma 14.9, the $\text{length}_0(E)$ is bounded by a linear function of $\text{size}_0(\mathcal{P})$, which depends on $|\mathcal{P}|$, $\text{genus}(\Sigma)$ and the parameters of (L1)–(L4).*

Proof. Let $P \in \mathcal{P}$ and F be a component of $P \cap E$. If P is non-rectangular, we see that $\text{length}_0(F)$ is uniformly bounded, whereas if P is rectangular, then $\text{length}_0(F)$ minus the length of P is similarly bounded. Clearly the length of each rectangle is bounded by $\text{size}_0(\mathcal{P})$. The result follows applying Lemma 14.9. \square

We can now prove Proposition 14.6:

Proof. By Lemma 14.8, we can assume that Λ is maximal. We construct \mathcal{P} as in Lemma 14.7 and apply Corollary 14.10. \square

15. OTHER PROPERTIES OF LAMINATIONS

In this section, we consider some further, more technical, properties of laminations. In particular, we (Lemma 15.7) show that for surfaces with a lower injectivity radius bound, a control on projected distances is sufficient to produce a uniform quasi-isometry between the universal

covers. This will be needed for the main result (Theorem 17.1) of Section 17. This is, in turn, key to Theorem 2.11.

Let $\Sigma = \mathbb{H}^2/\Gamma$ be a closed hyperbolic surface with a positive lower bound on injectivity radius, $\text{inj}(\Sigma) \geq \text{inj}_0 > 0$. This places an upper bound, d_0 , on the diameter, $\text{diam}(\Sigma)$, of Σ . Thus, the Γ -orbit, Γx , of any point, $x \in \mathbb{H}^2$ is d_0 -dense. We write d for both the metric on Σ and that on \mathbb{H}^2 .

If $\alpha = [x, y] \subseteq \Sigma$ is a geodesic segment and $0 \leq r \leq \frac{1}{2}d(x, y)$, we write $\text{mid}(\alpha, r) = [x', y']$ where $x', y' \in \alpha$ satisfy $d(x, x') = d(y, y') = r$. We can interpret the projected diameter, $\text{pdiam}(\alpha) = \sigma(x, y)$ as bounding the distance over which α can remain close to a leaf of Λ . More precisely, we have:

Lemma 15.1. *Given $k \geq 0$ there exist $\delta_1, \delta_2 > 0$ such that if $\alpha \subseteq \mathbb{H}^2$ is a geodesic segment with $\text{length}(\alpha) \geq 2k$ and $\text{pdiam}(\alpha) \leq k$, then any leaf of Λ that intersects $\text{int}(\alpha)$ does so at an angle at least δ_1 . Moreover, if $\lambda \in \Lambda$ satisfies $d(\lambda, \text{mid}(\alpha, k)) \leq \delta_2$, then $\lambda \cap \text{int}(\alpha) \neq \emptyset$.*

We deduce:

Lemma 15.2. *Given $k, \delta > 0$, there is some $\eta = \eta(k, \delta) > 0$ such that if $\alpha \subseteq \mathbb{H}^2$ is a geodesic segment with $\text{length}(\alpha) \geq 2k$ and $\text{pdiam}(\alpha) \leq k$, and $x \in \mathbb{H}^2 \setminus \bigcup \Lambda$ with $d(x, \text{mid}(\alpha, r)) \leq \eta$, then x can be connected to $\text{int}(\alpha)$ by a path, β , with $\text{length}(\beta) \leq \delta$ and $\beta \cap \bigcup \Lambda = \emptyset$.*

Proof. Let $\delta_1, \delta_2 > 0$ be the constants given by Lemma 15.1. We can assume that $\eta < \delta_2$. Suppose that x, α are as in the hypotheses. Choose $y \in \text{mid}(\alpha, k)$ with $d(x, y) \leq \eta$. If $[x, y] \cap \bigcup \Lambda = \emptyset$, we are done by setting $\eta \leq \delta$. If not, let λ be the first leaf of Λ that intersects $[x, y]$ starting from x . Let $\lambda \cap [x, y] = \{z\}$. Now, $d(\lambda, y) \leq \eta \leq \delta_2$, so by Lemma 15.1, λ meets $\text{int}(\alpha)$ at some point w . The angle of intersection of α and λ is bounded below by δ_1 . Considering the triangle xyw , we see that there is an upper bound on $d(w, z)$ which can be made arbitrarily small by ensuring that $d(x, z) \leq \eta$ is sufficiently small. In particular, we can choose $\eta \leq \delta/3$ sufficiently small so that $d(w, z) \leq \delta/3$. Now connect x to $\text{int}(\alpha)$ by a path β that first runs along $[x, y]$ nearly to z , and then follows close to λ in $\mathbb{H}^2 \setminus \bigcup \Lambda$ until it hits $\text{int}(\alpha)$ near w . \square

Suppose now that Λ is maximal. Each complementary region, T , is an ideal triangle, so we can associate to it its centre of symmetry, $\text{cent}(T)$. Let \mathcal{C} be the set of all such centres. Note that \mathcal{C} is a finite union of Γ -orbits, and is hence d_0 -dense in \mathbb{H}^2 . Moreover, it is d_1 -separated where $d_1 = \log(3)$ is a universal constant. Thus, \mathcal{C} is a (uniform) net in \mathbb{H}^2 , as defined in Section 3.

Lemma 15.3. *Given $\zeta_1 > 0$, there is some $\zeta_2 > 0$ such that if $x \in \mathbb{H}^2$ then there is some $y \in \mathbb{H}^2$ with $d(x, y) \leq \zeta_1$ and $d(y, \bigcup \Lambda) \geq \zeta_2$.*

Proof. Suppose, to the contrary, that there is some ζ_1 -disc, $N = N(x, \zeta_1)$, such that $N \cap (\bigcup \Lambda)$ is ζ -dense in N for $\zeta > 0$ arbitrarily small. There are two leaves λ, μ intersecting N such that all other leaves intersecting N lie between λ and μ . Now $d(\lambda \cap N, \mu \cap N)$ is bounded below, so we can find some $z \in \mathbb{H}^2$ a bounded distance from x such that $N' = N(x, d_0)$ lies in the region of \mathbb{H}^2 bounded by $\lambda \cup \mu$. By choosing ζ sufficiently small, we can arrange that $N' \cap (\bigcup \Lambda)$ is arbitrarily quasidense in N' . In particular, we can arrange that N' contains no centre of any complementary triangle. This contradicts the fact that \mathcal{C} is d_0 -dense. \square

Lemma 15.4. *Given $u > 0$, there is some $v = v(u)$, such that if $\alpha \subseteq \mathbb{H}^2$ is a geodesic segment with $\text{pdiam}(\alpha) \leq u$ and $\text{length}(\alpha) \geq 2u$, then for all $x \in \alpha$, there is some $y \in \mathcal{C}$ with $d(x, y) \leq v$ and such that the complementary triangle containing y intersects α .*

Proof. Choose any $\delta > 0$, and let $\eta = \eta(u, \delta)$ be the constant of Lemma 15.2. Let ζ_2 be the constant of Lemma 15.3 given $\zeta_1 = \eta$.

Now suppose that $x \in \alpha$. There is some $z \in \text{mid}(\alpha, u)$ with $\rho(x, z) \leq u$. By Lemma 15.3, there is some $w \in \mathbb{H}^2$ with $d(w, z) \leq \eta = \zeta_1$ and with $\rho(w, \bigcup \Lambda) \geq \zeta_2$. By Lemma 15.2, we can connect w to α by a path, β (with $\text{length}(\beta) \leq \delta$) so that $\beta \cap (\bigcup \Lambda) = \emptyset$. Let T be the complementary triangle containing w and let $y = \text{cent}(T)$. Since $\rho(y, \bigcup \Lambda) \geq \zeta_2$, we see that $d(w, y)$ and hence $d(x, y) \leq d(x, z) + d(z, w) + d(w, y)$ is bounded above. Moreover, since $\beta \subseteq T$, we see that $T \cap \alpha \neq \emptyset$ as required. \square

If $x, y \in \mathcal{C}$ let $\langle x, y \rangle \subseteq \mathcal{C}$ be the set of centres of triangles that intersect the geodesic segment $[x, y] \subseteq \mathbb{H}^2$. Note that this is determined entirely by the combinatorics of the lamination Λ .

Given $t \geq 0$, let $B_t(x, y) = \{z \in \langle x, y \rangle \mid \sigma(z, \{x, y\}) \leq t\}$ and let $\nu_t(x, y) = |B_t(x, y)|$.

Lemma 15.5. *Given $u \geq 0$, there is some $s \geq 0$ such that for all $n \geq 0$, there is some $l \geq 0$ such that for all $x, y \in \mathcal{C}$ with $\sigma(x, y) \leq u$ and $\nu_s(x, y) \leq n$, we have $d(x, y) \leq l$.*

Proof. We can suppose that $l \geq 2u$. Let $v = v(u)$ be as given by Lemma 15.4. Suppose that $\rho(x, y) \geq 2u$. Every point $w \in [x, y]$ is a distance at most v from a point $z \in \langle x, y \rangle$. Now $\sigma(w, z) \leq d(w, z) \leq v$ and so $\sigma(x, z) \leq \sigma(x, w) + \sigma(w, z) \leq \sigma(x, y) + \sigma(w, z) \leq u + v$. Thus, setting $s = u + v$, we see that $z \in B_s(x, y)$. By hypothesis, there are

at most $\nu_s(x, y) \leq n$ such points z , which gives us an upper bound on the length of the segment $[x, y]$ in terms of u and n as required. \square

Lemma 15.6. *Given $r, t \geq 0$, there is some $n \in \mathbb{N}$ such that if $x, y \in \mathcal{C}$ with $d(x, y) \leq r$, then $\nu_t(x, y) \leq n$.*

Proof. Suppose $z \in B_t(x, y)$. Thus z is the centre of a complementary triangle, T , with $T \cap [x, y] \neq \emptyset$. There is an edge λ of T with $\lambda \cap [x, y] \neq \emptyset$. Now $\sigma_\lambda(z, \{x, y\}) \leq \sigma(z, \{x, y\}) \leq t$. Since z is a bounded distance from λ , this places a uniform upper bound on $d(z, [x, y])$, and hence an upper bound on $d(x, z)$ in terms of r . Since \mathcal{C} is uniformly separated, this bounds $\nu_t(x, y) = |B_t(x, y)|$. \square

Suppose now that Σ' is another hyperbolic surface homeomorphic to Σ . We assume that Σ and Σ' are “marked” surfaces, so that there is a natural isomorphism, $\theta : \Gamma \rightarrow \Gamma'$, where $\Gamma' = \pi_1(\Sigma')$. Thus Γ' acts on \mathbb{H}^2 with $\Sigma' = \mathbb{H}^2/\Gamma'$. We shall assume that $\text{inj}(\Sigma') \geq \text{inj}_0$. There is a unique homeomorphism, $\psi : \partial\mathbb{H}^2 \rightarrow \partial\mathbb{H}^2$ which is equivariant with respect to the actions of Γ and Γ' via the isomorphism θ . Let Λ' be the (realisation of) the image of the lamination Λ under that map ψ . Thus Λ' descends to a lamination, $\Lambda'_{\Sigma'}$, on Σ' . Let \mathcal{C} be the set of centres of complementary triangles of Λ' . There is a canonical bijection from \mathcal{C} to \mathcal{C}' which we also denote by ψ . Let σ' denote projected distance on \mathbb{H}^2 with respect to Λ' .

Lemma 15.7. *Suppose there is an increasing function, $f : [0, \infty) \rightarrow [0, \infty)$ such that for all $x, y \in \mathcal{C}$, we have*

$$\sigma'(\psi(x), \psi(y)) \leq f(\sigma(x, y))$$

and

$$\sigma(x, y) \leq f(\sigma'(\psi(x), \psi(y))).$$

Then the map $\psi : \mathcal{C} \rightarrow \mathcal{C}'$ is bilipschitz (with respect to the hyperbolic metric d) with constants depending only on f , $\text{genus}(\Sigma)$ and inj_0 .

Proof. We know that \mathcal{C} and \mathcal{C}' are d_0 -dense, where d_0 depends only on $\text{genus}(\Sigma)$ and inj_0 . Moreover, since the situation is symmetrical with respect to interchanging Σ and Σ' , it suffices to place an upper bound on $d(\psi(x), \psi(y))$ for $x, y \in \mathcal{C}$ with $d(x, y) \leq 2d_0$.

Let $r = 2d_0$. Let s be the constant of Lemma 15.5 given $u = f(r)$, and let $t = f(s)$.

Suppose $x, y \in \mathcal{C}$ with $d(x, y) \leq r$. By Lemma 15.6 (applied to Σ), there is a bound, n , on $\nu_t(x, y)$ depending on r and t .

We claim that $B_s(\psi(x), \psi(y)) \subseteq \psi(B_t(x, y))$. To see this, suppose $z \in \mathcal{C}$ with $\psi(z) \in B_s(\psi(x), \psi(y))$. Thus $\psi(z) \in \langle \psi(x), \psi(y) \rangle$ and so

$z \in \langle x, y \rangle$. Moreover, $\sigma'(\psi(z), \{\psi(x), \psi(y)\}) \leq s$, and so $\sigma(z, \{x, y\}) \leq f(s) \leq t$. Thus $z \in B_t(x, y)$ so that $\psi(z) \in \psi(B_t(x, y))$ as claimed. Since ψ is a bijection, it follows that $\nu_s(\psi(x), \psi(y)) \leq \nu_t(x, y) \leq n$.

Now $\sigma'(\psi(x), \psi(y)) \leq f(\sigma(x, y)) \leq f(r) \leq u$. Thus, by Lemma 15.5 (applied to Σ') we see that $d(\psi(x), \psi(y)) \leq l$, where l depends on n and u , and hence ultimately on f , $\text{genus}(\Sigma)$ and inj_0 . \square

Note that \mathcal{C} and \mathcal{C}' are both uniform nets in \mathbb{H}^2 . We therefore deduce that $\psi : \mathcal{C} \rightarrow \mathcal{C}'$ is (or determines) an equivariant quasi-isometry from \mathbb{H}^2 to \mathbb{H}^2 whose parameters depend only on f , $\text{genus}(\Sigma)$ and inj_0 .

In summary, we have shown that a uniform control on projected distances gives us a uniform equivariant quasi-isometry.

16. STRATIFIED QUASI-ISOMETRIES

In this section, we return briefly to hyperbolic stacks, and define the notion of a ‘‘stratified quasi-isometry’’, which can be thought of as a quasi-isometry that sends sheets close to sheets. We begin by giving a more formal definition of this notion.

Suppose that (Ξ, ρ) and (Ξ', ρ') are hyperbolic stacks, with sheets $(X_i)_{i \in \mathcal{I}}$ and $(X'_j)_{j \in \mathcal{I}'}$ respectively. Here, the indexing sets \mathcal{I} and \mathcal{I}' have the form $\mathcal{I} = I \cap \mathbb{Z}$ and $\mathcal{I}' = J \cap \mathbb{Z}$ where I and J are subintervals of the real line. The sheets X_i and X'_j are assumed to have the induced path metrics ρ_i and ρ'_j respectively.

Suppose that $\sim \subseteq \Xi \times \Xi'$ is a relation between Ξ and Ξ' . Given $i \in \mathcal{I}$ and $j \in \mathcal{I}'$, we write $\sim_{i,j} = \sim \cap (X_i \times X'_j)$ for its restriction to a relation between X_i and X'_j . If we fix a set of parameters, we can define a relation \approx between \mathcal{I} and \mathcal{I}' by writing $i \approx j$ if $\sim_{i,j}$ is a quasi-isometry between X_i and X'_j with these parameters.

Suppose that $\psi : \Xi \rightarrow \Xi'$ is a quasi-isometry, formally defined as a relation, $\sim_\psi \subseteq \Xi \times \Xi'$.

Definition. We say that ψ is a *stratified quasi-isometry* if the relation \approx between \mathcal{I} and \mathcal{I}' defined as above is a quasi-isometry from \mathcal{I} to \mathcal{I}' (with their standard metrics as subsets of \mathbb{R}).

Given $i \in \mathcal{I}$ and $j \in \mathcal{I}'$ with $i \approx j$, we write $\psi_{i,j}$ for the quasi-isometry from X_i to X'_j given by the hypotheses.

It is easily seen that any weakening of a stratified quasi-isometry (as defined in Section 3) is a stratified quasi-isometry. Indeed, we can weaken any stratified quasi-isometry so that each point of \mathcal{I} is \approx -related to a point of \mathcal{I}' and vice versa.

We say that two stacks are *stratified quasi-isometric* if there is a stratified quasi-isometry between them. It is readily seen that this

relation is symmetric and transitive, and thus defines an equivalence between stacks.

A stratified quasi-isometry, $\psi : \Xi \longrightarrow \Xi'$ induces a homeomorphism from $\partial\Xi$ to $\partial\Xi'$ which we also denote by ψ . We can also construct a homeomorphism $\psi^0 : \partial^0\Xi \longrightarrow \partial^0\Xi'$ as follows.

Suppose $i, j \in \mathcal{I}$ and $i', j' \in \mathcal{I}'$ with $i \approx i'$ and $j \approx j'$. Let $\psi_1 = \psi_{i,i'} : X_i \longrightarrow X_{i'}$ and $\psi_2 = \psi_{j,j'} : X_j \longrightarrow X_{j'}$. There are natural quasi-isometries $\phi : X_i \longrightarrow X_j$ and $\phi' : X_{i'} \longrightarrow X_{j'}$. Recall this means that $\rho(x, y) \leq r$ whenever $x \sim_\phi y$ and $\rho'(x', y') \leq r'$ whenever $x' \sim_{\phi'} y'$, where r and r' are constants depending on the parameters, and, respectively $|i - j|$ and $|i' - j'|$. By definition, these natural quasi-isometries extend to the identities on $\partial^0\Xi$ and $\partial^0\Xi'$ via the identifications $\partial X_i \equiv \partial X_j \equiv \partial^0\Xi$ and $\partial X_{i'} \equiv \partial X_{j'} \equiv \partial^0\Xi'$.

Suppose that $x \in X_i$, $y \in X_j$, $x' \in X_{i'}$, $y', z \in X_{j'}$ with $x \sim_\phi y$, $x' \sim_{\phi'} y'$, $x \sim_\psi x'$ and $y \sim_\psi z$. Now $\rho(x, y) \leq r$ and $\rho'(x', y') \leq r'$. Since ψ is a quasi-isometry from Ξ to Ξ' , we see that $\rho'(x', z)$ is bounded in terms of r . Thus, $\rho(y', z)$ and hence $\rho_{j'}(y', z)$ is uniformly bounded. This means that the quasi-isometries $\phi' \circ \psi_1$ and $\psi_2 \circ \phi$ from X_i to $X_{j'}$ agree up to a bounded distance and hence induce the same map on boundaries. Put another way, up to the standard identifications, the maps ψ_1 and ψ_2 extend to identical maps from $\partial^0\Xi$ to $\partial^0\Xi'$.

In summary, we conclude:

Lemma 16.1. *If $\psi : \Xi \longrightarrow \Xi'$ is a stratified quasi-isometry, then there is a natural map $\psi^0 : \partial^0\Xi \longrightarrow \partial^0\Xi'$, such that for all $i \in \mathcal{I}$ and $j \in \mathcal{I}'$ with $i \approx j$, the quasi-isometry $\psi_{i,j} : X_i \longrightarrow X_j$ extends to the map $\phi^0 : \partial X_i \longrightarrow \partial X_j$ under the identifications $\partial X_i \equiv \partial^0\Xi$ and $\partial X_j \equiv \partial^0\Xi'$.*

Moreover, it is easily verified that these maps commute with the Cannon-Thurston maps. In other words, $\psi \circ \omega = \omega' \circ \psi^0$, where $\omega : \partial^0\Xi \longrightarrow \partial\Xi$ and $\omega' : \partial^0\Xi' \longrightarrow \partial\Xi'$ are the Cannon-Thurston maps.

Definition. We say that a stack is “taut” if its sheets are uniformly taut.

In other words every point of each sheet is a uniform distance from the centre of some three ideal points.

For taut stacks, the notion of a stratified quasi-isometry can be characterised sheetwise as follows. Suppose $\sim \subseteq (\bigcup_i X_i) \times (\bigcup_j X_j)$ is a relation between Ξ and Ξ' . We define a relation between \mathcal{I} and \mathcal{I}' by writing $i \approx j$ if $\sim_{i,j} \neq \emptyset$.

Lemma 16.2. *Suppose that \approx defined as above is a quasi-isometry between \mathcal{I} and \mathcal{I}' . Suppose also that there is a bijection, $\psi^0 : \partial^0\Xi \longrightarrow$*

$\partial^0 \Xi'$ such that for all $i \in \mathcal{I}$ and all $j \in \mathcal{I}'$ the relation $\sim_{i,j}$ is a quasi-isometry from X_i to X'_j , which induces the map $\psi^0 : \partial X_i \rightarrow \partial X'_j$, under the identifications $\partial^0 \Xi \equiv \partial X_i$ and $\partial^0 \Xi' \equiv \partial X'_j$. Then ψ is a stratified quasi-isometry.

Proof. We need only check that \sim is a quasi-isometry. For this in turn it suffices (by symmetry) to check that for all $k \geq 0$, there is some $k' \geq 0$ such that if $x, y \in \bigcup_i X_i$ and $x', y' \in \bigcup_j X'_j$ with $x \sim x'$, $y \in y'$ and $\rho(x, y) \leq k$ then $\rho'(x', y') \leq k'$. Let $x \in X_i$, $y \in X_j$, $x' \in X'_i$ and $y' \in X'_j$. Let $\phi : X_i \rightarrow X_j$ and $\phi' : X'_i \rightarrow X'_j$ be natural quasi-isometries. Thus, we can find constants r, r' depending only on $|i - j|$ and $|i' - j'|$ and hence ultimately only on k such that there are some $z \in X_j$ and $z' \in X'_j$, so that $x \sim_\phi z$ and $x' \sim_{\phi'} z'$ and $\rho(x, z) \leq r$ and $\rho'(y', z') \leq r'$. Let $\psi_1 : X_i \rightarrow X'_i$ and $\psi_2 : X_j \rightarrow X'_j$ be the quasi-isometries given by the hypotheses. Thus $\phi' \circ \psi_1 \circ \phi^{-1} : X_j \rightarrow X'_j$ is a quasi-isometry whose parameters depend only on k . Moreover, by hypothesis, ψ_2 and $\phi' \circ \psi_1 \circ \phi^{-1}$ induce the same map, namely ψ^0 , on the boundaries. Since X_j and X'_j are uniformly taut, it follows that they agree up to bounded distance depending on k (Lemma 4.8). Now, by definition of ψ_1 and ψ_2 , y is related to y' under ψ_2 , and z is related to z' under $\phi' \circ \psi_1 \circ \phi^{-1}$. Since $\rho(y, z) \leq \rho(x, y) + \rho(x, z) \leq k + r$ is bounded, $\rho_j(y, z)$ is bounded. Thus $\rho'(y', z') \leq \rho'_j(y', z')$ and hence $\rho'(x', y') \leq r' + \rho'(y', z')$ is bounded as required. \square

Reinterpreting all this in more familiar terms, suppose that $f : \mathcal{I} \rightarrow \mathcal{I}'$ is a quasi-isometry and that for all $i \in \mathcal{I}$, we have a uniform quasi-isometry from X_i to $X'_{f(i)}$. Suppose all these quasi-isometries induce the same map from $\partial^0 \Xi$ to $\partial^0 \Xi'$. Then we can assemble these quasi-isometries sheetwise to give us a stratified quasi-isometry from Ξ to Ξ' . Indeed, it suffices to define f on some quasidense subset of \mathcal{I} .

We finally mention the following easily verified result. Since we shall not be using it directly, we omit the proof.

Proposition 16.3. *Suppose that a group Γ acts by sheet preserving isometry on the stacks Ξ and Ξ' . Suppose that the diameters of the quotients X_i/Γ and X'_i/Γ are uniformly bounded. Then any equivariant quasi-isometry from Ξ to Ξ' is a stratified quasi-isometry.*

17. HYPERBOLIC SURFACE STACKS

In this section, we consider stacks of hyperbolic surfaces with a positive lower bound on injectivity radii. More formally we shall be dealing with stacks of hyperbolic planes admitting a cocompact action of a surface group. We shall eventually see that such stacks are hyperbolic if

and only if the surfaces stay close to a geodesic in Teichmüller space (Section 20). In this section, we take hyperbolicity (or equivalently, quasiconvexity of the distance function) as a hypothesis. This allows us to define the “ending lamination” of such a stack. From the results of Section 14, such a lamination is regular. The principal result (Theorem 17.1) will be that a hyperbolic surface stack with lower injectivity radius bound is determined up to equivariant quasi-isometry by its ending lamination.

Let Σ be a closed topological surface, and let $\Gamma = \pi_1(\Sigma)$.

Definition. A *hyperbolic surface stack*, Ξ_Σ , consists of a proper hyperbolic stack, Ξ , of hyperbolic planes, $X_i \cong \mathbb{H}^2$, together with a sheet-preserving isometric action of Γ , such that the induced action on each sheet is properly discontinuous and cocompact.

Thus, for each i , $\Sigma_i = X_i/\Gamma$ is a compact hyperbolic surface. For each i , the identification of $\pi_1(\Sigma_i) \equiv \Gamma \equiv \pi_1(\Sigma)$ gives a natural homotopy class of homeomorphism from Σ to Σ_i . We can thus view the Σ_i as “marked” hyperbolic surfaces. Note that any homotopy equivalence from Σ_i to Σ_j that respects this marking lifts to a quasi-isometry from X_i to X_j . This quasi-isometry extends to the natural identification of ∂X_i and ∂X_j with $\partial^0 \Xi$. We shall use the same notation, ρ_i , to denote the geodesic metric on either X_i or Σ_i .

Suppose that Ξ_Σ is a semi-infinite hyperbolic surface stack. Fix, for the moment, some $i \in \mathbb{N}$. We identify $X_i \equiv \mathbb{H}^2$ and $\partial X_i \equiv \partial^0 \Xi$. Proposition 12.3 now gives us a well defined abstract lamination on $\partial^0 \Xi$ which can be realised as a lamination on X_i . The latter descends to a (geometric) lamination, Λ_{Σ_i} , on Σ_i . This pulls back, via the marking of Σ_i to a lamination, Λ_Σ , on Σ (viewed as a topological object). Since everything is canonical, this is independent of the choice of $i \in \mathbb{N}$.

Definition. We refer to Λ_Σ as the *ending lamination* of Ξ_Σ .

Definition. The *injectivity radius*, $\text{inj}(\Xi_\Sigma)$, of the stack Ξ_Σ is defined as $\inf\{\text{inj}(\Sigma_i) \mid i \in \mathbb{N}\}$.

We shall see that the ending lamination is regular if $\text{inj}(\Xi_\Sigma) > 0$. There is a kind of converse, described in Section 18.

Suppose that Ξ_Σ and Ξ'_Σ are semi-infinite hyperbolic surface stacks of the same topological type. Suppose that $f_0 : \Sigma_0 \rightarrow \Sigma'_0$ is a homotopy equivalence that respects the markings. Thus, f_0 lifts to a quasi-isometry $\psi_0 : X_0 \rightarrow X'_0$, which in turn extends to an equivariant homeomorphism $\psi_0 : \partial^0 \Xi \rightarrow \partial^0 \Xi'$. The main result of this section gives a condition under which ψ_0 extends to a quasi-isometry.

Theorem 17.1. *Suppose $\text{inj}_0 > 0$. Let Ξ_Σ and Ξ'_Σ be hyperbolic surface stacks with $\text{inj}(\Xi_\Sigma) \geq \text{inj}_0$ and $\text{inj}(\Xi'_\Sigma) \geq \text{inj}_0$. Suppose that $\psi_0 : \Xi_\Sigma \rightarrow \Xi'_\Sigma$ is an equivariant quasi-isometry that sends the ending lamination of Ξ_Σ to that of Ξ'_Σ . Then, ψ_0 extends to an equivariant quasi-isometry $\Xi \rightarrow \Xi'$. Moreover, the parameters of this quasi-isometry depend only on the parameters of ψ_0 , Ξ_Σ , Ξ'_Σ (including $\text{genus}(\Sigma)$) and on inj_0 .*

Note that such a quasi-isometry is necessarily a stratified quasi-isometry with uniform parameters (by Proposition 16.3, since a lower bound on injectivity radius on hyperbolic surfaces places an upper bound on their diameters). In fact, our construction will give us a stratified quasi-isometry directly.

Conversely, it is easily seen that any equivariant (stratified) quasi-isometry must respect ending laminations. (See for example, Lemma 16.2 and the subsequent comments.) Whereas the assumption on injectivity radii plays an essential role in the proof, it is unclear to what extent it is necessary for the result to hold.

Before beginning the proof of Theorem 17.1, we give a brief outline of the overall strategy. We shall construct the quasi-isometry sheetwise, using Lemma 16.2 and subsequent remarks. This falls into two pieces. Firstly we need a quasi-isometry, $f : A \rightarrow \mathbb{N}$, where $A \subseteq \mathbb{N}$ is a quasidense subset. Secondly we will need uniform sheetwise quasi-isometries, $\psi_i : X_i \rightarrow X'_{f(i)}$ for all $i \in A$. Since the intrinsic geometry on the sheets, X_i and X'_j , is evolving exponentially in the indices i and j , we will need to determine f up to an additive bound in order to have any hope of constructing the quasi-isometries ψ_i . Determining how the levels of Ξ and Ξ' correspond with this precision constitutes the major part of the proof. Once this is achieved the maps ψ_i are constructed using Lemma 15.7. In other words, we only need to control projected lengths with respect to the ending laminations, Λ_Σ and Λ'_Σ , of the stacks. We are thus effectively reduced to analysis of the lengths of horizontal segments, and how these change as the level changes. We need to show that it does so in some kind of uniform way. For this, we can work in just one of the stacks, say Ξ .

For reasons of exposition, it will be convenient to shift perspective and imagine the spaces $X_i \equiv X$ as being fixed, while the intrinsic metric, ρ_i , changes. (This is justified by Lemma 13.8.) In particular, we are interested in the restriction of ρ_i to leaves of the lamination. In retrospect (given Theorem 17.1, and the ‘‘model case’’ described in Section 18), we know that, up to reparameterising within linear bounds, these

distances must decay uniformly exponential. To see how to reparameterise, we shall use the regularity of the lamination as discussed in Section 13.

Let Ξ_Σ be a semi-infinite hyperbolic surface stack with ending lamination Λ_Σ . Suppose $\text{inj}(\Xi_\Sigma) \geq \text{inj}_0 > 0$. By adding a finite number of isolated leaves, we can assume, for convenience that Λ_Σ is maximal (In fact, since Λ_Σ is regular, it necessarily fills Σ .) For each i , have a natural equivariant quasi-isometry $\phi_i : X_i \rightarrow X_{i+1}$. Indeed we can assume that ϕ_i satisfies the conclusion of Lemma 13.8. In particular, ϕ_i is equivariant, and sends the realisation of Λ in X_i to its realisation in X_{i+1} . For expositional convenience, we shall identify all the spaces X_i via compositions of the maps ϕ_i , to give as a fixed set X , with a sequence of metrics ρ_i . Note that (X, ρ_i) is always isometric to \mathbb{H}^2 , and the realisation of Λ is fixed as a collection of subsets of X . We can also assume that the centres of all complementary triangles are fixed. The metrics, ρ_i , descend to a sequence of hyperbolic metrics on Σ which we also denote by ρ_i . We are thus in the situation described in Section 14. We need to verify properties (L1)–(L4).

Property (L1) is immediate from the construction. Note that chain sequences as defined in Section 14 correspond exactly to chains in the stack Ξ_Σ . In particular, given $x, y \in X$, the function $[i \mapsto \rho_i(x, y)]$ is uniformly quasiconvex. If $x, y \in \lambda \in \Lambda$, then this function is bounded. We have thus verified property (L2). Property (L3) is given by the hypotheses. We thus need to show that (L4) follows from (or is indeed equivalent to) (L3).

Given a closed curve, γ in Σ , write γ_i for the homotopic closed geodesic in (Σ, ρ_i) . By definition, $l_i(\gamma) = \text{length}_i(\gamma)$. If $g \in \Gamma$ is any element in the conjugacy class representing γ , then $l_i(\gamma)$ is precisely the stable length, $\text{sl}_i(g)$, of g . By Lemma 6.6, we obtain:

Lemma 17.2. *The map $[i \mapsto l_i(\gamma)]$ is uniformly quasiconvex, with additive constant 0.*

The following is now an easy consequence of the discussion in Section 5, in particular, Lemma 5.5.

Proposition 17.3. *Let Ξ_Σ be a hyperbolic surface stack. The following are equivalent:*

- (1) $\text{inj}(\Xi_\Sigma) > 0$.
- (2) $(\forall l)(\exists p)$ if γ is an essential closed curve on Σ and $i, j \in \mathbb{N}$ with $l_i(\gamma) \leq l$ and $l_j(\gamma) \leq l$, then $|i - j| \leq p$.
- (3) $(\exists l, p)$ if γ is an essential closed curve on Σ and $i, j \in \mathbb{N}$ with $l_i(\gamma) \leq l$ and $l_j(\gamma) \leq l$, then $|i - j| \leq p$.

It follows by Proposition 14.6 that:

Proposition 17.4. *The ending lamination of a hyperbolic surface stack with lower injectivity radius bound is regular.*

Indeed with respect to any of the metrics, (Σ, ρ_i) , there is a linear regulating function depending only on the parameters of the stack.

The version we shall require in this section is that there is a function $f_0 : \mathbb{N} \rightarrow [0, \infty)$ depending on the parameters such that if $E \in \mathcal{H}_\Sigma$ and γ is any surrounding curve, then $\text{length}_i(E) \leq f(\text{plength}_i(\gamma))$. (For the moment, we not require that the function f_0 be linear, and so Proposition 14.2 suffices for this.)

It will be convenient to assume that (L2) is strengthened to assume that the distances are linearly controlled, i.e. that there exist $t < 1$ and $r \geq 0$ such that $\rho_{i+1}(x, y) \leq t\rho_i(x, y) + r$ for all $i \in \mathbb{N}$ and $x, y \in \lambda \in \Lambda$. This is justified by the remarks in Section 14.

We shall need the following combinatorial means of comparing the lengths of horizontal intervals. First, we choose a constant, r_1 , sufficiently large so that if α is a path in Σ , and $i \in \mathbb{N}$, then $\text{pdiam}_{i+1}(\alpha) \leq \max\{\text{pdiam}_i(\alpha), r_1\}$. Moreover, we can assume that any point of (Σ, ρ_i) is contained in a (simple) closed curve of length at most r_1 . (This depends only on $\text{genus}(\Sigma)$ and $\text{inj}(\Sigma) \geq \text{inj}_0$.) Given $E \in \mathcal{H}_\Sigma$, we write $h(E) = \min\{i \mid \text{length}_i(E) \leq r_1\}$. We can suppose that r_1 is large enough so that $h(E) > 0$, then $\text{length}_{h(E)}(E) \geq r_2$ for some uniform constant, $r_2 > 0$.

We define a *shortcut* of a horizontal interval $E \in \mathcal{H}_\Sigma$ to be an arc σ such that $\alpha \cap E = \partial\alpha$ and $\alpha \cup E_\alpha$ is essential in Σ , where E_α is the subarc of α lying between the endpoints $\partial\alpha$, and such that $\text{pdiam}(\alpha) \leq r_1$. By the choice of r_1 , every horizontal interval admits a shortcut, and $\text{pdiam}_i(\alpha) \leq r_1$ for all $i \in \mathbb{N}$.

Given $E, F \in \mathcal{H}_\Sigma$, let $n(E, F)$ be the minimal value of $|\alpha \cap F|$, where α is a shortcut of E .

Lemma 17.5. *There is a function, $f_1 : \mathbb{N} \rightarrow \mathbb{N}$ depending on the parameters of Ξ_Σ such that if $E, F \in \mathcal{H}_\Sigma$, with $n(E, F) \leq p$, then $h(F) \leq h(E) + f_1(p)$.*

Proof. Let $i = h(E)$. Suppose first that $n(E, F) = 0$. There is a shortcut, α , such that $\gamma = \alpha \cup E_\alpha$ is a surrounding curve for F . Now, $\text{pdiam}_i(E_\alpha) = \text{length}_i(\alpha) \leq \text{length}_i(E) \leq r_1$. Thus, by Lemma 13.7, $\text{plength}_i(\gamma) \leq 2(\text{pdiam}_i(E_\alpha) + \text{pdiam}_i(\alpha)) \leq 4r_1$, and so $\text{length}_i(F) \leq f_0(4r_1)$.

In the general case, we can cut F into (at most) $p + 1$ subintervals F_0, \dots, F_p with $n(E, F_j) = 0$. Thus $\text{length}_i(F) \leq (p + 1)f_0(4r_1)$. This places an upper bound on $h(F) - i$ in terms of p . \square

Conversely, we have:

Lemma 17.6. *There is a function $f_2 : \mathbb{N} \rightarrow \mathbb{N}$ depending on the parameters of Ξ_Σ such that if $E, F \in \mathcal{H}_\Sigma$ with $h(F) \leq h(E) + p$, then $n(E, F) \leq f_2(p)$.*

Proof. Let α be a shortcut of E with $|\alpha \cap F|$ minimal. Let F_1, \dots, F_q be the subintervals into which $\alpha \cap F$ cuts F . For each j , let α_j be the subarc of α between the endpoints of F_j , and let $\gamma_j = \alpha_j \cup F_j$. Now γ_j is essential in Σ , for otherwise, we could reduce $|\alpha \cap F|$ by replacing α by an arc following close to F_j . (This would increase $\text{pdiam}_0(\alpha)$ by at an arbitrarily small amount, so it would remain a shortcut of E .) Thus, γ_j is a surrounding curve for E .

Now choose j so that $h(F_j)$ is minimal, and let $i = h(F_j)$. Now $\text{pdiam}(F_j) = \text{length}(F_j) \leq r_1$ and $\text{pdiam}(\alpha_j) \leq \text{pdiam}(\alpha) \leq r_1$. Thus by Lemma 13.7, $\text{plength}(\gamma_j) \leq 4r_1$, and so $\text{length}_i(E) \leq f_0(4r_1)$. This places an upper bound, say t , on $h(E)$ in terms of the parameters of Ξ_Σ .

Case (1): $i = 0$. Now $h(F) \leq t + p$, and so $\text{length}_0(F)$ is bounded above in terms of p . Let β be the shortest shortcut of E . Thus, β is a geodesic segment of length at most r_1 . But there is a bound, in terms of $\text{inj}(\Sigma, \rho_0) \geq \text{inj}_0$, on the number of points in which two geodesic segments of bounded length, in particular, β and F , can intersect. This puts a bound on $n(E, F) \leq |\beta \cap F|$ as required.

Case (2): $i > 0$. In this case, there is positive lower bound, r_2 , on $\text{length}_i(F_k)$ for all k . Thus $\text{length}_i(F) \geq (q-1)r_2$. Moreover, $h(F) - i \leq h(E) + p - i \leq t + p$. This places an upper bound on $\text{length}_i(F)$ and hence on q in terms of p . But $n(E, F) \leq q$, and the result follows. \square

Lemma 17.7. *For each $i \in \mathbb{N}$, there is an interval $E_i \in \mathcal{H}_\Sigma$ with $h(E_i) = i$. Moreover we can assume that $E_i \subseteq E_{i+1}$ for all i .*

Proof. Choose any $x_0 \in \lambda \in \Lambda$. Given any $t \in [0, \infty)$ let E_t be the horizontal interval of ρ_0 -length t lying on one side of x_0 , and let $f(t, i) = \text{length}_i(E_t)$. For each i , the map $[t \mapsto f(t, i + 1)]$ is continuous and strictly increasing in t . Moreover, for all i , $f(t, i + 1) < \max\{f(t, i), r_1\}$. It follows that $h(E_t) = \min\{i \mid f(t, i) \leq r_1\}$ is non-decreasing in t , and takes all possible values in \mathbb{N} . \square

We are now ready to embark on the proof of Theorem 17.1. Let Ξ_Σ and Ξ'_Σ be as in the hypotheses. In each stack we identify the sheets as

sets, and write them as (X, ρ_i) and (X', ρ'_i) . We write Λ and Λ' for the ending laminations realised in X and X' . Let $\Sigma = X/\Gamma$ and $\Sigma' = X'/\Gamma'$, and let $\psi_0 : (X, \rho_i) \rightarrow (X', \rho'_i)$ be the equivariant quasi-isometry as given by the hypotheses.

We first observe, that (up to stratified quasi-isometry), there is no loss in assuming that $\rho_0 = \rho'_0$, and that ψ_0 is the identity. To see this, note that we can replace Ξ' by the stack $\bigvee_{i \in \mathbb{N}} Y_i$, where $Y_0 = X_0$ and $Y_i = X'_{i-1}$ for $i \geq 1$, and where the connecting quasi-isometries between Y_i and Y_{i+1} are taken to be ψ_0 if $i = 0$ and to be the natural quasi-isometry between X'_{i-1} and X_i in Ξ' for $i > 0$. (In other words, we have essentially added another sheet, X_0 , to Ξ' via the quasi-isometry ψ_0 .) Note that the new parameters of Ξ' depend on the old, and those of the quasi-isometry ψ_0 .

From this point on, we can identify (X, ρ_0) with (X', ρ'_0) , and hence Λ with Λ' . We want to construct a sequence of uniform quasi-isometries between the sheets (X, ρ_i) and $(X, \rho'_{i'})$. First we need to work out the correspondence between the i 's and i' 's.

Let $h' : \mathcal{H}_\Sigma \rightarrow \mathbb{N}$ be the function corresponding to h for the stack Ξ' . We can apply Lemmas 17.5 and 17.6 equally well to h' . We thus have $h(F) \leq h(E) + p \Rightarrow n(E, F) \leq f_2(p) \Rightarrow h'(F) \leq h'(E) + f_1 \circ f_2(1)$, and similarly interchanging h and h' . Setting $q = f_1 \circ f_2(1)$, we deduce that:

$$\begin{aligned} h(F) \leq h(E) + 1 &\Rightarrow h'(F) \leq h'(E) + q \\ h'(F) \leq h'(E) + 1 &\Rightarrow h(F) \leq h(E) + q. \end{aligned}$$

Now let E_i be the sequence of intervals given by Lemma 17.7 for the stack Ξ . Given $i \in \mathbb{N}$, let $f(i) = h'(E_i)$. This defines a non-increasing function $f : \mathbb{N} \rightarrow \mathbb{N}$ with $f(0) = 0$ and $f(i+1) \leq f(i) + q$. Moreover, if $f(i) = f(j)$, then $|i - j| \leq q$. It follows that f is a quasi-isometry from \mathbb{N} to \mathbb{N} .

By Lemma 16.2, it remains to construct a uniform quasi-isometry from (X, ρ_i) to $(X, \rho'_{f(i)})$.

Lemma 17.8. *There is a function $g : [0, \infty) \rightarrow [0, \infty)$ with the following property. Suppose that $i \in A$ and let $j = f(i)$. If $E \in \mathcal{H}_\Sigma$, then $\rho'_j(E) \leq g(\rho_i(E))$ and $\rho_i(E) \leq g(\rho'_j(E))$.*

Proof. By symmetry, it is sufficient to verify the first inequality. Let r_1 be the constant featuring in the definitions of h and h' (i.e. $h(E) = \min\{i \mid \rho_i(E) \leq r_1\}$ and $h'(E) = \min\{j \mid \rho'_j(E) \leq r_1\}$). We can subdivide any interval, E , into subintervals of ρ_i -length at most r_1 , and so it suffices to bound the ρ'_j -length of E for any $E \in \mathcal{I}_\Sigma$ with $\rho_i(E) \leq r_1$. To do this, note that $h(E) \leq i = h(E_i)$. Thus $h'(E) \leq$

$h'(E_i) + q$. By definition of f , we have $h'(E_i) = f(i) = j$, and so $h'(E) - j \leq q$. But this now places a bound on $\rho'_j(E)$ as required. \square

The following is now an immediate consequence of the definition of projected distances, σ and σ' .

Lemma 17.9. *There is a function $g : [0, \infty) \rightarrow [0, \infty)$ with the following property. Suppose $i \in A$ and $j = f(i)$. Then for all $x, y \in X$, we have $\sigma'_j(x, y) \leq g(\sigma_i(x, y))$ and $\sigma_i(x, y) \leq g(\sigma'_j(x, y))$.*

We have thus verified the hypotheses of Lemma 15.7. This shows that for each $i \in A$, there is a uniform quasi-isometry from (X, ρ_i) to $(X, \rho'_{f(i)})$. Assembling these quasi-isometries, as described by Lemma 13.2, we obtain a stratified quasi-isometry from Ξ to Ξ' .

This proves Theorem 17.1.

18. SINGULAR EUCLIDEAN METRICS AND SOL GEOMETRY

In this section we shall discuss singular euclidean metrics on surfaces, and describe how these give rise to stacks with a singular Sol geometry. Spaces of this type are constructed in [CannT], and play an important role in the work of Minsky [Min2]. In Section 20, we shall explain how such stacks arise from Teichmüller geodesics. For some discussion of singular euclidean surfaces, see [EP, Bow1], and the references therein. We are primarily interested in non-positively curved metrics. These are locally CAT(0) in the sense of Aleksandrov as defined by Gromov [Gr1]. For further references on CAT(0), and more generally, CAT(κ) spaces, see for example Ballmann's article in [GhH] or [BriH].

Definition. A *singular euclidean surface* is a complete locally compact path-metric space, (X, d) , together with a discrete subset, $\mathcal{S} \subseteq X$, such that $X \setminus \mathcal{S}$ is locally modelled on the euclidean plane.

Thus, X is topologically a surface. Each point, x , of \mathcal{S} is a *cone point*, i.e. has a neighbourhood isometric to the completion of the quotient of the universal cover of $\mathbb{R}^2 \setminus \{0\}$ by a rotation through an angle of $\theta(x)$. We refer to $\theta(x)$ as the *cone angle* at x . (We allow for X to be non-singular at x , i.e. $\theta(x) = 2\pi$.)

Remark. For most of what we do in this section, one can drop the assumption of local compactness, thereby allowing for “cone angles” of ∞ , as for example, in the completion of the universal cover of $\mathbb{R} \setminus \{0\}$.

Definition. We say that X is *non-positively* curved if $\theta(x) \geq 2\pi$ for all $x \in \mathcal{S}$.

This is equivalent to asserting that X is locally CAT(0). If X is also simply connected, then it will be globally CAT(0).

If we sharpen this condition slightly, we obtain:

Lemma 18.1. *Given $r, \epsilon > 0$, there is some $k \geq 0$ such that if X is a simply connected singular euclidean surface such that \mathcal{S} is r -dense in S and such that $\theta(x) \geq 2\pi + \epsilon$ for all $x \in \mathcal{S}$, then X is k -hyperbolic.*

We note that in this case, the Gromov boundary, ∂X , is a topological circle. There are a number of ways to see this. For example, note that any pair of points of ∂X are connected by a bi-infinite geodesic which cuts X into two convex sets. It follows that the pair separates ∂X . We see that ∂X is a metrisable continuum in which every pair of points separates, and thus homeomorphic to a circle.

The idea of the proof of Lemma 18.1 is to “diffuse” the negative curvature away from the cone points so as to obtain a space that is globally CAT($-\eta^2$) for some $\eta > 0$ depending only on r and ϵ . Such a space is k -hyperbolic, where k depends only on η . One way to achieve this is to combine the following two lemmas:

Lemma 18.2. *Suppose that X is a singular euclidean surface such that \mathcal{S} is r -dense in X . Then we can triangulate X so that each edge is geodesic of length at most $2r$.*

Proof. (The technicalities of the proof are simpler if we assume that X is CAT(0), as will be the case in applications. In this case we get a genuine triangulation with no loops or multiple edges etc.) One can start with the Delaunay cellulation of S , (see [EP, Bow1]). Thus, we join $x, y \in \mathcal{S}$ by an edge if there is a closed euclidean disc $D \subseteq X$ such that $D \cap \mathcal{S} = \{x, y\}$. This cuts X into cyclic polygons of circumradius at most r . We then subdivide each polygon into triangles. \square

Given a euclidean triangle, T , and $\eta > 0$ we write $T(\eta)$ for the comparison triangle in the space of constant curvature $-\eta^2$ (i.e. having the same side lengths). The following is a simple geometric exercise:

Lemma 18.3. *Given $R \geq 0$, there is a continuous function, $\lambda_R : [0, \infty) \rightarrow \mathbb{R}$, with $\lambda_R(0) = 1$ with the following property. Suppose that $\eta > 0$, and that T is a euclidean triangle with side lengths at most R . Suppose that θ and $\theta(\eta)$ are corresponding angles in T and $T(\eta)$. Then $\theta(\eta) \geq \lambda_R(\eta)\theta$.*

We can now prove Lemma 18.1:

Proof. We take the triangulation given by Lemma 18.2, and replace each triangle by the comparison triangle in the space of constant curvature $-\eta^2$. This gives a surface locally modelled on this space with

cone angles at least $\lambda_{2r}(2\pi + \epsilon)$. By choosing η sufficiently small in relation to ϵ , we can arrange that this is always at least 2π . Such a space is locally $\text{CAT}(-\eta^2)$. Since X is simply connected, it is also globally $\text{CAT}(-\eta^2)$, and hence k -hyperbolic, where k depends on η and hence on r and ϵ . \square

Definition. By a *grid structure* we mean a parallel line field on $X \setminus \mathcal{S}$.

Such a structure is locally modelled on a field of parallel lines in the euclidean plane. We refer to such lines as *horizontal*. A grid structure gives rise to a *horizontal foliation*. Such a leaf is generically locally geodesic, with tangent vectors horizontal, though we need to allow leaves to have finite valence branch points at the cone points.

We can rotate a grid structure through any angle. In particular, rotating by $\pi/2$, we obtain an orthogonal grid structure. We refer to the horizontal foliation of the latter as the *vertical foliations* of the former.

A singular euclidean surface admits an grid structure if and only if all the cone angles are integral multiples of π , in which case we get a whole circle of grid structures.

Definition. By a *grid surface*, we mean a singular euclidean surface together with a grid structure.

Note that if X is a non-positively curved grid surface, then the set of cone angles lies in $\{2\pi, 3\pi, 4\pi, \dots\}$. (A non-positively curved grid surface is essentially the same as a Riemann surface together with a quadratic differential [Ga].) If X is also simply connected, then every non-singular horizontal leaf is a bi-infinite geodesic which cuts X into two convex half-spaces. A singular leaf is a geodesically embedded tree whose complement is a (possibly infinite) union of convex half-spaces.

In a neighbourhood of a non-singular point, we can choose horizontal and vertical coordinates, (ξ, ζ) , so that ξ -axes and ζ -axes are respectively horizontal and vertical. The infinitesimal metric, ds , is given by $ds^2 = d\xi^2 + d\zeta^2$. It will sometimes also be convenient to consider the L^1 -metric, \bar{d} , given infinitesimally by the Finsler metric $d\xi + d\zeta$. The identity map from (X, d) to (X, \bar{d}) is thus $\sqrt{2}$ -bilipschitz.

Given a real number, t , we can deform the metric d by contracting by a factor e^{-t} in the horizontal direction, and expanding by a factor e^t in the vertical direction. In other words, the infinitesimal metric ds_t is given by $ds_t^2 = e^{-2t}d\xi^2 + e^{2t}d\zeta^2$. We write d_t for the induced path metric. This is also a grid surfaces, whose horizontal and vertical foliations are identical with the original. Note that cone angles, and thus non-positive curvature, are preserved by such a deformation.

A path $\alpha \subseteq X$ is a local geodesic with respect to the metric d if and only if each component of $\alpha \setminus \mathcal{S}$ is a euclidean geodesic and the exterior angles at each cone point are at least π . This property is clearly invariant under the deformation under the deformation described above. If X is non-positively curved, then any local geodesic is a global geodesic. We conclude:

Lemma 18.4. *Let (X, d) be a non-positively curved grid surface. If α is geodesic with respect to the metric d , then it is geodesic with respect to the metric d_t for all $t \in \mathbb{R}$.*

Lemma 18.5. *Let (X, d) be a simply connected non-positively curved grid surface. If $x, y \in X$, then the map $[t \mapsto d_t(x, y)] : \mathbb{R} \rightarrow [0, \infty)$ is uniformly quasiconvex.*

Proof. Let α be the d -geodesic connecting x to y . By Lemma 18.4, α is a d_t -geodesic for all t , and so $d_t(x, y)$ is the d_t -length of α . Now the L^1 -length, \bar{d}_t -length(α) of α has the form $Ae^t + Be^{-t}$ for constants $A, B \geq 0$. Thus the map $[t \mapsto d_t$ -length(α)] is a uniformly quasiconvex function, being a positive linear combination of the quasiconvex functions $[t \mapsto e^t]$ and $[t \mapsto e^{-t}]$. Now the d_t -length of α is $\sqrt{2}$ -bilipschitz related to its d_t -length, and so it follows that $[t \mapsto d_t$ -length(α)] is uniformly quasiconvex as required. \square

Suppose that $\beta : \mathbb{R} \rightarrow X$ is a path. We say that β has *speed bounded* by $\mu \geq 0$ if for all $t \in \mathbb{R}$, $\limsup_{u \rightarrow 0} d_t(\beta(t), \beta(t+u)) \leq \mu$. (Thus if β is smooth, then the tangent vectors $\partial\beta/\partial t$ have length at most μ with respect to the metric d_t .) We note the following generalisation of Lemma 18.5:

Lemma 18.6. *Suppose that $\alpha, \beta : \mathbb{R} \rightarrow X$ are paths with speed bounded by $\mu \geq 0$. Then the map $[t \mapsto d_t(\alpha(t), \beta(t))]$ is quasiconvex with parameters depending only on μ .*

Proof. First note that $d_t(\alpha(t), \alpha(u)) \leq h(t-u)$ where $h(x) = \mu(e^{|x|} - 1)$. We see that if $t, u \in \mathbb{R}$, then $d_u(\alpha(u), \beta(t)) \leq d_u(\alpha(u), \alpha(t)) + d_u(\alpha(t), \beta(t)) + d_u(\beta(t), \beta(u)) \leq d_u(\alpha(t), \beta(t)) + 2h(t-u)$. Now the map $[u \mapsto d_u(\alpha(t), \beta(t))]$ is uniformly quasiconvex by Lemma 18.5, so the result follows by Lemma 5.3. \square

Definition. We say that grid surface is *normal* if it satisfies the hypotheses of Lemma 18.1.

In other words, X is simply connected, \mathcal{S} is quasidense in X , and the set of cone angles lies in $\{3\pi, 4\pi, 5\pi, \dots\}$.

Now, X is hyperbolic, and ∂X is a circle. Every horizontal leaf of X is convex, and so determines a closed subset of ∂X . This is a pair in the generic case of a non-singular leaf.

Note that for all $t \in \mathbb{R}$, the identity map from (X, d) to (X, d_t) is bilipschitz, and hence a quasi-isometry. Note that (X, d_t) is also normal (perhaps with larger quasidensity parameter), and we get a natural identification of ∂X with the Gromov boundary of (X, d_t) . Under this identification, each horizontal leaf gives rise to the same closed subset of ∂X .

Definition. If $I \subseteq \mathbb{R}$, we say that X is I -normal if the grid surfaces, (X, d_t) are uniformly normal for $t \in I$.

Suppose that X is $[0, \infty)$ -normal. In this case, any two parallel horizontal rays must eventually coincide. For otherwise they would bound a euclidean strip foliated by horizontal rays. For sufficiently large t , such a strip would contain arbitrarily large euclidean d_t discs, contradicting uniform normality.

As a consequence, we see that the closed subsets of ∂X determined by distinct leaves are disjoint. We thus obtain an equivalence relation \sim on ∂X , by writing $x \sim y$ if $x = y$ or if x and y are the endpoints of a horizontal bi-infinite geodesic. This relation is closed and unlinked.

We shall need the following observation:

Lemma 18.7. *Suppose that $A, B \subseteq X$ are disjoint closed horizontal half-spaces bounded by bi-infinite geodesics lying in different horizontal leaves. Then $d_t(A, B) \rightarrow \infty$ as $t \rightarrow \infty$.*

Proof. From the above observations, the boundaries of A and B must diverge in both directions. There is thus a shortest path in (X, d) connecting A to B . Such a geodesic meets the boundaries of A and B in angles at least $\pi/2$. Since this property is preserved under the stretching deformation, we see that it remains a shortest path between A and B in (X, d_t) for all $t \geq 0$. Now this path cannot be entirely horizontal, and so its length must tend to ∞ . \square

Suppose that X is a grid surface, and $I \subseteq \mathbb{R}$ is a closed interval. We can put a singular riemannian metric on $\Xi = X \times I$ as follows. Let (ξ, ζ) be horizontal and vertical coordinates in a neighbourhood of a point of $X \setminus \mathcal{S}$. Let t be the I -coordinate. We define an infinitesimal metric, ds , by $ds^2 = e^{-2t}d\xi^2 + e^{2t}d\zeta^2 + dt^2$. Let ρ be the induced path metric on $(X \setminus \mathcal{S}) \times I$. This extends to a complete path metric on $X \times I$, which we also denote by ρ . This is a ‘singular Sol’ metric in the sense that away from $\mathcal{S} \times I$ it is locally modelled on 3-dimensional Sol geometry

(see for example [Sc]). Note that paths of the form $\{x\} \times I$ for $x \in X$ are all global geodesics.

Given $t \in I$, let $X_t = X \times \{t\} \subseteq \Xi$, and write ρ_t for the induced path-metric on X_t . Thus, the map $[x \mapsto (x, t)]$ is an isometry from (X, d_t) to (X_t, ρ_t) . We can thus think of Ξ as a “continuous stack”. To relate this to our discrete formulation, we can take the sheets to be $(X_i)_{i \in I \cap \mathbb{Z}}$. (We can assume that $I \cap \mathbb{Z} \neq \emptyset$.)

Suppose that now that X is non-positively curved, and that $\alpha \subseteq X$ is a horizontal segment. Then $\alpha \times I$ is a totally geodesic surface locally modelled on the hyperbolic plane. The arcs $\alpha \times \{t\}$ are horocyclic. In particular, note that if α is bi-infinite and $I = [0, \infty)$, then we get a horodisc foliated by horocircles.

Proposition 18.8. *Suppose $I \subseteq \mathbb{R}$ is an interval and that X is an I -normal grid surface. Let $\Xi = X \times I$ be the singular Sol space constructed above. Then Ξ is hyperbolic where the hyperbolicity constant depends only on the quasidensity parameter of the normality hypothesis.*

Proof. Indeed, Ξ is a hyperbolic stack with sheets $(X_t)_{t \in I \cap \mathbb{Z}}$. This follows easily by applying Lemma 18.1, Lemma 18.6 and Theorem 6.4. \square

Now suppose that $I = [0, \infty)$. It follows by Proposition 10.2 that $\partial\Xi$ is a dendrite. As usual we may identify $\partial^0\Xi$ with ∂X . We have a Cannon-Thurston map $\omega : \partial X \rightarrow \partial\Xi$. This is surjective. This can be seen by applying essentially the same argument as Proposition 11.1, Alternatively, it can be seen more directly as follows. Suppose $(x_i, t_i)_{i \in \mathbb{N}}$ converges on $y \in \partial\Xi$, where $x_i \in X$ and $t_i \in [0, \infty)$. If (x_i) is unbounded, then it subconverges on a point $a \in \partial X$, and we see that $\omega(a) = y$. If not, then $t_i \rightarrow \infty$, and x_i subconverges on some $x \in X$. Now x lies in some horizontal bi-infinite geodesic α . In this case, $\alpha \times [0, \infty)$ is a geodesically embedded horodisc, and (x_i, t_i) converges on its limit point. This limit point is equal to $\omega(a)$ where $a \in \partial X$ is either of the ideal points of α .

We have already defined a closed unlinked equivalence relation, \sim , on ∂X , where equivalent points are the endpoints of a horizontal geodesic. We note:

Lemma 18.9. *If $a, b \in X \cup \partial X$, then $\omega(a) = \omega(b)$ if and only if $a \sim b$.*

Proof. Suppose first that $a \sim b$. Let α be the horizontal geodesic connecting a to b . Now $\alpha \times [0, \infty)$ is a totally geodesic horodisc, and so the both rays of α converge on the same point of $\partial\Xi$. Thus (using Lemma 7.4) we have $\omega(a) = \omega(b)$.

Conversely, suppose that $a \not\sim b$. Let $\gamma \subseteq X$ be a bi-infinite geodesic connecting a to b . Since γ is not horizontal, it must cross two distinct horizontal bi-infinite geodesics, α and β . Let A and B be closed half-spaces bounded by α and β , so that $A \cap B = \emptyset$, $a \in \partial A \subseteq \partial X$ and $b \in \partial B \subseteq \partial X$.

Let $P = A \times [0, \infty) \subseteq \Xi$ and $Q = B \times [0, \infty) \subseteq \Xi$. Thus P and Q are convex subsets. By Lemma 18.7, $d_t(A, B) \rightarrow \infty$ as $t \rightarrow \infty$, from which it follows that $N(P, r) \cap N(Q, r)$ is compact for all $r \geq 0$. Thus $\partial P \cap \partial Q = \emptyset$. By Lemma 7.4, $\omega(a) \in \partial P$ and $\omega(b) \in \partial Q$, and so $\omega(a) \neq \omega(b)$, as required. \square

In other words, we conclude that the lamination on X determined by the equivalence relation \sim on ∂X is the ending lamination of the stack Ξ .

We now consider a closed non-positively curved surface (Σ, d) . Let $A = \text{area}(\Sigma, d)$. Let $I \subseteq \mathbb{R}$ be an interval.

Lemma 18.10. *Suppose that $\alpha \subseteq \Sigma$ is an arc of length ϵ . Then there is a horizontal interval $E \subseteq \Sigma$ of length at least A/ϵ and with $\alpha \cap \text{int}(E) = \emptyset$.*

(For the moment, we need to allow for the possibility that E wraps around a closed horizontal geodesic disjoint from α .)

Proof. Given $r \geq 0$, let $S(\alpha, r)$ be the set of points of Σ that are contained in an interval of length less than r that intersects α . We see that $\text{area}(S(\alpha, r)) \leq 2\epsilon r$. Setting $r = A/2\epsilon$, we see that there is some point $y \in \Sigma \setminus S(\alpha, r)$. Let E be the horizontal interval of length $2r$ centred on y . \square

As for hyperbolic surfaces, the *injectivity radius*, $\text{inj}(\Sigma, d)$ can be defined as half of the length of the shortest closed geodesic. We note that $\text{diam}(\Sigma, d)$ is bounded above in terms of $\text{area}(\Sigma, d)$ and $\text{inj}(\Sigma, d)$ (see Lemma 19.1). If Σ has negative Euler characteristic, then it follows that the universal cover, $X = \tilde{\Sigma}$ is uniformly normal.

Note that $\text{area}(\Sigma, d_t)$ is constant. As above, we can put a singular Sol geometry on $\Sigma \times [0, \infty)$. Let Ξ be its universal cover. Applying Lemma 18.1 and Proposition 18.8, we see:

Lemma 18.11. *Suppose that Σ is a closed grid surface with $\text{genus}(\Sigma) \geq 2$. Suppose that for all $t \in I$, we have $\text{inj}(\Sigma, d_t) \geq \eta > 0$. Then the singular Sol space $\Xi = \tilde{\Sigma} \times I$ is hyperbolic, with constant depending on $\text{area}(\Sigma)$, $\text{genus}(\Sigma)$ and η .*

Given a grid surface, we can define the notion of a “surrounding curve” of a horizontal interval, and hence a “regulating function” exactly as for laminations in Section 13. We thus have a notion of “regularity” for a grid surface. We now suppose again that $I = [0, \infty)$. We note:

Lemma 18.12. *If (Σ, d) is a compact non-positively curved regular grid surface, then there is a positive lower bound on the injectivity radii of the surfaces (Σ, d_t) for $t \in [0, \infty)$ which depends only on the regularity function and $\text{area}(A)$.*

Proof. We show that for any $\nu \geq 0$, there is some $\eta > 0$ such that if $\text{inj}(\Sigma, d_t) \leq \eta$ for some t , then we can find a simple closed curve (in fact, a closed geodesic) and a disjoint horizontal interval, E , with $d(E) > \nu d(\gamma)$. Moreover, if η is sufficiently small, then E can be chosen arbitrarily long.

To this end, choose $\eta < \sqrt{A/2\nu}$. Suppose that $\text{inj}(\Sigma, d_t) \leq \eta$. Let γ be a simple closed geodesic of d_t -length at most 2η . By Lemma 18.11 (applied to (Σ, d_t)), there is a horizontal interval disjoint from γ of d_t -length at least $A/2\eta > \nu\eta$. Now $\text{length}(\gamma) \leq e^t \text{length}_t(\gamma) \leq \eta e^t$ and $\text{length}(E) = e^t \text{length}_t(E) > \nu\eta e^t$ and the result follows. \square

In fact, we can elaborate on Proposition 18.12. We claim that we can choose γ to consist of an alternating sequence of horizontal and vertical segments, where the length of each vertical segment is bounded in terms of A , and the total number of such segments is bounded in terms of $\text{genus}(\Sigma)$.

To this end, we define a *rectilinear* path to be one consisting only of horizontal and vertical segments. Fix a structure, (Σ, d_t) , and in the proof of Proposition 18.12. Clearly we can find a rectilinear closed curve of length at most $\sqrt{2}$ times the injectivity radius. Among such curves, choose one of minimal length. Note that such a curve must be simple. Now among such shortest curves, choose one with a minimal number of segments. We can assume that each vertical segment of such a curve contains at least one singular point, since by pushing it as far as we can in a horizontal direction, we either run into a singular point, or reduce the total number of segments, contrary to our assumption. Since the curve is simple, the number of vertical segments is bounded by the number of singular points, and hence in terms of $\text{genus}(\Sigma)$.

We now use this curve in the proof of Proposition 18.2. We see that the d_t -lengths and hence the d -lengths of the vertical segments are bounded in terms of A as claimed.

19. NON-POSITIVELY CURVED METRICS ON SURFACES

We want to relate singular euclidean metrics of the type described in Section 18 to hyperbolic metrics on a surface, Σ . We shall do this by placing them in the more general context of non-positively curved metrics. We aim to show that two such surfaces, of given area, are a bounded distance apart in Teichmüller space if and only if their universal covers are uniformly equivariantly quasi-isometric, given a lower bound on injectivity radius. Indeed, we show that it is enough that one such surface has a lower bound on injectivity radius.

Let Σ be a closed surface with $\text{genus}(\Sigma) \geq 2$, and let $\Gamma = \pi_1(\Sigma)$. Let $\mathcal{T} = \mathcal{T}(\Sigma)$ be the associated Teichmüller space. We can thus think of \mathcal{T} as the space of marked conformal structures, or equivalently marked hyperbolic structures, on Σ . Let $d_{\mathcal{T}}$ be the Teichmüller metric. The mapping class group of Σ acts properly discontinuously on \mathcal{T} , with quotient the “moduli space” of Σ .

By a “non-positively curved metric” on Σ we shall mean a non-positively curved riemannian metric, possibly with a finite number of cone singularities each with cone-angle at least 2π . The induced conformal structure gives us a point in \mathcal{T} .

We note:

Lemma 19.1. *Given $\eta > 0$ and $A \geq 0$, there is some $L \geq 0$ such that if Σ is a non-positively curved surface with $\text{inj}(\Sigma) \geq \eta$, $\text{area}(\Sigma) \leq A$ then $\text{diam}(\Sigma) \leq L$.*

Proof. This is standard argument. Given any $x, y \in \Sigma$, we connect x to y by a shortest geodesic, α . The $(\eta/2)$ -neighbourhood of α must then be a regular neighbourhood, whose area is at least $\eta d(x, y)$. It follows that $d(x, y) \leq A/\eta$. \square

Recall that $d_{\mathcal{T}}(\Sigma, \Sigma')$ is defined as logarithm of the minimal quasi-conformal deformation of a quasiconformal homeomorphism from Σ to Σ' respecting the marking. We note:

Lemma 19.2. *Given $A, \eta > 0$ and $k \geq 1$, there is some $\eta' > 0$, such that if Σ, Σ' are non-positively curved surfaces each of area A , with $d_{\mathcal{T}}(\Sigma, \Sigma') \leq k$ and $\text{inj}(\Sigma) \geq \eta$, then $\text{inj}(\Sigma') \geq \eta'$. Here, η' depends on A, η, k and $\text{genus}(\Sigma) = \text{genus}(\Sigma')$.*

Proof. This will be a simple consequence of Lemma 19.6, as we note after the proof thereof. \square

Lemma 19.3. *Suppose $A, \eta > 0$ and $k \geq 0$. Suppose Σ, Σ' are non-positively curved surfaces of area A and injectivity radius at least η .*

Suppose $f : \Sigma \rightarrow \Sigma'$ is a k -quasiconformal map. Let $\tilde{f} : \tilde{\Sigma} \rightarrow \tilde{\Sigma}'$ be a lift of f to the universal covers. Then \tilde{f} is a quasi-isometry whose parameters depend only on A, η, k and $\text{genus}(\Sigma) = \text{genus}(\Sigma')$.

This will be proven at the end of this section.

Conversely, we note:

Lemma 19.4. *Suppose that $\eta, A > 0$. Suppose that Σ and Σ' are non-positively curved surfaces of area A and $\text{inj}(\Sigma) \geq \eta'$ and that there is an equivariant quasi-isometry from the universal cover of Σ to the universal cover of Σ' . Then, $d_{\mathcal{T}}(\Sigma, \Sigma') \leq k$, where η and k depend on $\eta, A, \text{genus}(\Sigma)$ and the parameters of the quasi-isometry.*

Proof. Replacing Σ and Σ' by the hyperbolic surfaces in the same conformal class and applying Lemmas 19.2 and 19.3, we see that it is enough to prove Lemma 19.4 for hyperbolic surfaces.

Now, the length of a closed geodesic on a surface is the stable length of the corresponding element of Γ . The ratio of such lengths under quasi-isometry remains bounded (depending on the parameters of quasi-isometry). In other words, the ratio of lengths of the corresponding closed geodesics in Σ and Σ' remain bounded, and the result follows easily. \square

In summary, we have seen that two non-positively curved surfaces of the same area are a bounded distance apart in Teichmüller space if and only if the universal covers are quasi-isometric, where we have uniform control over the various parameters involved. This control depends on the injectivity radius of one of the surfaces (and hence both).

We next recall some basic facts regarding moduli and extremal length [LV].

Let Σ be any Riemann surface (not necessarily compact). Let Π be any set of paths in Σ . The *extremal length* of Π can be defined as $\sup\{\inf\{\text{length}(\gamma)^2 \mid \gamma \in \Pi\}\}$ where the supremum is taken over all Riemannian pseudometrics compatible with the given conformal structure. (We allow for isolated cone singularities.) If R is a conformal annulus, we define its *modulus*, $\text{mod}(R)$, as the reciprocal of the extremal length of the set of essential closed curves. Clearly if $Q \subseteq R$ is an essential subannulus, then $\text{mod}(Q) \leq \text{mod}(R)$.

If $A, B \subseteq \Sigma$ are closed and non-empty we define the *capacity*, $\text{cap}_{\Sigma}(A, B)$, to be the extremal length of the set of curves connecting A to B . If $R = \Sigma \setminus (A \cup B)$ is an annulus, then $\text{mod}(R) = \text{cap}_{\Sigma}(A, B)$. If $\Sigma = \mathbb{C}$, $A = N(0, r)$ and $B = \mathbb{C} \setminus N(0, s)$, then $\text{mod}(R) = \frac{1}{2\pi} \log(s/r)$.

If R is the quotient of the hyperbolic plane by a loxodromic element of translation distance l , then $\text{mod}(R) = \pi/l$. As a corollary, we note

that if Σ is any complete hyperbolic surface, and $R \subseteq \Sigma$ is an essential annulus homotopic to a closed geodesic of length l , then $\text{mod}(R) \leq \pi/l$. We write $M(\Sigma) = \max\{\text{mod}(R)\}$, where R varies over essential annuli in Σ . We see that $M(\Sigma) \leq \pi/2 \text{inj}(\Sigma)$.

In fact, we get similar inequalities in the case of any complete non-positively curved surface of finite area A . In this case, $\text{mod}(R) \leq A/l^2$, and so $M(\Sigma) \leq A/4 \text{inj}(\Sigma)^2$.

Suppose now that Σ is a non-positively curved surface (riemannian away from isolated cone singularities). Given $x \in \Sigma$, let $C(x, r) = \{y \in \Sigma \mid d(x, y) = r\}$. Thus, $C(x, r)$ is a rectifiable (indeed piecewise smooth) set. In what follows, we shall assume that there is some $\lambda \geq 0$ such that $\text{length}(C(x, r)) \leq \lambda r$. Integrating, we see that $\text{area}(N(x, r)) \leq \frac{1}{2}\lambda r^2$. In particular, $\text{area}(\Sigma) \leq \frac{1}{2}\lambda \text{diam}(\Sigma)^2$. This situation arises, for example, if the total curvature of Σ (i.e. the integral of scalar curvature with respect to area) no less than some negative constant, μ . In this case, we can take $\lambda = 2\pi - \mu$. (The worst case being a euclidean cone with cone angle $2\pi - \mu$ at x .) Note that if Σ is compact, then the total curvature is $2\pi\chi(\Sigma)$, where $\chi(\Sigma)$ is the Euler characteristic. and so we can take $\lambda = 2\pi(1 - \chi(\Sigma))$.

Lemma 19.5. *Suppose Σ is as above, and that $\alpha, \beta \subseteq \Sigma$ are simple closed curves bounding an annulus, R , and that $x \in \Sigma$ with $\alpha \subseteq N(x, r)$ and $d(x, \beta) \geq s$. Then $\text{mod}(R) \geq \frac{1}{\lambda} \log(s/r)$.*

Proof. Given $y \in R$, let $\nu(y) = 1/d(x, y)$ if $r \leq d(x, y) \leq s$ and $\nu(y) = 0$ otherwise. We obtain a (singular) riemannian pseudometric by scaling the original metric locally by the function ν . With respect to the new pseudometric, we have $\text{area}(R) \leq \int_r^s (\lambda t)(1/t^2) dt = \lambda \log(s/r)$. Moreover, any path from α to β in R has length at least $\int_r^s (1/t) dt = \log(s/r)$. Thus, $\text{mod}(R) = \text{cap}_R(\alpha, \beta) \geq (\log(s/r))^2 / \lambda \log(s/r) = \frac{1}{\lambda} \log(s/r)$. \square

Lemma 19.6. *Suppose $M > 0$ and $g \geq 2$. There is a function $D : [0, \infty) \rightarrow [0, \infty)$ with $D(0) = 0$ depending on M and g which have the following property. Suppose (Σ, d) is a closed non-positively curved surface with $\text{genus}(\Sigma) = g$, $\text{mod}(\Sigma) \leq M$ and $\text{inj}(\Sigma) \leq \eta$. Then $\text{diam}(\Sigma) \leq D(\eta)$.*

Proof. There is a bound, n , depending on $\text{genus}(\Sigma)$, on the number of pairwise non-homotopic simple closed curves one can embed disjointly into Σ . Let $\lambda = 2\pi(1 - \chi(\Sigma)) = 2\pi(\text{genus}(\Sigma) + 3)$, and let k be any constant bigger than $e^{\lambda M}$. We claim that $\text{diam}(\Sigma) \leq D(\eta) = 2(1 + \lambda)\eta k^n$.

To this end, choose $x \in \Sigma$ so that $N(x, \eta)$ is not a disc. If the claim is false, then there is some $y \in \Sigma$ with $d(x, y) > (1 + \lambda)\eta k^n$. For each $i \in \{0, \dots, n\}$, let $r_i = \eta k^i$. Now, y lies in a component, B_i , of $\Sigma \setminus N(x, r_i)$. Clearly $\text{fr}(B_n) \subseteq C(x, r_n)$, and so $\text{length}(\text{fr}(B_n)) \leq \lambda r_n$. It follows that C_n cannot be a disc. For otherwise (using standard facts of non-positive curvature) we would have $d(y, \text{fr}(B_n)) \leq \text{length}(\text{fr}(B_n)) \leq \eta r_n$, so that $d(x, y) \leq r_n + \lambda r_n = (1 + \lambda)\eta k^n$, contradicting the choice of y .

Thus, for each y , we know that neither B_n nor its complement is a disc. We can therefore choose an essential simple closed curve β_i arbitrarily close to $\text{fr}(B_i)$. From the choice of n we can find $0 \leq i < j \leq n$, so that β_i and β_j are homotopic and hence bound an annulus, R . But now $r_j/r_i \geq k > e^{\lambda M}$ and so, by Lemma 19.2, we must have $\text{mod}(R) > \frac{1}{\lambda} \log(e^{\lambda M}) = M \geq M(\Sigma)$ giving a contradiction. \square

Note that $\text{area}(\Sigma) \leq \lambda(\text{diam}(\Sigma))^2 \leq \lambda D(\eta)^2$. If Σ' is another such surface, and $f: \Sigma \rightarrow \Sigma'$ is k -quasiconformal, then $\frac{1}{k}M(\Sigma) \leq M(\Sigma') \leq kM(\Sigma)$. Lemma 19.2 is now an immediate consequence.

We now prove Lemma 19.3:

Proof of Lemma 19.3. Again, we let $\lambda = 2\pi(1 - \chi(\Sigma))$.

It will be enough to place a bound on $l = d(\tilde{f}(x), \tilde{f}(y))$ for $x, y \in \tilde{\Sigma}$ with $d(x, y) < \eta$. The same statement then follows for arbitrary x, y , any by symmetry, on swapping Σ and Σ' .

To this end, let D be the disc of radius η centred on the midpoint of $[x, y]$ and let R be the annulus $D \setminus [x, y]$. By Lemma 19.5, $\text{mod}(R) \geq \frac{1}{\lambda} \log 2$, and so $\text{mod}(\tilde{f}(R)) \geq \frac{1}{\lambda k} \log 2$. Since $\text{inj}(\Sigma) \geq \eta$, D and hence R meets no image of itself under the action of Γ . We see that the same must be true of $\tilde{f}(R)$ and so $\text{area}(\tilde{f}(R)) \leq A$.

Suppose γ is an essential curve in $\tilde{f}(R)$. Since the projection to $[\tilde{f}(x), \tilde{f}(y)]$ in $\tilde{\Sigma}$ is 1-lipschitz, we must have $\text{length}(\gamma) \leq 2d(\tilde{f}(x), \tilde{f}(y)) = 2l$. Now, $1/\text{mod}(\tilde{f}(R))$ is the extremal set of such curves. Thus, $1/\text{mod}(\tilde{f}(R)) \geq (2l)^2/A$, and so l^2 is bounded above by $A/4 \text{mod}(\tilde{f}(R)) \leq A\lambda k/4 \log 2$ as required. \square

20. TEICHMÜLLER GEODESICS

In this section, we give one of the main results of this paper, namely Theorem 20.1.

Given $\eta > 0$, and let $\text{thick}_\eta(\mathcal{T})$ be the η -thick part of Teichmüller space, \mathcal{T} , that is, the set of points corresponding to hyperbolic surfaces of injectivity radius at least η . Let $I \subseteq \mathbb{R}$ be an interval.

Definition. A *stable* path is a bilipschitz path from I to $\text{thick}_\eta(\mathcal{T})$ for some $\eta > 0$.

(Note that mapping into the thick part is the same as mapping into a bounded subset of moduli space in the quotient.)

Definition. A Lipschitz path, $\beta : I \rightarrow \mathcal{T}$, is a *pseudogeodesic* if there is a bilipschitz homeomorphism, $s : I \rightarrow J$, to another interval, J , and a Teichmüller (geodesic) ray $\alpha : J \rightarrow \mathcal{T}$ such that $d_{\mathcal{T}}(\alpha(s(t)), \beta(t))$ is bounded for all $t \in I$.

The principal cases of interest are when I is either \mathbb{R} or $[0, \infty)$, in which case, we can take $J = I$.

Out of a stable path, $\beta : I \rightarrow \mathcal{T}$, we will construct a stack, $\Xi(\beta)$, out of the hyperbolic surfaces corresponding to $\beta(i)$ for $i \in \mathcal{I} = I \cap \mathbb{Z}$. (A reinterpretation of this, in terms of metrics on $\Sigma \times I$ will be described in Section 22.) First, we describe this in general terms, using non-positively curved surfaces, since these will be needed later.

Suppose $(\Sigma_i)_{i \in \mathcal{I}}$ is a stable sequence of non-positively curved surfaces of fixed area A . We have $\text{inj}(\Sigma_i) \geq \eta > 0$ for all i . Let $X_i = \tilde{\Sigma}_i$ be the universal cover of Σ_i . By Lemma 19.4, there are uniform equivariant quasi-isometries, $\phi : X_i \rightarrow X_{i+1}$. We can thus construct the stack, $\Xi = \bigvee_{i \in \mathcal{I}} X_i$ as described in Section 6. The construction can be assumed to be equivariant. Any other choice of quasi-isometries, $\phi' : X_i \rightarrow X_{i+1}$ will be equivalent to ϕ (i.e. ϕ and ϕ' remain a bounded distance apart, see Lemma 4.8). Thus, the resulting stack will be quasi-isometric to Ξ by an equivariant sheet-preserving quasi-isometry.

Suppose Σ'_i is another sequence of non-positively curved surfaces with $d_{\mathcal{T}}(\Sigma_i, \Sigma'_i)$ bounded above. It follows that $(\Sigma'_i)_i$ is uniformly stable and that $d_{\mathcal{T}}(\Sigma'_i, \Sigma'_{i+1})$ is bounded above. We can similarly construct a stack, Ξ' , from the sequence $(\Sigma'_i)_i$. By Lemma 19.4, there is an equivariant quasi-isometry from X_i to $X'_i = \tilde{\Sigma}'_i$. As with Lemma 16.2, we can assemble these to give an equivariant sheet-preserving quasi-isometry from Ξ to Ξ' .

Suppose that $\alpha : I \rightarrow \mathcal{T}$ is a stable Lipschitz path. We apply the above construction to the sequence $(\alpha(i))_i$ to give a stack $\Xi(\alpha)$, well defined up to equivariant sheet-preserving quasi-isometry. If we reparameterise α by precomposing by a bilipschitz homeomorphism of I , then we get another stack equivariantly quasi-isometric to $\Xi(\alpha)$. Such a quasi-isometry is equivalent to an equivariant quasi-isometry.

We can now state:

Theorem 20.1. *Suppose that $I = \mathbb{R}$ or $I = [0, \infty)$, and that $\beta : I \rightarrow \mathcal{T}$ is a stable geodesic path. Then β is a pseudogeodesic if and only if the stack $\Xi(\beta)$ is hyperbolic.*

We shall see that the parameters of hyperbolicity depend on those of quasigeodesicity and vice versa (as well as $\text{genus}(\Sigma)$, the lower bound on injectivity radius, the area of the surfaces in the construction of $\Xi(\beta)$, and the bilipschitz constant).

As an immediate corollary, we have:

Corollary 20.2. *Any hyperbolic surface stack of positive injectivity radius is equivariantly quasi-isometric to the singular Sol stack constructed from a Teichmüller ray.*

Such stacks were described in Section 18, and are necessarily hyperbolic (Lemma 18.11). We already have enough information to prove one direction of Theorem 20.1. Suppose that $\beta : I \rightarrow \mathcal{T}$ is a pseudogeodesic, and α a Teichmüller ray given by this assumption. From the above discussion, we know that $\Xi(\alpha)$ and $\Xi(\beta)$ are quasi-isometric. Since $\Xi(\alpha)$ is hyperbolic, it follows that $\Xi(\beta)$ is hyperbolic.

We next move on to prove the other direction of Theorem 20.1. We first deal with the case where $I = [0, \infty)$.

Let us suppose therefore that $\beta : [0, \infty) \rightarrow \mathcal{T}$ is a stable path. We assume $\Xi(\beta)$ to be hyperbolic. Let Λ_Σ be the ending lamination (as defined in Section 17). From the theory of quadratic differentials [Ga], we know that there is a grid surface (Σ, d) which induces the conformal structure, $\beta(0)$, on Σ and whose horizontal foliation determines the lamination, Λ_Σ , in the manner described in Section 18. We shall normalise so that $\text{area}(\Sigma, d) = 1$. (Note that this involves a choice of transverse measure to Λ_Σ — a grid surface structure is equivalent to a conformal structure together with a quadratic differential, which is in turn equivalent to a conformal structure together with a transversely measured foliation. However by Proposition 17.4, Λ_Σ is regular, and hence uniquely ergodic, so there is a unique measure up to scale.) Now, let $\alpha(t)$ be the structure (Σ, d_t) . Thus α is a Teichmüller geodesic.

Let us suppose, for the moment, that:

Lemma 20.3. *For all $t \in [0, \infty)$, $\text{inj}(\Sigma, d_t)$ is bounded below by a uniform constant $\eta > 0$.*

It now follows, by Proposition 18.8, that $\Xi(\alpha)$ is uniformly hyperbolic. By Lemma 18.9, its ending lamination is precisely Λ_Σ . By Theorem 15.1, there is an equivariant stratified quasi-isometry from $\Xi(\alpha)$ to $\Xi(\beta)$. Thus, we get a quasi-isometry $s : \mathbb{N} \rightarrow \mathbb{N}$, and for all $i \in \mathbb{N}$, an equivariant quasi-isometry from the universal cover of (Σ, d_i) to that of

the hyperbolic surface corresponding to $\beta(s(i))$. Thus, by Lemma 19.4, $d_{\mathcal{T}}(\alpha(i), \beta(s(i)))$ is uniformly bounded. Now, we can assume that s is strictly increasing and that $s(0) = 0$. Thus by interpolating to a bilipschitz homeomorphism of $[0, \infty)$, we deduce that β is a pseudogeodesic as required.

We have thus proven Theorem 20.1 modulo Lemma 20.2. Lemma 20.3 relies on the following result:

Lemma 20.4. *Let Σ be a compact Riemann surface. Let d be a grid structure, and ρ a hyperbolic metric inducing the given conformal structure on Σ . Let Λ_{Σ} be the hyperbolic lamination corresponding to the horizontal foliation of (Σ, d) . Then Λ_{Σ} is regular in (Σ, ρ) if and only if the horizontal foliation is regular in (Σ, d) .*

This is similar to, and can be proven by similar means to Proposition 13.8. Here we only need and prove the “if” part. The converse follows by similar arguments. Let (X, d) be the universal cover of (Σ, d) and let (\mathbb{H}^2, ρ) be the universal cover of (Σ, ρ) . By Lemma 19.3, there is a uniform equivariant quasi-isometry from X to \mathbb{H}^2 , which extends to a homeomorphism $\phi : \partial X \rightarrow \partial \mathbb{H}^2$. Each horizontal leaf of X gives us a stratum of the lamination Λ . After making minor adjustments if necessary, we can confine ourselves to considering non-singular leaves. Thus, to each non-singular leaf, $\lambda \in \Lambda$, we can associate a (generic) leaf $\phi(\lambda)$ in \mathbb{H}^2 . Given $x \in \lambda$ we can associate a point, denoted $\phi(x)$, in $\phi(\lambda)$ exactly as in the proof of Proposition 13.8. The resulting map ϕ need not be continuous as we move from one leaf to another, however, nearby points will be mapped to points a uniformly bounded distance apart. If $E \subseteq \lambda$ is a horizontal interval with endpoints $x, y \in \lambda$, we write $\phi(E)$ for the horizontal interval with endpoints $\phi(x)$ and $\phi(y)$. If α is a vertical arc with endpoints w and z and with $\alpha \cap E = \emptyset$, then $\phi(\alpha) \cap \text{mid}(E, k) = \emptyset$ for some uniform $k \geq 0$, where $\phi(\alpha)$ denotes the geodesic $[\phi(w), \phi(z)]$.

Suppose now that E is a horizontal interval in (Σ, d) and that γ is a closed curve in Σ with $E \cap \gamma = \emptyset$. By a construction similar to that of Proposition 13.10, we can assume that γ consists of an alternating sequence of a bounded number of horizontal and vertical intervals, F_i and α_i respectively. The vertical intervals are of uniformly bounded length. We now lift everything to (X, d) . We get an equivariant set of lifts, \tilde{E} of E and \tilde{F}_i of F_i and $\tilde{\alpha}_i$ of α_i . We can assume, for convenience, that the \tilde{E} and \tilde{F}_i all lie in different non-singular leaves. We map them across to disjoint horizontal intervals $\phi(\tilde{E})$ and $\phi(\tilde{F}_i)$ in (\mathbb{H}^2, ρ) . We connect together the $\phi(\tilde{F}_i)$ in \mathbb{H}^2 with geodesic segments of the form $\phi(\tilde{\alpha}_i)$ to give us a path which projects to a closed curve γ' in (Σ, ρ) .

Now γ' does not intersect $\text{mid}(E', k)$, where E' is the projection of (any copy of) $\phi(\tilde{E})$ to (Σ, ρ) . From the stability condition on (Σ, ρ) , we see that the ρ -length of E' is bounded by a linear function of the ρ -length of γ' . We have a linear control on the deformation of lengths under ϕ , and so the stability condition for (Σ, d) now follows.

Proof of Theorem 20.1 for $I = [0, \infty)$. By Proposition 17.4, we know that the lamination Λ_Σ is regular, as defined in Section 13. Lemma 20.4 now tells us that the horizontal foliation of the corresponding grid surface is regular, as described in Section 18. It now follows by Lemma 18.2 that $\text{inj}(\Sigma_t, d)$ is bounded below by a positive constant, as required. \square

We finally need to deduce the case of bi-infinite geodesics from the semi-infinite case. To this end, let $\beta : \mathbb{R} \rightarrow \mathcal{T}$ be a stable path such that $\Xi(\beta)$ is hyperbolic. Let $\beta^+ = \beta|_{[0, \infty)}$ and $\beta^- = \beta|_{(-\infty, 0]}$. Now, $\Xi(\beta) = \Xi(\beta^+) \cup \Xi(\beta^-)$ and $X_0 = \Xi(\beta^+) \cap \Xi(\beta^-)$. By the Proposition 6.7, $\Xi(\beta^+)$ is hyperbolic, and so by the above, we see that β^+ is pseudogeodesic. The same applies to $\beta|_{[t, \infty)}$ for all $t \in \mathbb{R}$. In particular, for each $n \in \mathbb{N}$, we see that after uniform bilipschitz reparameterisation, $\beta|_{[-n, \infty)}$ remains a uniformly bounded distance from a geodesic, $\alpha_n : [-n, \infty) \rightarrow \mathcal{T}$. By a diagonal sequence argument we can suppose that (after bilipschitz reparameterisation) we have $\alpha_n(i) \rightarrow \alpha(i) \in \mathcal{T}$ as $n \rightarrow \infty$ for all $i \in \mathbb{Z}$. We can extend $\alpha : \mathbb{Z} \rightarrow \mathcal{T}$ to a geodesic $\alpha : \mathbb{R} \rightarrow \mathcal{T}$, and we see that β remains a bounded distance from α . This proves Theorem 20.1 in the bi-infinite case.

21. A CRITERION FOR RECOGNISING STABLE PSEUDOGEODESICS

In this section, we give a simple geometrical criterion for recognising Teichmüller pseudogeodesics among stable sequences of surfaces (or stable paths in Teichmüller space).

Let Σ be a compact surface, and let $(d_i)_{i \in \mathbb{N}}$ be a sequence of hyperbolic metrics on Σ . We write $(\Sigma_i)_i$ for the surface Σ with the metric d_i . We say that the sequence is *stable* if $\text{inj}(\Sigma_i)$ is bounded below and the Teichmüller distance between Σ_i and Σ_{i+1} is bounded above. It is a *pseudogeodesic* if it remains a bounded distance from a Teichmüller ray and the Teichmüller distance between Σ_i and Σ_j is bounded below by a linear function of $|i - j|$. (This is a reformulation in terms of sequences of the definition of Section 20.)

Given an essential closed curve, α , in Σ , we write $l_i(\alpha)$ for length of the simple closed geodesic in the same homotopy class. This is equal

to the stable length (here the minimal translation distance), $\text{sl}(g)$, of the corresponding element, $g \in \Gamma$.

We show:

Theorem 21.1. *Let \mathcal{I} be either \mathbb{Z} or \mathbb{N} . Let $(\Sigma_i)_{i \in \mathcal{I}}$ be a stable sequence of surfaces. Then $(\Sigma_i)_i$ is pseudogeodesic if and only if for each essential simple closed curve α in Σ , the function $[i \mapsto l_i(\alpha)]$ is uniformly quasiconvex.*

As usual the parameters of pseudogeodesicity depend on the constants of quasiconvexity and vice versa. We can also rephrase the statement in terms of stable paths. Theorem 2.12 is an immediate consequence, on taking $\Sigma_i = \beta(i)$, where β is a stable path in \mathcal{T} .

The proof relies on the following observation:

Lemma 21.2. *Let $\Sigma = \mathbb{H}^2/\Gamma$ be a compact hyperbolic surface. There is a constant, k , depending only on $\text{inj}(\Sigma)$ and $\text{genus}(\Sigma)$ with the following property. If $x \in \mathbb{H}^2$ and $g \in \Gamma$, there is some $h \in \Gamma$ such that $\rho(x, hx) \leq k$ and $\rho(hx, gx) \leq \text{sl}(h^{-1}g) + k$.*

Proof. First note that, at the cost of increasing k , it is enough to prove this for x lying in a given Γ -orbit. Now it is a simple exercise to show that we can find constants, $\theta > 0$ and $l \geq 0$, depending only on $\text{inj}(\Sigma)$ and $\text{genus}(\Sigma)$, so that Σ contains two closed geodesics, α and β , each of length at most l , and meeting at an angle of at least θ at some point $y \in \Sigma$. As observed above, we can suppose that y is the projection of x .

Now elementary hyperbolic geometry gives us another constant, r , such that any broken geodesic path whose segments all have length at least r and where adjacent segments meet at angles at least $\theta/3$ remains a uniformly bounded distance from a geodesic. We can suppose that $r \geq l$. Thus, there are multiples m, n respectively of the curves α, β , so that α^m and β^n each have length between r and $2r$.

Now consider the geodesic path, γ , based at y , that corresponds to the element $g \in \Gamma$. We can suppose that $\text{length}(\gamma) = \rho(x, gx) \geq r$ (otherwise we simply take $h = 1$ and $k \geq r$). Let ϕ be the angle of γ at y .

If $\phi \geq \theta/3$, then γ lifts to a broken geodesic a bounded distance from the axis of g . It follows that $\text{length}(\gamma) - \text{sl}(g)$ is uniformly bounded, so we again take $h = 1$.

If on the other hand, $\phi < \theta/3$, then both endpoints of γ meet either α or β at an angle at least $\theta/3$. Suppose, without loss of generality, it is α . Let $h \in \Gamma$ be the element represented by α^m with basepoint at y , so that $\rho(x, hx) \leq 2r$. The closed path $\gamma \cup \alpha^{-m}$ lifts to a broken

geodesic a bounded distance from the axis of $h^{-1}g$. Now $\rho(hx, gx) = \rho(x, h^{-1}gx) \leq \text{length}(\gamma \cup \alpha^{-m})$, and so $\rho(hx, gx) - \text{sl}(h^{-1}g)$ is uniformly bounded, as required. \square

Now let $(\Sigma_i)_i$ be a stable sequence of surfaces, and let $\Xi = (X_i)_i$ be the associated stack of hyperbolic planes.

One direction of Theorem 21.1 is fairly elementary. If $(\Sigma_i)_i$ is pseudogeodesic, then by Theorem 20.1, Ξ is hyperbolic. If α is an essential closed curve on Σ and g is the corresponding element of Γ , then we can find a uniform chain $(x_i)_i$ so that x_i lies on the axis of g in X_i . Thus $l_i(\alpha) = \text{sl}(g) = \rho(x_i, gx_i)$, so $[i \mapsto l_i(\alpha)]$ is uniformly quasiconvex.

For the converse, let $(x_i)_i$ be any chain in Ξ . It is enough to show that the function $[i \mapsto \rho_i(x_i, gx_i)]$ is uniformly quasiconvex for all $g \in \Gamma$, since by Corollary 6.5, it then follows that Ξ is hyperbolic and we can apply Theorem 20.1.

Suppose $g \in \Gamma$ is fixed and $i \in \mathcal{I}$. By Lemma 21.2 applied to Σ_i , there is some $h \in \Gamma$ such that $\rho_i(x_i, hx_i)$ and $\rho_i(hx_i, gx_i) - \text{sl}_i(h^{-1}g)$ are uniformly bounded above. Now $\rho_j(x_j, hx_j) \leq h(|i - j|)$ for some uniform non-decreasing function, $h : \mathcal{I} \rightarrow \mathbb{R}$.

Now $\text{sl}_j(h^{-1}g) \leq \rho_j(hx_j, gx_j) \leq \rho_j(x_j, gx_j) + \rho_j(x_j, hx_j) \leq \rho_j(x_j, gx_j) + h(|i - j|)$. Moreover, $|\text{sl}_i(g^{-1}h) - \rho_i(x_i, gx_i)|$ is uniformly bounded (without loss of generality by $h(0)$). Thus, by Lemma 5.3, the function $[i \mapsto \rho_i(x_i, gx_i)]$ satisfies condition (Q1) of quasiconvexity. Since it automatically satisfies (Q2), we see that it is uniformly quasiconvex as required.

This proves Theorem 21.1.

22. ENDS OF 3-MANIFOLDS

In this section, we return in more detail to the subject discussed in the introduction, namely the geometry of hyperbolic 3-manifolds. Let $M = \mathbb{H}^3/\Gamma$ be a complete orientable 3-manifold, with $\Gamma \equiv \pi_1(M)$ finitely generated. It follows by tameness [Bon, Ag, CalG], that M is topologically finite, i.e. homeomorphic to the interior of a compact manifold. Moreover, the Ending Lamination Theorem [Min6, BroCM1, BroCM2] together with [O] or [NS] gives a complete description of M in terms of its topology and end invariants. For many purposes, one can effectively reduce to the case where $\Gamma = \pi_1(\Sigma)$ for a closed orientable surface, Σ , so that, by tameness [Bon], M is homeomorphic to $\Sigma \times \mathbb{R}$. We consider here the case where M has no cusps. Let H be the convex core of M (the quotient of the convex hull of the limit set). Then, up to isotopy in M , we can assume that H has the form $H \equiv \Sigma \times I \subseteq M$, where $I \subseteq \mathbb{R}$ is a closed interval. If I is compact,

then M is quasifuchsian, and this case is understood by earlier work of Ahlfors, Bers, Marden etc. (see, for example, [Mar]). Otherwise, we can assume, without loss of generality, that $I = [0, \infty)$ or $I = \mathbb{R}$. In the former “singly degenerate” case, one of the end invariants lies in \mathcal{T} , and the other, corresponding to the infinite end of H , is a lamination. The former invariant lies a bounded distance in \mathcal{T} from the hyperbolic structure induced on the convex surface, $\partial H \equiv \Sigma \times \{0\}$. In second “doubly degenerate” case, where $H = M$, there are two (necessarily distinct) ending laminations. These end invariants were described in [T, Bon]. The Ending Lamination Theorem in this case tells us that these end invariants determine M up to isometry.

The general strategy for proving the Ending Lamination Theorem, as used in [Min1, Min6, BroCM1], is to construct a “model space”, P , with a bilipschitz map to H . In fact, all we really require is a Γ -equivariant quasi-isometry between the universal covers, \tilde{P} of P , and $\tilde{H} \subseteq \mathbb{H}^3$. The model space, P , can be chosen to depend only on the end invariants of M . By a result of Sullivan [Su], this is sufficient to prove the Ending Lamination Theorem for such manifolds. Here we will consider only the bounded geometry case, as described in [Min1], that is, where M has positive injectivity radius. We consider how models are constructed out of a path in Teichmüller spaces.

Let $\eta > 0$ be some constant, less than the 2-dimensional Margulis constant. Let $\text{thick}_\eta(\mathcal{T})$ be the η -thick part of \mathcal{T} , that is, the closed subset corresponding to hyperbolic surfaces of injectivity radius at least η .

Lemma 22.1. *Given $\mu > 0$, there is some $\nu > 0$, depending only on η and μ , such that given any μ -lipschitz path, $\beta : I \longrightarrow \text{thick}_\eta(\mathcal{T})$, there is a smooth riemannian metric, ρ , on $\Sigma \times I$ with the following properties.*

- (1) *For all $x \in \Sigma$, the map $[t \mapsto (x, t)] : I \longrightarrow \{x\} \times I$ is ν -bilipschitz.*
- (2) *For all $t, u \in I$ with $|t - u| \leq 1$, the map $[(x, t) \mapsto (x, u)] : \Sigma \times \{t\} \longrightarrow \Sigma \times \{u\}$ is ν -bilipschitz.*
- (3) *For all t , there is a ν -bilipschitz map, in the natural homotopy class, from $\Sigma \times \{t\}$ to the hyperbolic structure on Σ represented by $\beta(t)$.*

Moreover, if ρ' is another such metric, then there is a uniformly bilipschitz diffeomorphism from $(\Sigma \times I, \rho)$ to $(\Sigma \times I, \rho')$, which is homotopic to the identity, and which preserves setwise each surface, $\Sigma \times \{t\}$.

We will write $P = P(\beta)$ for $\Sigma \times I$ with a riemannian metric of the sort given by Lemma 22.1. This is well defined up to bilipschitz diffeomorphism as described. One can construct such a metric by approximating β by a smooth path, and pulling back the canonical Σ -bundle over \mathcal{T} . Since we do not formally need Lemma 22.1 for the proofs of the main theorems, we will not give details here. In practice, we can work instead with equivariant coarse quasi-isometries involving stacks, though the description in terms of the model space $P(\beta)$ is more intuitive.

Let $\mathcal{I} = \mathbb{Z} \cap I$. We assume that $\mathcal{I} \neq \emptyset$. Given $i \in \mathcal{I}$, let $\Sigma_i = \Sigma \times \{i\} \subseteq P$, and let $X_i = \tilde{\Sigma}_i$. Then \tilde{P} is a stack with sheets, $(X_i)_{i \in \mathcal{I}}$. Note that the map $[(x, i) \mapsto (x, i + 1)] : X_i \rightarrow X_{i+1}$ is bilipschitz. Up to bounded distance, it agrees with the quasi-isometry from X_i to X_{i+1} , described for a general stack in Section 6. Note that we can reconstruct a stack for the sheets, $(X_i)_i$, and these quasi-isometries. Up to sheet-preserving quasi-isometry, this agrees with \tilde{P} . We have seen (Theorem 2.11 proven in Section 20), that if $I = [0, \infty)$ or $I = \mathbb{R}$, then \mathcal{P} is Gromov hyperbolic if and only if β is pseudogeodesic, that is, a bounded distance from a Teichmüller geodesic.

Let us return to the case of a hyperbolic 3-manifold, $M \cong \Sigma \times \mathbb{R}$, with no parabolics. Suppose that the positive end of M is simply degenerate. That is, $\Sigma \times [0, \infty)$ can be assumed to be contained in the convex core, H . This implies that there is a sequence of uniformly Lipschitz, in fact, 1-Lipschitz, homotopy equivalences, with $f_i(\Sigma_i)$ going out the end (see [T, Bon]).

Each Σ_i has a marking, determined by the induced homotopy equivalence with Σ , and hence determines a point in \mathcal{T} . The structure tends (in an appropriate sense) to an “ending lamination”. We can assume that for all $i > 0$, $f_i(\Sigma_i)$ homologically separates $f_{i-1}(\Sigma_{i-1})$ from $f_{i+1}(\Sigma_{i+1})$.

Suppose that $\text{inj}(M) \geq \eta$ for some $\eta > 0$. Then we will have $\Sigma_i \in \text{thick}_\eta(\mathcal{T})$. In particular, it follows that the diameters of the Σ_i are bounded, hence the diameters of the $f_i(\Sigma_i)$ are bounded above in terms of η and $\text{genus}(\Sigma)$. Moreover (by interpolation of pleated surfaces) we can assume that the distances between consecutive images, $f_i(\Sigma_i)$ and $f_{i+1}(\Sigma_{i+1})$ are bounded above and below, again in terms of η and $\text{genus}(\Sigma)$.

Suppose that M is singly degenerate. In this case, we can suppose, in addition, that f_0 is an isometry from Σ_0 to $\partial H \equiv \{0\}$. In this case, every point of H is a bounded distance from some $f_i(\Sigma_i)$. Let $\tilde{f}_i : X_i = \tilde{\Sigma}_i \rightarrow \tilde{H} \subseteq \mathbb{H}^3$ be the lift. Now a geometric limit argument (see [Min1] or [Bow3]) shows that $\tilde{f}_i : X_i \rightarrow H$ is a uniform

embedding. In particular, the relation defined by $x \sim y$ if and only if $\rho(\tilde{f}_i(x), \tilde{f}_{i+1}(y)) \leq r$, for $x \in X_i, y \in X_{i+1}$ is a uniform equivariant quasi-isometry from X_i to X_{i+1} , provided r is chosen large enough. This enables us to construct a stack, Ξ , with sheets $(X_i)_i$. We also note that the existence of these quasi-isometries implies that distance in \mathcal{T} between Σ_i and Σ_{i+1} is bounded above, so we can interpolate to give a lipschitz path $\beta : [0, \infty) \rightarrow \mathcal{T}$, which converges on the ending lamination. This enables us to construct a riemannian manifold, $P(\beta) \cong \Sigma \times [0, \infty)$, as in Lemma 22.1, with $\mathcal{P}(\beta)$ quasi-isometric to Ξ . We define a map $f : P \rightarrow H$, by setting $f(x, i) = f_i(x)$ for each $i \in \mathbb{N}$, and then interpolate linearly between (x, i) and $(x, i + 1)$ along a geodesic segment in H between $f_i(x)$ and $f_{i+1}(x)$ in the appropriate homotopy class. (This can be done equivariantly in the universal covers.) This gives us a uniformly lipschitz map $f : P \rightarrow H$, which lifts to $\tilde{f} : \tilde{P} \rightarrow \tilde{H} \subseteq \mathbb{H}^3$. It is easily checked that \tilde{f} is a quasi-isometry, and we can deduce that \tilde{P} is hyperbolic. (One can arrive at the same conclusion by constructing a quasi-isometry from the stack, Ξ , based on $(X_i)_i$ to \tilde{H} , which would bypass the construction of a riemannian manifold.) It follows by Theorem 2.11 that β is a bounded distance from a geodesic ray, α , and so we can replace $P(\beta)$ by $P(\alpha)$. (In this case, $P(\alpha)$ will be bilipschitz equivalent to the singular Sol geometry manifold described in Section 18.) We can now proceed as described in the introduction.

A similar construction works in the doubly degenerate case. In this case, we see that M is bilipschitz equivalent to $P(\alpha)$, where α is a bi-infinite geodesic whose limit points in $\partial\mathcal{T}$ are the end invariants of M . In particular, suppose that $M' = \mathbb{H}^3/\Gamma'$ is another hyperbolic 3-manifold with positive injectivity radius, and the same end invariants, then passing via $P(\alpha)$, we get a quasi-isometry from \mathbb{H}^3 to itself, which is equivariant with respect to the actions of Γ and Γ' . This gives rise to an equivariant quasiconformal map from $\partial\mathbb{H}^3$ to itself. From the result of Sullivan [Su], it follows that this must in fact be conformal. Therefore the actions of Γ and Γ' are conjugate by an isometry of \mathbb{H}^3 , and so M and M' are isometric, by an isometry in the natural proper homotopy class. This proves the Ending Lamination Theorem for doubly degenerate manifolds with positive injectivity radius.

In the singly degenerate case, we get an equivariant quasi-isometry between the convex hull of the limit sets. This time we need to deal with the geometrically finite end. One approach to this is described in [Min1] and [Min6]. A different approach is described in [Bow4]. The

quasifuchsian case is well known from earlier work of Ahlfors, Bers, Marden etc., see for example, [Mar].

REFERENCES

- [Ag] I. Agol, *Tameness and hyperbolic 3-manifolds* : preprint, Chicago (2004).
- [AlDP] R.C. Alperin, W. Dicks, J. Porti, *The boundary of the Gieseking tree in hyperbolic three-space* : Topology Appl. **93** (1999) 219–259.
- [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, *Singular euclidean structures on surfaces* : J. London Math. Soc. **44** (1991) 553–565.
- [Bow2] 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.
- [Bow3] B.H. Bowditch, *The Cannon-Thurston map for punctured-surface bundles* : Math. Z. **255** (2007) 35–76.
- [Bow4] B.H. Bowditch, *The ending lamination theorem* : preprint, Warwick 2011.
- [BriH] M.R. Bridson, A. Haefliger, *Metric spaces of non-positive curvature* : Grundlehren der Mathematischen Wissenschaften No. 319, Springer-Verlag (1999).
- [BroCM1] J.F. Brock, R.D. Canary, Y.N. Minsky, *The classification of Kleinian surface groups II: The Ending Lamination Conjecture* : preprint, 2004.
- [BroCM2] J.F. Brock, R.D. Canary, Y.N. Minsky, *The classification of finitely generated Kleinian groups* : in preparation.
- [CalG] D. Calegari, D. Gabai, *Shrinkwrapping and the taming of hyperbolic 3-manifolds* : J. Amer. Math. Soc. **19** (2006) 385–446
- [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.
- [CannT] J.W. Cannon, W.P. Thurston, *Group invariant Peano curves* : preprint (1989).
- [CasB] A.J. Casson, S.A. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston* : London Math. Soc. Student Texts No. 9, Cambridge University Press (1988).
- [CoDP] M. Coornaert, T. Delzant, A. Papadopoulos, *Les groupes hyperboliques de Gromov* : Lecture Notes in Mathematics No. 1441, Springer Verlag (1990).
- [Ga] F.P. Gardiner, *Teichmüller theory and quadratic differentials* : John Wiley & Sons (1987).
- [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 Riemannian and non-Riemannian spaces* : Progress in Mathematics No. 152, Birkhäuser (1998).
- [EP] D.B.A.Epstein, R.C.Penner, *Euclidean decompositions of noncompact hyperbolic manifolds* : J. Differential Geom. **27** (1988) 67-80.
- [IT] Y.Imayoshi, M.Taniguchi, *An introduction to Teichmüller spaces* : Springer-Verlag (1992).
- [K] S.P.Kerckhoff, *The asymptotic geometry of Teichmüller space* : Topology **19** (1980) 23-41.
- [LV] O.Lehto, K.I.Virtanen, *Quasiconformal mappings in the plane* (translated from German by K.W.Lucas) : Grundlehren der Mathematischen Wissenschaften No. 126, Springer-Verlag (1974).
- [Mar] A.Marden, *The geometry of finitely generated kleinian groups* : Ann. of Math. **99** (1974) 383-462.
- [Mas] H.Masur, *Uniquely ergodic quadratic differentials* : Comment. Math. Helv. **55** (1980) 255-266.
- [MasM] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103-149.
- [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, *Kleinian groups and the complex of curves* : Geom. Topol. **4** (2000) 117-148.
- [Min4] Y.N.Minsky, *Short geodesics and end invariants* : in “Comprehensive Research in Complex Dynamics and Related Fields”, (eds. M.Kisaka, S.Morosawa) RIMS Kokyuroka No. 1153 (2000) 1-20.
- [Min5] Y.N.Minsky, *Bounded geometry for Kleinian groups* : Invent. Math. **146** (2001) 143-192.
- [Min6] Y.N.Minsky, *The classification of Kleinian surface groups I : Models and bounds* : Ann. Math. **171** (2010) 1-107.
- [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 extentions* : Topology **37** (1998) 527-538.
- [Mj1] M.Mj, *Cannon-Thurston maps for pared manifolds of bounded geometry* : Geom. Topol. **13** (2009) 189-245.
- [Mj2] M.Mj, *Cannon-Thurston maps for surface groups I: amalgamation geometry and split geometry* : preprint, Belur Math (2006).
- [Mj3] M.Mj, *Cannon-Thurston maps for surface groups E: split geometry and the Minsky model* : preprint, Belur Math (2010).
- [Mo] L.Mosher, *Stable Teichmüller quasigeodesics and ending laminations* : Geom. Topol. **7** (2003) 33-90.
- [NS] H.Namazi, J.Souto, *Non-realizability and ending laminations proof of the density conjecture* : preprint (2010).
- [O] K.Ohshika, *Realising end invariants by limits of minimally parabolic geometrically finite groups* : Geom. Topol. **15** (2011) 827-890.

- [Sc] P.Scott, *The geometries of 3-manifolds* : Bull. London Math. Soc. **15** (1983) 401–487.
- [Sh] H.Short et al., *Notes on word hyperbolic groups* : in “Group theory from a geometrical viewpoint” (ed. E.Ghys, A.Haefliger, A.Verjovsky), World Scientific (1991) 3–63.
- [Su] D.Sullivan, *Quasiconformal homeomorphisms and dynamics E. Structural stability implies hyperbolicity for kleinian groups* : Acta Math. **155** (1985) 243–260.
- [T] W.P.Thurston, *The geometry and topology of 3-manifolds* : notes, Princeton (1979).

MATHEMATICS INSTITUTE, UNIVERSITY OF WARWICK, COVENTRY CV4 7AL,
GREAT BRITAIN