Annals of Mathematics, **167** (2008), 1055–1077

Dimension and rank for mapping class groups

By Jason A. Behrstock and Yair N. Minsky*

Dedicated to the memory of Candida Silveira.

Abstract

We study the large scale geometry of the mapping class group, $\mathcal{MCG}(S)$. Our main result is that for any asymptotic cone of $\mathcal{MCG}(S)$, the maximal dimension of locally compact subsets coincides with the maximal rank of free abelian subgroups of $\mathcal{MCG}(S)$. An application is a proof of Brock-Farb's Rank Conjecture which asserts that $\mathcal{MCG}(S)$ has quasi-flats of dimension N if and only if it has a rank N free abelian subgroup. (Hamenstadt has also given a proof of this conjecture, using different methods.) We also compute the maximum dimension of quasi-flats in Teichmuller space with the Weil-Petersson metric.

Introduction

The coarse geometric structure of a finitely generated group can be studied by passage to its *asymptotic cone*, which is a space obtained by a limiting process from sequences of rescalings of the group. This has played an important role in the quasi-isometric rigidity results of [DS], [KaL] [KlL], and others. In this paper we study the asymptotic cone $\mathcal{M}_\omega(S)$ of the mapping class group of a surface of finite type. Our main result is

Dimension Theorem. *The maximal topological dimension of a locally compact subset of the asymptotic cone of a mapping class group is equal to the maximal rank of an abelian subgroup.*

Note that [BLM] showed that the maximal rank of an abelian subgroup of a mapping class group of a surface with negative Euler characteristic is $3g-3+p$ where g is the genus and p the number of boundary components. This is also the number of components of a pants decomposition and hence the largest rank of a pure Dehn twist subgroup.

*First author supported by NSF grants DMS-0091675 and DMS-0604524. Second author supported by NSF grant DMS-0504019.

As an application we obtain a proof of the "geometric rank conjecture" for mapping class groups, formulated by Brock and Farb [BF], which states:

RANK THEOREM. *The geometric rank of the mapping class group of a surface of finite type is equal to the maximal rank of an abelian subgroup.*

Hamenstädt had previously announced a proof of the rank conjecture for mapping class groups, which has now appeared in [Ham]. Her proof uses the geometry of train tracks and establishes a homological version of the dimension theorem. Our methods are quite different from hers, and we hope that they will be of independent interest.

The geometric rank of a group G is defined as the largest n for which there exists a quasi-isometric embedding $\mathbb{Z}^n \to G$ (not necessarily a homomorphism), also known as an n-dimensional quasi-flat. It was proven in [FLM] that, in the mapping class group, maximal rank abelian subgroups are quasi-isometrically embedded—thereby giving a lower bound on the geometric rank. This was known when the Rank Conjecture was formulated; thus the conjecture was that the known lower bound for the geometric rank is sharp. The affirmation of this conjecture follows immediately from the dimension theorem and the observation that a quasi-flat, after passage to the asymptotic cone, becomes a bi-Lipschitz-embedded copy of $\mathbb{R}^n$.

We note that in general the maximum rank of (torsion-free) abelian subgroups of a given group does not yield either an upper or a lower bound on the geometric rank of that group. For instance, nonsolvable Baumslag-Solitar groups have geometric rank one [Bur], but contain rank two abelian subgroups. To obtain groups with geometric rank one, but no subgroup isomorphic to $\mathbb{Z}$, one may take any finitely generated infinite torsion group. The n-fold product of such a group with itself has n-dimensional quasi-flats, but no copies of $\mathbb{Z}^n$.

Similar in spirit to the above results, and making use of Brock's combinatorial model for the Weil-Petersson metric [Bro], we also prove:

DIMENSION THEOREM FOR TEICHMÜLLER SPACE. *Every locally compact subset of an asymptotic cone of Teichmüller space with the Weil-Petersson metric has topological dimension at most* $\lfloor \frac{3g+p-2}{2} \rfloor$.

The dimension theorem implies the following, which settles another conjecture of Brock-Farb.

RANK THEOREM FOR TEICHMÜLLER SPACE. *The geometric rank of the Weil-Petersson metric on the Teichmüller space of a surface of finite type is equal to* $\lfloor \frac{3g+p-2}{2} \rfloor$.

This conjecture was made by Brock-Farb after proving this result in the case $\lfloor \frac{3g+p-2}{2} \rfloor \leq 1$, by showing that in such cases Teichmüller space is δ-hyperbolic [BF]. (Alternate proofs of this result were obtained in [Be] and

[Ara].) We also note that the lower bound on the geometric rank of Teichmüller space is obtained in [BF].

Outline of the proof. For basic notation and background see Section 1. We will define a family $\mathcal{P}$ of subsets of $\mathcal{M}_\omega(S)$ with the following properties: Each $P \in \mathcal{P}$ comes equipped with a bi-Lipschitz homeomorphism to a product $F \times \mathcal{A}$, where

(1) F is an $\mathbb{R}$-tree;
(2) $\mathcal{A}$ is the asymptotic cone of the mapping class group of a (possibly disconnected) proper subsurface of S.

There will also be a Lipschitz map $\pi_P \colon \mathcal{M}_\omega(S) \to F$ such that:

(1) The restriction of π_P to P is projection to the first factor.
(2) π_P is locally constant in the complement of P.

These properties immediately imply that the subsets $\{t\} \times \mathcal{A}$ in $P = F \times \mathcal{A}$ separate $\mathcal{M}_\omega(S)$ globally.

The family $\mathcal{P}$ will also have the property that it *separates points*, that is: for every $x \neq y$ in $\mathcal{M}_\omega(S)$ there exists $P \in \mathcal{P}$ such that $\pi_P(x) \neq \pi_P(y)$.

Using induction, we will be able to show that locally compact subsets of $\mathcal{A}$ have dimension at most $r(S) - 1$, where $r(S)$ is the expected rank for $\mathcal{M}_\omega(S)$. The separation properties above together with a short lemma in dimension theory then imply that locally compact subsets of $\mathcal{M}_\omega(S)$ have dimension at most $r(S)$.

Section 1 will detail some background material on asymptotic cones and on the constructions used in Masur-Minsky [MM1, MM2] to study the coarse structure of the mapping class group. Section 2 introduces product regions in the group and in its asymptotic cone which correspond to cosets of curve stabilizers.

Section 3 introduces the $\mathbb{R}$-trees F, which were initially studied by Behrstock in [Be]. The regions $P \in \mathcal{P}$ will be constructed as subsets of the product regions of Section 2, in which one factor is restricted to a subset which is one of the $\mathbb{R}$-trees. The main technical result of the paper is Theorem 3.5, which constructs the projection maps π_P and establishes their locally constant properties. An almost immediate consequence is Theorem 3.6, which gives the family of separating sets whose dimension will be inductively controlled.

Section 4 applies Theorem 3.6 to prove the Dimension Theorem.

Section 5 applies the same techniques to prove a similar dimension bound for the asymptotic cone of a space known as the pants graph and to deduce a corresponding geometric rank statement there as well. These can be translated into results for Teichmüller space with its Weil-Petersson metric, by applying Brock's quasi-isometry [Bro] between the Weil-Petersson metric and the pants graph.

Acknowledgements. The authors are grateful to Lee Mosher for many insightful discussions, and for a simplification to the original proof of Theorem 3.5. We would also like to thank Benson Farb for helpful comments on an earlier draft.

1. Background

1.1. *Surfaces.* Let $S = S_{g,p}$ be a orientable compact connected surface of genus g and p boundary components. The mapping class group, $\mathcal{MCG}(S)$, is defined to be $\mathrm{Homeo}^+(S)/\mathrm{Homeo}_0(S)$, the orientation-preserving homeomorphisms up to isotopy. This group is finitely generated [Deh], [Bir] and for any finite generating set one considers the word metric in the usual way [Gro2], whence yielding a metric space which is unique up to quasi-isometry.

Throughout the remainder, we tacitly exclude the case of the closed torus $S_{1,0}$. Nonetheless, the Dimension Theorem does hold in this case since $\mathcal{MCG}(S_{1,0})$ is virtually free so that its asymptotic cones are all one dimensional and the largest rank of its free abelian subgroups is one.

Let $r(S)$ denote the largest rank of an abelian subgroup of $\mathcal{MCG}(S)$ when S has negative Euler characteristic. In [BLM], it was computed that $r(S) = 3g - 3 + p$ and it is easily seen that this rank is realized by any subgroup generated by Dehn twists on a maximal set of disjoint essential simple closed curves. Moreover, such subgroups are known to be quasi-isometrically embedded by results in [Mos], when S has punctures, and by [FLM] in the general case.

For an annulus let $r = 1$. For a disconnected subsurface $W \subset S$, with each component homotopically essential and not homotopic into the boundary, and no two annulus components homotopic to each other, let $r(W)$ be the sum of $r(W_i)$ over the components of W. We note that r is automatically additive over disjoint unions, and is monotonic with respect to inclusion.

1.2. *Quasi-isometries.* If (X_1, d_1) and (X_2, d_2) are metric spaces, a map $\phi\colon X_1 \to X_2$ is called a (K, C)-*quasi-isometric embedding* if for each $y, z \in X_1$ we have:

$$d_2(\phi(y), \phi(z)) \approx_{K,C} d_1(y, z). \tag{1.1}$$

Here the expression $a \approx_{K,C} b$ means $a/K - C \leq b \leq Ka + C$. We sometimes suppress K, C, writing just $a \approx b$ when this will not cause confusion.

We call ϕ a *quasi-isometry* if, additionally, there exists a constant $D \geq 0$ so that each $q \in X_2$ satisfies $d_2(q, \phi(X_1)) \leq D$, i.e., ϕ is almost onto. The property of being quasi-isometric is an equivalence relation on metric spaces.

1.3. *Subsurface projections and complexes of curves.* On any surface S, one may consider the *complex of curves of* S, denoted $\mathcal{C}(S)$. The complex of curves is a finite dimensional flag complex whose vertices correspond to nontrivial homotopy classes of nonperipheral, simple, closed curves and with edges between any pair of such curves which can be realized disjointly on S. In the cases where $r(S) \leq 1$ the definition must be modified slightly. When S is a one-holed torus or 4-holed sphere, any pair of curves intersect, so edges are placed between any pair of curves which realize the minimal possible intersection on S (1 for the torus, 2 for the sphere). With this modified definition, these curve complexes are the Farey graph. When S is the 3-holed sphere its curve complex is empty since S supports no simple closed curves. Finally, the case when S is an annulus will be important when S is a subsurface of a larger surface S'. We define $\mathcal{C}(S)$ by considering the annular cover $\tilde{S}'$ of S' in which S lifts homeomorphically. Now $\tilde{S}'$ has a natural compactification to a closed annulus, and we let vertices be paths connecting the boundary components of this annulus, up to homotopy rel endpoints. Edges are pairs of paths with disjoint interiors. With this definition, one obtains a complex quasi-isometric to $\mathbb{Z}$. (See [MM1] for further details.)

The following basic result on the curve complex was proved by Masur-Minsky [MM1]. (See also Bowditch [Bow] for an alternate proof.)

THEOREM 1.1. *For any surface S, the complex of curves is an infinite diameter δ-hyperbolic space* (*as long as it is nonempty*).

Given a subsurface $Y \subset S$, one can define a *subsurface projection* which is a map $\pi_{\mathcal{C}(Y)} \colon \mathcal{C}(S) \to 2^{\mathcal{C}(Y)}$. Suppose first that Y is not an annulus. Given any curve $\gamma \in \mathcal{C}(S)$ intersecting Y essentially, we define $\pi_{\mathcal{C}(Y)}(\gamma)$ to be the collection of vertices in $\mathcal{C}(Y)$ obtained by surgering the essential arcs of $\gamma \cap Y$ along ∂Y to obtain simple closed curves in Y. It is easy to show that $\pi_{\mathcal{C}(Y)}(\gamma)$ is nonempty and has uniformly bounded diameter. If Y is an annulus and γ intersects it transversely essentially, we may lift γ to an arc crossing the annulus $\tilde{S}'$ and let this be $\pi_{\mathcal{C}(Y)}(\gamma)$. If γ is a core curve of Y or fails to intersect it, we let $\pi_{\mathcal{C}(Y)}(\gamma) = \emptyset$ (this holds for general Y too).

When measuring distance in the image subsurface, we usually write $d_{\mathcal{C}(Y)}(\mu, \nu)$ as shorthand for $d_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(Y)}(\mu), \pi_{\mathcal{C}(Y)}(\nu))$.

Markings. The curve complex can be used to produce a geometric model for the mapping class group as done in [MM2]. This model is a graph called the *marking complex*, $\mathcal{M}(S)$, and is defined as follows.

We define vertices $\mu \in \mathcal{M}(S)$ to be pairs $(\mathrm{base}(\mu), \text{transversals})$ for which:

- The set of *base curves of* μ, denoted $\mathrm{base}(\mu)$, is a maximal simplex in $\mathcal{C}(S)$.
- The *transversals of* μ consist of one curve for each component of $\mathrm{base}(\mu)$, intersecting it transversely.

Further, the markings are required to satisfy the following two properties. First, for each $\gamma \in \text{base}(\mu)$, we require the transversal curve to γ, denoted t, to be disjoint from the rest of the base(μ). Second, given γ and its transversal t, we require that $\gamma \cup t$ fill a nonannular surface W satisfying $r(W) = 1$ and for which $d_{\mathcal{C}(W)}(\gamma, t) = 1$.

The edges of $\mathcal{M}(S)$ are of two types:

(1) *Twist*: Replace a transversal curve by another obtained by performing a Dehn twist along the associated base curve.

(2) *Flip*: Swap the roles of a base curve and its associated transversal curve. (After doing this move, the additional disjointness requirement on the transversals may not be satisfied. As shown in [MM2], one can surger the new transversal to obtain one that does satisfy the disjointness requirement. The additional condition that the new and old transversals intersect minimally restricts the surgeries to a finite number, and we obtain a finite set of possible flip moves for each marking. Each of these moves gives rise to an edge in the marking graph, and the naturality of the construction makes it invariant by the mapping class group.)

It is not hard to verify that $\mathcal{M}(S)$ is a locally finite graph on which the mapping class group acts cocompactly and properly discontinuously. As observed by Masur-Minsky [MM2], this yields:

LEMMA 1.2. *$\mathcal{M}(S)$ is quasi-isometric to the mapping class group of S.*

The same definitions apply to essential subsurfaces of S. For an annulus W, we let $\mathcal{M}(W)$ just be $\mathcal{C}(W)$.

Note that the above definition of marking makes no requirement that the surface S be connected. In the case of a disconnected surface $W = \sqcup_{i=1}^{n} W_i$, it is easy to see that $\mathcal{M}(W) = \prod_{i=1}^{n} \mathcal{M}(W_i)$.

Projections and distance. We now recall several ways in which subsurface projections arise in the study of mapping class groups.

First, note that for any $\mu \in \mathcal{M}(S)$ and any $Y \subseteq S$ the above projection maps extend to $\pi_{\mathcal{C}(Y)} \colon \mathcal{M}(S) \to 2^{\mathcal{C}(Y)}$. This map is simply the union over $\gamma \in \text{base}(\mu)$ of the usual projections $\pi_{\mathcal{C}(Y)}(\gamma)$, unless Y is an annulus about an element of base(μ). When Y is an annulus about $\gamma \in \text{base}(\mu)$, then we let $\pi_{\mathcal{C}(Y)}(\mu)$ be the projection of γ's transversal curve in μ. As in the case of curve complex projections, we write $d_{\mathcal{C}(Y)}(\mu, \nu)$ as shorthand for $d_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(Y)}(\mu), \pi_{\mathcal{C}(Y)}(\nu))$.

Remark 1.3. An easy, but useful, fact is that if a pair of markings $\mu, \nu \in \mathcal{M}(S)$ share a base curve γ and $\gamma \cap Y \neq \emptyset$, then there is a uniform bound on the diameter of $\pi_{\mathcal{C}(Y)}(\mu) \cup \pi_{\mathcal{C}(Y)}(\nu)$.

We say a pair of subsurfaces *overlap* if they intersect, and neither is nested in the other. The following is proven in [Be]:

THEOREM 1.4. *Let Y and Z be a pair of subsurfaces of S which overlap. There exists a constant M_1 depending only on the topological type of S, such that for any $\mu \in \mathcal{M}(S)$:*

$$\min\left\{d_{\mathcal{C}(Y)}(\partial Z, \mu), d_{\mathcal{C}(Z)}(\partial Y, \mu)\right\} \leq M_1.$$

Another application of the projection maps is the following distance formula of Masur-Minsky [MM2]:

THEOREM 1.5. *If $\mu, \nu \in \mathcal{M}(S)$, then there exists a constant $K(S)$, depending only on the topological type of S, such that for each $K > K(S)$ there exists $a \geq 1$ and $b \geq 0$ for which*:

$$d_{\mathcal{M}(S)}(\mu, \nu) \approx_{a,b} \sum_{Y \subseteq S} \left\{\!\!\left\{ d_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(Y)}(\mu), \pi_{\mathcal{C}(Y)}(\nu)) \right\}\!\!\right\}_K.$$

Here we define the expression $\{\!\{N\}\!\}_K$ to be N if $N > K$ and 0 otherwise — hence K functions as a "threshold" below which contributions are ignored.

Hierarchy paths. In fact, the distance formula of Theorem 1.5 is a consequence of a construction in [MM2] of a class of quasi-geodesics in $\mathcal{M}(S)$ which we call *hierarchy paths*, and which have the following properties.

Any two points $\mu, \nu \in \mathcal{M}(S)$ are connected by at least one hierarchy path γ. Each hierarchy path is a quasi-geodesic, with constants depending only on the topological type of S. The path γ "shadows" a $\mathcal{C}(S)$-geodesic β joining base(μ) to base(ν), in the following sense: There is a monotonic map $v\colon \gamma \to \beta$, such that $v(\gamma_n)$ is a vertex in base(γ_n) for every γ_n in γ.

(Note: the term "hierarchy" refers to a long combinatorial construction which yields these paths, and whose details we will not need to consider here.)

Furthermore the following criterion constrains the makeup of these paths. It asserts that subsurfaces of S which "separate" μ from ν in a significant way must play a role in the hierarchy paths from μ to ν:

LEMMA 1.6. *There exists a constant $M_2 = M_2(S)$ such that, if W is an essential subsurface of S and $d_{\mathcal{C}(W)}(\mu, \nu) > M_2$, then for any hierarchy path γ connecting μ to ν, there exists a marking γ_n in γ with $[\partial W] \subset \mathrm{base}(\gamma_n)$. Furthermore there exists a vertex v in the geodesic β shadowed by γ such that $W \subset S \setminus v$.*

This follows directly from Lemma 6.2 of [MM2].

Marking projections. We have already defined two types of subsurface projections; we end by mentioning one more which we shall use frequently.

Given a subsurface $Y \subset S$, we define a projection

$$\pi_{\mathcal{M}(Y)} \colon \mathcal{M}(S) \to \mathcal{M}(Y)$$

using the following procedure: If Y is an annulus $\mathcal{M}(Y) = \mathcal{C}(Y)$, we let $\pi_{\mathcal{M}(Y)} = \pi_{\mathcal{C}(Y)}$. For nonannular Y: given a marking μ we intersect its base curves with Y and choose a curve $\alpha \in \pi_Y(\mu)$. We repeat the construction with the subsurface $Y \setminus \alpha$, continuing until we have found a maximal simplex in $\mathcal{C}(Y)$. This will be the base of $\pi_{\mathcal{M}(Y)}(\mu)$. The transversal curves of the marking are obtained by projecting μ to each annular complex of a base curve, and then choosing a transversal curve which minimizes distance in the annular complex to this projection. (In case a base curve of μ already lies in Y, this curve will be part of the base of the image, and its transversal curve in μ will be used to determine the transversal for the image.)

This definition involved arbitrary choices, but it is shown in [Be] that the set of all possible choices form a uniformly bounded diameter subset of $\mathcal{M}(Y)$. Moreover, it is shown there that:

LEMMA 1.7. *$\pi_{\mathcal{M}(Y)}$ is coarsely Lipschitz with uniform constants.*

Similarly to the case of curve complex projections, we write $d_{\mathcal{M}(Y)}(\mu, \nu)$ as shorthand for $d_{\mathcal{M}(Y)}(\pi_{\mathcal{M}(Y)}(\mu), \pi_{\mathcal{M}(Y)}(\nu))$.

1.4. *Asymptotic cones.* The asymptotic cone of a metric space is roughly defined to be the limiting view of that space as seen from an arbitrarily large distance. This can be made precise using ultrafilters:

By a (*nonprincipal*) *ultrafilter* we mean a finitely additive probability measure ω defined on the power set of the natural numbers and taking values only 0 or 1, and for which every finite set has zero measure. The existence of nonprincipal ultrafilters depends in a fundamental way on the Axiom of Choice.

Given a sequence of points (x_n) in a topological space X, we say $x \in X$ is its *ultralimit*, or $x = \lim_\omega x_n$, if for every neighborhood U of x the set $\{n : x_n \in U\}$ has ω-measure equal to 1. We note that ultralimits are unique when they exist, and that when X is compact every sequence has an ultralimit.

The ultralimit of a sequence of based metric spaces $(X_n, x_n, \mathrm{dist}_n)$ is defined as follows: Using the notation $\boldsymbol{y} = (y_n \in X_n) \in \Pi_{n\in\mathbb{N}} X_n$ to denote a sequence, define $\mathrm{dist}(\boldsymbol{y}, \boldsymbol{z}) = \lim_\omega (y_n, z_n)$, where the ultralimit is taken in the compact set $[0, \infty]$. We then let

$$\lim_\omega (X_n, x_n, \mathrm{dist}_n) \equiv \{\boldsymbol{y} : \mathrm{dist}(\boldsymbol{y}, \boldsymbol{x}) < \infty\} / \sim,$$

where we define $\boldsymbol{y} \sim \boldsymbol{y}'$ if $\mathrm{dist}(\boldsymbol{y}, \boldsymbol{y}') = 0$. Clearly dist makes this quotient into a metric space.

Given a sequence of positive constants $s_n \to \infty$ and a sequence (x_n) of basepoints in a fixed metric space (X, dist), we may consider the rescaled space $(X, x_n, \mathrm{dist}/s_n)$. The ultralimit of this sequence is called the *asymptotic cone of* (X, dist) *relative to the ultrafilter* ω*, scaling constants* s_n*, and basepoint* $\boldsymbol{x} = (x_n)$:

$$\mathrm{Cone}_\omega(X, (x_n), (s_n)) = \lim_\omega (X, x_n, \frac{\mathrm{dist}}{s_n}).$$

(For further details see [dDW], [Gro1].)

For the remainder of the paper, let us fix a nonprincipal ultrafilter ω, a sequence of scaling constants $s_n \to \infty$, and a basepoint μ_0 for $\mathcal{M}(S)$. We write $\mathcal{M}_\omega = \mathcal{M}_\omega(S)$ to denote an asymptotic cone of $\mathcal{M}(S)$ with respect to these choices. Note that since $\mathcal{M}$ is quasi-isometric to a word metric on $\mathcal{MCG}$, the space $\mathcal{M}_\omega$ is homogeneous and thus the asymptotic cone is independent of the choice of basepoint. Further, since on a given group any two finitely generated word metrics are quasi-isometric, fixing an ultrafilter and scaling constants we have that different finitely generated word metrics on $\mathcal{MCG}$ have bi-Lipschitz homeomorphic asymptotic cones. Also, we note that in general the asymptotic cone of a geodesic space is a geodesic space. Thus, $\mathcal{M}_\omega$ is a geodesic space, and in particular is locally path connected.

Any essential connected subsurface W inherits a basepoint $\pi_{\mathcal{M}(W)}(\mu_0)$, canonical up to bounded error by Lemma 1.7, and we can use this to define its asymptotic cone $\mathcal{M}_\omega(W)$. For a disconnected subsurface $W = \sqcup_{i=1}^k W_i$ we have $\mathcal{M}(W) = \Pi_{i=1}^k \mathcal{M}(W_i)$ and we may similarly construct $\mathcal{M}_\omega(W)$ which can be identified with $\Pi_{i=1}^k \mathcal{M}_\omega(W_i)$ (this follows from the general fact that the process of taking asymptotic cones commutes with finite products). Note that for an annulus A we've defined $\mathcal{M}(A) = \mathcal{C}(A)$ which is quasi-isometric to $\mathbb{Z}$, so that $\mathcal{M}_\omega(A)$ is $\mathbb{R}$.

It will be crucial to generalize this to *sequences* of subsurfaces in S. Let us note first the general fact that any sequence in a finite set A is ω-a.e. constant. That is, given $(a_n \in A)$ there is a unique $a \in A$ such that $\omega(\{n : a_n = a\}) = 1$. Hence for example if $\boldsymbol{W} = (W_n)$ is a sequence of essential subsurfaces of S then the topological type of W_n is ω-a.e. constant and we call this the topological type of $\boldsymbol{W}$. Similarly the topological type of the pair (S, W_n) is ω-a.e. constant. We can moreover interpret expressions like $\boldsymbol{U} \subset \boldsymbol{W}$ for sequences $\boldsymbol{U}$ and $\boldsymbol{W}$ of subsurfaces to mean $U_n \subset W_n$ for ω-a.e. n, and so on. We say that two sequences (α_n), (α'_n) are equivalent mod ω if $\alpha_n = \alpha'_n$ for ω-a.e. n, and note that topological type, containment, etc. are invariant under this equivalence relation. Throughout, we adopt the convention of using boldface to denote sequences. We will always consider such sequences mod ω, unless they are sequences of markings $\boldsymbol{\mu} \in \mathcal{M}_\omega$, in which case they are considered modulo the weaker equivalence $\sim$ from the definition of asymptotic cones.

If $\boldsymbol{W} = (W_n)$ is a sequence of subsurfaces, we let $\mathcal{M}_\omega(\boldsymbol{W})$ denote the ultralimit of $\mathcal{M}(W_n)$ with metrics rescaled by $\frac{1}{s_n}$ and with basepoints $\pi_{\mathcal{M}(W_n)}(\mu_0)$. Note that $\mathcal{M}_\omega(\boldsymbol{W})$ can be identified with $\mathcal{M}_\omega(W)$, where W is a surface homeomorphic to W_n for ω-a.e. n.

2. Product regions

In this section we will describe the geometry of the set of markings containing a prescribed set of base curves. Equivalently, in the mapping class group such a set corresponds to the coset of the stabilizer of a simplex in the complex of curves. Not surprisingly, these regions coarsely decompose as products.

Let Δ be a simplex in the complex of curves, i.e., a multicurve in S. We may partition S into subsurfaces isotopic to complementary components of Δ, and annuli whose cores are elements of Δ. After throwing away components homeomorphic to $S_{0,3}$ we obtain what we call the "partition" of Δ, and denote it by $\sigma(\Delta)$.

Let $\mathcal{Q}(\Delta) \subset \mathcal{M}(S)$ denote the set of markings whose bases contain Δ. There is a natural (coarse) identification

$$\mathcal{Q}(\Delta) \approx \prod_{U \in \sigma(\Delta)} \mathcal{M}(U) \tag{2.1}$$

where if U is an annulus we take $\mathcal{M}(U)$ to mean the annulus complex of U. This identification is obtained simply by restriction (or equivalently by subsurface projection) for each nonannulus component, and by associating transversals with points in annulus complexes for the annular components.

Theorem 1.5 yields the following basic lemmas. When A is a subsurface and B is a collection of curves, we write $A \pitchfork B \neq \emptyset$ to mean that B cannot be deformed away from A.

LEMMA 2.1. *The identification* (2.1) *is a quasi-isometry with uniform constants.*

LEMMA 2.2. *If* $\mu \in \mathcal{M}(S)$ *then*

$$d(\mu, \mathcal{Q}(\Delta)) \approx \sum_{W \pitchfork \Delta \neq \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(W)}(\mu, \Delta) \right\}\!\!\right\}_K .$$

Proof of Lemma 2.1. If $\mu, \nu \in \mathcal{Q}(\Delta)$, the distance formula in Theorem 1.5 gives

$$d(\mu, \nu) \approx \sum_{W} \left\{\!\!\left\{ d_{\mathcal{C}(W)}(\mu, \nu) \right\}\!\!\right\}_K$$

where the constants in $\approx$ depend on the threshold K. Now if $W \pitchfork \Delta \neq \emptyset$, then Remark 1.3 implies that $\pi_W(\mu)$ and $\pi_W(\nu)$ are each a bounded distance from

$\pi_W(\Delta)$, and hence the W term in the sum is bounded by twice this. Raising K above this constant means that all such terms vanish and the sum is only over surfaces W disjoint from Δ, or annuli whose cores are components of Δ. But this is estimated by the distance in $\prod_{U\in\sigma(\Delta)}\mathcal{M}(U)$, when we use Theorem 1.5 in each U separately. $\square$

Proof of Lemma 2.2. Let $\mu\in\mathcal{M}(S)$. For any $\nu\in\mathcal{Q}(\Delta)$, we note that, if $W \pitchfork \Delta\neq\emptyset$, then

$$|d_{\mathcal{C}(W)}(\mu,\nu)-d_{\mathcal{C}(W)}(\mu,\Delta)|\le c$$

for some constant c, by Remark 1.3. If K_0 is the minimal threshold that can be used in the distance formula of Theorem 1.5, let $K=K_0+2c$. We then see that for any W contributing to the sum

$$\sum_{W\pitchfork\Delta\neq\emptyset}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\Delta)\right\}\!\!\right\}_K$$

we must have

$$d_{\mathcal{C}(W)}(\mu,\nu)\ge d_{\mathcal{C}(W)}(\mu,\Delta)-c>K_0$$

and, since our choice of K yields $\frac{1}{2}d_{\mathcal{C}(W)}(\mu,\Delta)>c$, we furthermore have

$$d_{\mathcal{C}(W)}(\mu,\nu)\ge \tfrac{1}{2}d_{\mathcal{C}(W)}(\mu,\Delta).$$

It follows then that

$$\begin{aligned}\sum_{W}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\nu)\right\}\!\!\right\}_{K_0} &\ge \sum_{W\pitchfork\Delta\neq\emptyset}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\nu)\right\}\!\!\right\}_{K_0}\\ &\ge \tfrac{1}{2}\sum_{W\pitchfork\Delta\neq\emptyset}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\Delta)\right\}\!\!\right\}_{K}.\end{aligned}$$

This gives one direction of the desired inequality.

To obtain the other direction, we fix $\mu\in\mathcal{M}(S)$ and let $\nu\in\mathcal{Q}(\Delta)$ be the marking whose restriction to each $U\in\sigma(\Delta)$ is just $\pi_{\mathcal{M}(U)}(\mu)$. With this choice,

$$d_{\mathcal{C}(W)}(\mu,\nu)\le c$$

for a uniform constant c whenever $W\pitchfork\Delta=\emptyset$, since the intersections of μ and ν with W are essentially the same. Setting our threshold $K\ge K_0+2c$ again we see that these terms all vanish, and

$$\begin{aligned}\sum_{W}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\nu)\right\}\!\!\right\}_{K} &= \sum_{W\pitchfork\Delta\neq\emptyset}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\nu)\right\}\!\!\right\}_{K}\\ &\le 2\sum_{W\pitchfork\Delta\neq\emptyset}\left\{\!\!\left\{d_{\mathcal{C}(W)}(\mu,\Delta)\right\}\!\!\right\}_{K_0}\end{aligned}$$

where the last inequality is obtained using the same threshold trick as above (we can assume it is the same value of c). $\square$

Product regions in the asymptotic cone. Consider a sequence $\boldsymbol{\Delta} = \{\Delta_n\}$ such that $\lim_\omega \frac{1}{s_n} d(\mu_0, \mathcal{Q}(\Delta_n)) < \infty$. We can take the ultralimit of $\mathcal{Q}(\Delta_n)$, with metrics rescaled by $1/s_n$, obtaining a subset of $\mathcal{M}_\omega(S)$ which we denote $\mathcal{Q}_\omega(\boldsymbol{\Delta})$. Lemma 2.1 and the fact that ultralimits commute with finite products implies that there is a bi-Lipschitz identification

$$\mathcal{Q}_\omega(\boldsymbol{\Delta}) \cong \prod_{\boldsymbol{U} \in \sigma(\boldsymbol{\Delta})} \mathcal{M}_\omega(\boldsymbol{U}). \tag{2.2}$$

Here $\sigma(\boldsymbol{\Delta})$ is defined as follows: As in Section 1.4, the topological type of $\sigma(\Delta_n)$ is ω-a.e. constant, and so there is a set $J \subset \mathbb{N}$ with $\omega(J) = 1$, a partition $\sigma' = \{U^1, \dots, U^k\}$ of S, and a sequence of homeomorphisms $f_n \colon S \to S$ taking σ' to $\sigma(\Delta_n)$ for each $n \in J$. We then let $\sigma(\boldsymbol{\Delta}) = \{\boldsymbol{U}^1, \dots, \boldsymbol{U}^k\}$ where $\boldsymbol{U}^i = (f_n(U^i))$ for $n \in J$ (it doesn't matter, mod ω, how we define it for $n \notin J$). Any nonuniqueness of f_n, up to isotopy, corresponds to a symmetry of σ', and hence to a permutation of the indices of elements of $\sigma(\boldsymbol{\Delta})$.

Moreover, Lemma 2.2 implies that distance to $\mathcal{Q}_\omega(\boldsymbol{\Delta})$ can be estimated, up to bounded ratio, by:

$$\rho(\boldsymbol{\mu}, \Delta) \equiv \lim_\omega \frac{1}{s_n} \sum_{W \pitchfork \Delta_n \neq \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(W)}(\mu_n, \Delta_n) \right\}\!\!\right\}_K . \tag{2.3}$$

3. Separating product regions and locally constant maps

In this section we will define the family of product regions equipped with locally constant maps (denoted as $\mathcal{P}$ in the outline in the introduction). Each region will be determined by a sequence $\boldsymbol{W} = (W_n)$ of connected subsurfaces of S, and a choice $\boldsymbol{x} = (x_n)$ of basepoint in $\mathcal{M}_\omega(\boldsymbol{W})$. Theorem 3.5, which defines the projection map associated to each region and establishes its properties, is the main result of this section.

3.1. *Sublinear growth sets.* In Behrstock [Be], a family of subsets of $\mathcal{M}_\omega(S)$ is introduced, and defined as follows: for $\boldsymbol{x} \in \mathcal{M}_\omega(S)$, let

$$F(\boldsymbol{x}) = \left\{ \boldsymbol{y} : \lim_\omega \frac{1}{s_n} \sup_{U \subsetneq S} d_{\mathcal{M}(U)}(x_n, y_n) = 0 \right\}.$$

That is, the distance between x_n and y_n, projected to the marking graph of any proper subsurface, is vanishingly small compared to their distance in $\mathcal{M}(S)$. We note that, because the subsurface projections are uniformly Lipschitz, this condition is well-defined, i.e., does not depend on the choice of y_n representing $\boldsymbol{y}$.

Behrstock proved that $F(\boldsymbol{x})$ is an $\mathbb{R}$-tree, and more strongly that for any two points in $F(\boldsymbol{x})$ there is a unique embedded arc in $\mathcal{M}_\omega(S)$ connecting them. We can generalize this construction slightly as follows:

First, for a sequence $\boldsymbol{U} = (U_n)$ of connected subsurfaces and $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{M}_\omega(S)$ we have

$$d_{\mathcal{M}_\omega(\boldsymbol{U})}(\boldsymbol{x}, \boldsymbol{y}) = \lim_\omega \frac{1}{s_n} d_{\mathcal{M}(U_n)}(x_n, y_n).$$

Now if $\boldsymbol{W} = (W_n)$ is a sequence of connected subsurfaces (considered mod ω) and $\boldsymbol{x} \in \mathcal{M}_\omega(\boldsymbol{W})$, we define $F_{\boldsymbol{W},\boldsymbol{x}} \subset \mathcal{M}_\omega(\boldsymbol{W})$ to be:

$$F_{\boldsymbol{W},\boldsymbol{x}} = \{\boldsymbol{y} \in \mathcal{M}_\omega(\boldsymbol{W}) : d_{\mathcal{M}_\omega(\boldsymbol{U})}(\boldsymbol{x}, \boldsymbol{y}) = 0 \text{ for all } \boldsymbol{U} \subsetneq \boldsymbol{W}\}.$$

If $W_n \equiv S$, this is equivalent to the definition of $F(\boldsymbol{x})$ above. Note also that if $\boldsymbol{W} = \mathbf{collar}(\boldsymbol{\alpha})$ then $F_{\boldsymbol{W},\boldsymbol{x}}$ is just the asymptotic cone of the annulus complex of $\boldsymbol{W}$, which is a copy of $\mathbb{R}$.

Let us restate and discuss Behrstock's theorem from [Be]:

THEOREM 3.1. *Let $\boldsymbol{W} = (W_n)$ be a sequence of connected subsurfaces of S, and $\boldsymbol{x} \in \mathcal{M}_\omega(\boldsymbol{W})$. Any two points $\boldsymbol{y}, \boldsymbol{z} \in F_{\boldsymbol{W},\boldsymbol{x}}$ are connected by a unique embedded path in $\mathcal{M}_\omega(\boldsymbol{W})$, and this path lies in $F_{\boldsymbol{W},\boldsymbol{x}}$.*

In particular, it follows that $F_{\boldsymbol{W},\boldsymbol{x}}$ is an $\mathbb{R}$-tree. Here is a brief outline of the proof: The annular case is trivial because $F_{\boldsymbol{W},\boldsymbol{x}} = \mathcal{M}_\omega(\boldsymbol{W}) \cong \mathbb{R}$. Hence, we assume W_n are not annuli for ω-a.e. n. In each W_n, connect y_n to z_n with a hierarchy path γ_n (see §1.3). Since γ_n are uniform quasi-geodesics, after rescaling, their ultralimit gives a path $\boldsymbol{\gamma}$ in $\mathcal{M}_\omega(\boldsymbol{W})$. Using the tools of [MM2] together with the assumption that $\boldsymbol{y}, \boldsymbol{z} \in F_{\boldsymbol{W},\boldsymbol{x}}$, one can show that $\boldsymbol{\gamma}$ lies in $F_{\boldsymbol{W},\boldsymbol{x}}$.

Let β_n be a $\mathcal{C}(W_n)$-geodesic shadowed by γ_n. One can see that the length $|\beta_n| \to_\omega \infty$ as follows: Suppose instead that $|\beta_n| < L$ for ω-a.e. n. Choose the threshold in the distance formula large enough so that the nonzero terms in

$$\sum_{V \subset W_n} \left\{\!\!\left\{ d_{\mathcal{C}(V)}(y_n, z_n) \right\}\!\!\right\}_K$$

are proper subsurfaces in W_n which play the role in γ_n determined by Lemma 1.6 — that is, each one is disjoint from some $v \in \beta_n$. But since β_n has at most L vertices, there must be one, v_n, which is disjoint from enough surfaces to contribute at least $1/L$ times the sum. But this means, by the distance formula within $Y_n = S \setminus v_n$, that $d_{\mathcal{M}_\omega(\boldsymbol{Y})}(\boldsymbol{y}, \boldsymbol{z}) > 0$, which contradicts the assumption that $\boldsymbol{y}, \boldsymbol{z} \in F_{\boldsymbol{W},\boldsymbol{x}}$.

Consider the map $p_n \colon \mathcal{M}(W_n) \to \beta_n$ which takes a marking μ to a vertex $v \in \beta_n$ of minimal $\mathcal{C}(W_n)$-distance to the base of μ. We promote p_n to a map $q_n \colon \mathcal{M}(W_n) \to \gamma_n$ by letting $q_n(\mu)$ be a marking of γ_n which shadows $v = p_n(\mu)$.

The ultralimit of q_n yields a map $\boldsymbol{q} \colon \mathcal{M}_\omega(\boldsymbol{W}) \to \boldsymbol{\gamma} \subset F_{\boldsymbol{W},\boldsymbol{x}}$. Furthermore one can show using hyperbolicity of $\mathcal{C}(W_n)$ (Masur-Minsky [MM1]) and properties of the subsurface projection maps that q_n has coarse contraction properties that, in the limit, imply that $\boldsymbol{q}$ is locally constant in the complement of $\boldsymbol{\gamma}$. It then easily follows that $\boldsymbol{y}$ and $\boldsymbol{z}$ cannot be connected in the complement of any point of $\boldsymbol{\gamma}$, and hence any path between them must contain $\boldsymbol{\gamma}$, and any embedded path must equal $\boldsymbol{\gamma}$.

3.2. *Definition of* $P_{\boldsymbol{W},\boldsymbol{x}}$. Given $\boldsymbol{W}$ and $\boldsymbol{x}$ as above, our separating product regions, denoted $P_{\boldsymbol{W},\boldsymbol{x}}$, will be subsets of $\mathcal{Q}_\omega(\partial\boldsymbol{W})$ defined as follows:

In the product structure (2.2) for $\mathcal{Q}_\omega(\partial\boldsymbol{W})$, $\boldsymbol{W}$ is a member of $\sigma(\partial\boldsymbol{W})$, and hence $\mathcal{M}_\omega(\boldsymbol{W})$ appears as a factor. We let $P_{\boldsymbol{W},\boldsymbol{x}}$ be the subset of $\mathcal{Q}_\omega(\partial\boldsymbol{W})$ consisting of points whose coordinate in the $\mathcal{M}_\omega(\boldsymbol{W})$ factor lies in $F_{\boldsymbol{W},\boldsymbol{x}}$.

Since the identification of $\mathcal{Q}_\omega(\partial\boldsymbol{W})$ with the product structure is made using the subsurface projections, we have this characterization:

LEMMA 3.2. *$P_{\boldsymbol{W},\boldsymbol{x}}$ is the set of points $\boldsymbol{y} \in \mathcal{M}_\omega(S)$ such that*:

(1) $\pi_{\mathcal{M}_\omega\boldsymbol{W}}(\boldsymbol{y}) \in F_{\boldsymbol{W},\boldsymbol{x}}$, *and*

(2) $\rho(\boldsymbol{y}, \partial\boldsymbol{W}) = 0$.

Here $\rho(\boldsymbol{y}, \partial\boldsymbol{W})$ is an estimate for the distance of $\boldsymbol{y}$ from $\mathcal{Q}_\omega(\partial\boldsymbol{W})$, as defined in (2.3). Also, the ultralimit of the rescaled marking projection maps $\mathcal{M}(S) \to \mathcal{M}(W_n)$ is denoted by:

$$\pi_{\mathcal{M}_\omega\boldsymbol{W}} \colon \mathcal{M}_\omega(S) \to \mathcal{M}_\omega(\boldsymbol{W}).$$

Define W_n^c to be the union of the components of $\sigma(\partial W_n)$ not equal to W_n (so W_n^c includes annuli around ∂W_n, unless W_n itself is an annulus). Let $\boldsymbol{W}^c = (W_n^c)$. Then $\mathcal{M}_\omega(\boldsymbol{W}^c)$ is the asymptotic cone of $(\mathcal{M}(W_n^c))$, and can be identified with the product of the remaining factors in $\mathcal{Q}_\omega(\partial\boldsymbol{W})$:

$$\mathcal{M}_\omega(\boldsymbol{W}^c) \equiv \prod_{\substack{\boldsymbol{U}\in\sigma(\partial\boldsymbol{W})\\ U\neq\boldsymbol{W}}} \mathcal{M}_\omega(\boldsymbol{U}).$$

We can summarize this in the following:

LEMMA 3.3. *There exists a bi-Lipschitz identification of $P_{\boldsymbol{W},\boldsymbol{x}}$ with*

$$F_{\boldsymbol{W},\boldsymbol{x}} \times \mathcal{M}_\omega(\boldsymbol{W}^c).$$

3.3. *Projection maps.* The following projection theorem is a small improvement on Theorem 3.1 from Behrstock [Be].

THEOREM 3.4. *Given $\boldsymbol{x} \in \mathcal{M}_\omega(\boldsymbol{W})$, there is a continuous map*

$$\wp = \wp_{\boldsymbol{W},\boldsymbol{x}} \colon \mathcal{M}_\omega(\boldsymbol{W}) \to F_{\boldsymbol{W},\boldsymbol{x}}$$

with these properties:

(1) $\wp$ *is the identity on* $F_{\boldsymbol{W},\boldsymbol{x}}$.

(2) $\wp$ *is locally constant in* $\mathcal{M}_\omega(\boldsymbol{W}) \setminus F_{\boldsymbol{W},\boldsymbol{x}}$.

Note that in the proof of Theorem 3.1 a projection to individual paths was shown to have locally constant properties. In this theorem we construct a projection from $\mathcal{M}_\omega(\boldsymbol{W})$ onto $F_{\boldsymbol{W},\boldsymbol{x}}$.

Proof. For any $\boldsymbol{y} \in \mathcal{M}_\omega(\boldsymbol{W})$ let α be a path connecting $\boldsymbol{y}$ to any point in $F_{\boldsymbol{W},\boldsymbol{x}}$. Let α_1 be the first point in α that is in $F_{\boldsymbol{W},\boldsymbol{x}}$. We claim that α_1 depends only on $\boldsymbol{y}$. For otherwise let β be another path with $\beta_1 \neq \alpha_1$. Then segments of α and β form a path connecting two points of $F_{\boldsymbol{W},\boldsymbol{x}}$ outside of $F_{\boldsymbol{W},\boldsymbol{x}}$ — this contradicts Theorem 3.1.

We can then define $\wp(\boldsymbol{y}) \equiv \alpha_1$. This is locally constant at $\boldsymbol{y} \notin F_{\boldsymbol{W},\boldsymbol{x}}$ because for a sufficiently small neighborhood U of $\boldsymbol{y}$, every $\boldsymbol{z} \in U$ can be connected to $F_{\boldsymbol{W},\boldsymbol{x}}$ by a path going first through $\boldsymbol{y}$ (since $\mathcal{M}_\omega(\boldsymbol{W})$ is locally path-connected).

Continuity of $\wp$ at points of $F_{\boldsymbol{W},\boldsymbol{x}}$ follows immediately from the definition of $\wp$ and the fact that $\mathcal{M}_\omega(\boldsymbol{W})$ is a locally path connected geodesic space. $\square$

We can now construct our global projection map for $F_{\boldsymbol{W},\boldsymbol{x}}$:

THEOREM 3.5. *Given* $\boldsymbol{x} \in \mathcal{M}_\omega(\boldsymbol{W})$, *there is a continuous map*

$$\Phi = \Phi_{\boldsymbol{W},\boldsymbol{x}} \colon \mathcal{M}_\omega(S) \to F_{\boldsymbol{W},\boldsymbol{x}}$$

with these properties:

(1) Φ *restricted to* $P_{\boldsymbol{W},\boldsymbol{x}}$ *is projection to the first factor in the product structure* $P_{\boldsymbol{W},\boldsymbol{x}} \cong F_{\boldsymbol{W},\boldsymbol{x}} \times \mathcal{M}_\omega(\boldsymbol{W}^c)$.

(2) Φ *is locally constant in the complement of* $P_{\boldsymbol{W},\boldsymbol{x}}$.

Proof. We define the map simply by

$$\Phi_{\boldsymbol{W},\boldsymbol{x}} = \wp_{\boldsymbol{W},\boldsymbol{x}} \circ \pi_{\mathcal{M}_\omega \boldsymbol{W}}.$$

Property (1) follows from the definition, and from the way that the identification of $P_{\boldsymbol{W},\boldsymbol{x}}$ with the product in Lemma 3.3 is constructed via subsurface projections.

We divide the proof of property (2) into two cases:

Case 1. $\pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{y}) \notin F_{\boldsymbol{W},\boldsymbol{x}}$. In this case the desired fact follows immediately from the locally constant property of $\wp$ shown in Theorem 3.4, and the continuity of $\pi_{\mathcal{M}_\omega \boldsymbol{W}}$.

Case 2. $\pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{y}) \in F_{\boldsymbol{W},\boldsymbol{x}}$. Since $\boldsymbol{y} \notin P_{\boldsymbol{W},\boldsymbol{x}}$ and $\pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{y}) \in F_{\boldsymbol{W},\boldsymbol{x}}$, Lemma 3.2 implies that $\rho(\boldsymbol{y}, \partial \boldsymbol{W}) > 0$.

Let $\boldsymbol{z} \in \mathcal{M}_\omega(S)$, with $\Phi(\boldsymbol{z}) \neq \Phi(\boldsymbol{y})$. We will derive a lower bound for $d(\boldsymbol{y}, \boldsymbol{z})$, and this will prove the theorem.

Let $\boldsymbol{z}' = \pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{z})$ and $\boldsymbol{y}' = \pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{y})$. Since Case 1 has already been handled, we may assume $\boldsymbol{y}' \in F_{\boldsymbol{W},\boldsymbol{x}}$, so that $\boldsymbol{y}' = \wp(\boldsymbol{y}') = \Phi(\boldsymbol{y})$. As in Theorem 3.4, any path from $\boldsymbol{z}'$ to $\boldsymbol{y}'$ must pass through $\wp(\boldsymbol{z}')$ first. Note that $\wp(\boldsymbol{z}') = \Phi(\boldsymbol{z}) \neq \boldsymbol{y}'$. Now let γ_n be hierarchy paths in $\mathcal{M}(W_n)$ connecting z_n' to y_n'. Since γ_n are quasigeodesics, their ultralimit after rescaling gives rise to a path in $\mathcal{M}_\omega(\boldsymbol{W})$ connecting $\boldsymbol{z}'$ to $\boldsymbol{y}'$ and hence there must exist $\delta_n \in \gamma_n$ such that (δ_n) represents $\wp(\boldsymbol{z}')$. As remarked in the outline of the proof of Theorem 3.1, $d_{\mathcal{C}(W_n)}(\delta_n, y_n') \to_\omega \infty$ since $\wp(\boldsymbol{z}')$ and $\boldsymbol{y}'$ are distinct points in $F_{\boldsymbol{W},\boldsymbol{x}}$. Now since γ_n monotonically shadows a $\mathcal{C}(W_n)$ geodesic from z_n' to y_n', we conclude that

$$d_{\mathcal{C}(W_n)}(y_n', z_n') \to_\omega \infty.$$

Since $\pi_{\mathcal{C}(W_n)} \circ \pi_{\mathcal{M}(W_n)}$ and $\pi_{\mathcal{C}(W_n)}$ differ by a bounded constant (immediate from the definitions), we conclude that

$$d_{\mathcal{C}(W_n)}(y_n, z_n) \to_\omega \infty.$$

Now by the definition of $\rho(\boldsymbol{y}, \partial \boldsymbol{W})$, we know that

$$\frac{1}{s_n} \sum_{U \pitchfork \partial W_n \neq \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(U)}(y_n, \partial W_n) \right\}\!\!\right\}_K \to_\omega c > 0. \tag{3.1}$$

Let U be a subsurface participating in this sum for some n, so that $d_{\mathcal{C}(U)}(y_n, \partial W_n) > K$. We want to show that

$$d_{\mathcal{C}(U)}(y_n, z_n) \geq d_{\mathcal{C}(U)}(y_n, \partial W_n) - K' \tag{3.2}$$

for some K'.

We assume that K is larger than the constant M_1 from Theorem 1.4, and recall that this theorem states that

$$\min\{d_{\mathcal{C}(V)}(\mu, \partial V'), d_{\mathcal{C}(V')}(\mu, \partial V)\} \leq M_1 \tag{3.3}$$

for any marking μ and subsurfaces V, V' with $\partial V \pitchfork \partial V' \neq \emptyset$.

Since U meets ∂W_n, we have either $\partial U \pitchfork W_n \neq \emptyset$, in which case the subsurfaces W_n and U overlap, or $W_n \subsetneq U$.

Suppose first that $\partial U \pitchfork W_n \neq \emptyset$. Now we have $d_{\mathcal{C}(U)}(y_n, \partial W_n) > K > M_1$, since W_n and U overlap and (3.3) implies

$$d_{\mathcal{C}(W_n)}(y_n, \partial U) \leq M_1.$$

Now by the triangle inequality

$$d_{\mathcal{C}(W_n)}(\partial U, z_n) \geq d_{\mathcal{C}(W_n)}(y_n, z_n) - M_1 - D$$

(where D is a bound for $\operatorname{diam}_{\mathcal{C}(W_n)}(\mu)$ of any marking, as given by Remark 1.3). Since $d_{\mathcal{C}(W_n)}(y_n, z_n) \to_\omega \infty$, we may assume that this gives

$$d_{\mathcal{C}(W_n)}(\partial U, z_n) > M_1.$$

Now again by (3.3) we have

$$d_{\mathcal{C}(U)}(\partial W_n, z_n) \le M_1$$

and again by the triangle inequality

$$d_{\mathcal{C}(U)}(y_n, z_n) \ge d_{\mathcal{C}(U)}(y_n, \partial W_n) - M_1 - D$$

which establishes (3.2) when $\partial U \pitchfork W_n \ne \emptyset$.

Next, let us establish (3.2) when $W_n \subsetneq U$. Since $d_{\mathcal{C}(W_n)}(y_n, z_n) \to_\omega \infty$, we may assume that this distance is larger than the constant M_2 in Lemma 1.6. Let γ_n be a hierarchy path in $\mathcal{M}(U)$ connecting $\pi_{\mathcal{M}U}(y_n)$ to $\pi_{\mathcal{M}U}(z_n)$, and let β_n be the $\mathcal{C}(U)$-geodesic from $\pi_{\mathcal{C}(U)}(y_n)$ to $\pi_{\mathcal{C}(U)}(z_n)$ that γ_n shadows. Lemma 1.6 implies that ∂W_n appears in the base of at least one marking in γ_n, and hence $[\partial W_n]$ is $\mathcal{C}(U)$-distance at most one from a vertex of β_n. This means that the length of β_n is at least $d_{\mathcal{C}(U)}(\partial W_n, y_n) - 2$, in particular:

$$d_{\mathcal{C}(U)}(z_n, y_n) \ge d_{\mathcal{C}(U)}(y_n, \partial W_n) - 2.$$

Thus, we have established (3.2) with $K' = \max\{M_1 + D, 2\}$.

Now applying this to all the terms in the sum of (3.1), we would like to obtain a lower bound (for ω-a.e. n)

$$\frac{1}{s_n} \sum_{U \pitchfork \partial W_n \ne \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(U)}(y_n, z_n) \right\}\!\!\right\}_K > c' > 0. \tag{3.4}$$

To do this we apply the same threshold trick we used in the proof of Lemma 2.2. Since Theorem 1.5 applies to any sufficiently large threshold, we may choose $K'' = 2K' + K$ to replace the threshold K in the sum in (3.1), and obtain

$$\frac{1}{s_n} \sum_{U \pitchfork \partial W_n \ne \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(U)}(y_n, \partial W_n) \right\}\!\!\right\}_{K''} \to_\omega c' > 0. \tag{3.5}$$

Now, for a given, n if U contributes to this sum then by (3.2), we have $d_{\mathcal{C}(U)}(y_n, z_n) \ge K'' - K' > K$, and moreover

$$d_{\mathcal{C}(U)}(y_n, z_n) \ge d_{\mathcal{C}(U)}(y_n, \partial W_n) - K' \ge \tfrac{1}{2} d_{\mathcal{C}(U)}(y_n, \partial W_n).$$

This implies that

$$\sum_{U \pitchfork \partial W_n \ne \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(U)}(y_n, z_n) \right\}\!\!\right\}_K \ge \tfrac{1}{2} \sum_{U \pitchfork \partial W_n \ne \emptyset} \left\{\!\!\left\{ d_{\mathcal{C}(U)}(y_n, \partial W_n) \right\}\!\!\right\}_{K''}.$$

In other words, again by the distance formula, this gives us a lower bound of the form

$$d_{\mathcal{M}_\omega(S)}(\boldsymbol{y}, \boldsymbol{z}) > c'' > 0.$$

The conclusion is that if $d(\boldsymbol{y}, \boldsymbol{z}) < c''$ then $\Phi(\boldsymbol{y}) = \Phi(\boldsymbol{z})$, which is what we wanted. ∎

3.4. *Separators.* In [Be], it was shown that mapping class groups have global cut-points in their asymptotic cones; cf. Theorem 3.1. Since mapping class groups are not δ-hyperbolic, except in a few low complexity cases, it clearly cannot hold that arbitrary pairs of points in the asymptotic cone are separated by a point. Instead we identify here a larger class of subsets which do separate points:

THEOREM 3.6. *There is a family $\mathcal{L}$ of closed subsets of $\mathcal{M}_\omega(S)$ such that any two points in $\mathcal{M}_\omega(S)$ are separated by some $L \in \mathcal{L}$. Moreover each $L \in \mathcal{L}$ is isometric to $\mathcal{M}_\omega(Z)$, where Z is some proper essential* (*not necessarily connected*) *subsurface of S, with $r(Z) < r(S)$.*

We will see as part of an inductive argument in the next section that these separators L all have (locally compact) dimension at most $r(S)-1$; this bound is sharp since $\mathcal{M}_\omega$ contains $r(S)$-dimensional bi-Lipschitz flats which, of course, can not be separated by any subset of dimension less than $r(S)-1$.

Proof. Fix $\boldsymbol{x} \neq \boldsymbol{y} \in \mathcal{M}_\omega(S)$. We claim that there exists a subsurface sequence $\boldsymbol{W} = (W_n)$ such that:

(1) $d_{\mathcal{M}_\omega(\boldsymbol{W})}(\boldsymbol{x}, \boldsymbol{y}) > 0$, and

(2) For any $\boldsymbol{Y} = (Y_n)$ with $\boldsymbol{Y} \subsetneq \boldsymbol{W}$, $d_{\mathcal{M}_\omega(\boldsymbol{Y})}(\boldsymbol{x}, \boldsymbol{y}) = 0$.

Indeed, $\boldsymbol{W} = (S)$ satisfies the first condition. If it fails the second, we may choose $\boldsymbol{W'} \subsetneq \boldsymbol{W}$ with $d_{\mathcal{M}_\omega(\boldsymbol{W'})}(\boldsymbol{x}, \boldsymbol{y}) > 0$, and continue. This terminates since the complexity of the subsurface sequence decreases.

Let $\boldsymbol{x}' = \pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{x})$ and $\boldsymbol{y}' = \pi_{\mathcal{M}_\omega \boldsymbol{W}}(\boldsymbol{y})$. The choice of $\boldsymbol{W}$ implies that $\boldsymbol{x}' \neq \boldsymbol{y}'$ and that $\boldsymbol{y}' \in F_{\boldsymbol{W},\boldsymbol{x}'}$. (Note that the second condition implies $F_{\boldsymbol{W},\boldsymbol{x}'} = F_{\boldsymbol{W},\boldsymbol{y}'}$.) Let $\boldsymbol{z}$ be a point in $F_{\boldsymbol{W},\boldsymbol{x}'}$ in the interior of the path from $\boldsymbol{x}'$ to $\boldsymbol{y}'$. Since $F_{\boldsymbol{W},\boldsymbol{x}'}$ is an $\mathbb{R}$-tree (by Theorem 3.1), $\boldsymbol{z}$ separates $\boldsymbol{x}'$ from $\boldsymbol{y}'$ in $F_{\boldsymbol{W},\boldsymbol{x}'}$.

Let L be the subset of $P_{\boldsymbol{W},\boldsymbol{x}'}$ identified with $\{\boldsymbol{z}\} \times \mathcal{M}_\omega(\boldsymbol{W}^c)$ by Lemma 3.3. Certainly L separates $P_{\boldsymbol{W},\boldsymbol{x}'}$. We claim L also separates $\mathcal{M}_\omega(S)$, with $\boldsymbol{x}$ and $\boldsymbol{y}$ on different sides. This follows immediately from Theorem 3.5:

Recall the map $\Phi = \Phi_{\boldsymbol{W},\boldsymbol{x}'} \colon \mathcal{M}_\omega(S) \to F_{\boldsymbol{W},\boldsymbol{x}'}$, and, also, that $\boldsymbol{x}' = \Phi(\boldsymbol{x})$ and $\boldsymbol{y}' = \Phi(\boldsymbol{y})$. Divide $F_{\boldsymbol{W},\boldsymbol{x}'} \setminus \{\boldsymbol{z}\}$ into two disjoint open sets $E_{\boldsymbol{x}}$ and $E_{\boldsymbol{y}}$ containing $\boldsymbol{x}'$ and $\boldsymbol{y}'$, respectively. $\Phi^{-1}(E_{\boldsymbol{x}})$ and $\Phi^{-1}(E_{\boldsymbol{y}})$ are open sets containing $\boldsymbol{x}$ and $\boldsymbol{y}$ respectively. The remainder $\Phi^{-1}(\{\boldsymbol{z}\})$ consists of L union an open set V, by the locally constant property. Hence we have divided $\mathcal{M}_\omega(S) \setminus L$ into three disjoint open sets two of which contain $\boldsymbol{x}$ and $\boldsymbol{y}$ respectively. This proves L separates $\boldsymbol{x}$ and $\boldsymbol{y}$.

The construction exhibits L as an asymptotic cone $\mathcal{M}_\omega(\boldsymbol{W}^c)$, from which it follows that L is closed (cf. [dDW]). Since the topological type of $\boldsymbol{W}^c$ is ω-a.e. constant, this is isometric to $\mathcal{M}_\omega(W^c)$ for some fixed surface W^c. $\square$

4. The dimension theorem

In this section we will apply the separation Theorem 3.6 to prove the main theorem on dimension in $\mathcal{M}_\omega(S)$. We begin with some terminology:

Historically, topologists have studied three different versions of dimension: small inductive dimension, ind, large inductive dimension, Ind, and covering dimension, dim (the covering dimension is also called the topological dimension). Dimension theory grew out of the development of these various definitions and studies the interplay and applications of the various versions of dimension [Eng2]. For a topological space X, let $\widehat{\mathrm{ind}}(X)$ denote the supremum of $\mathrm{ind}(X')$ over all locally compact subsets $X' \subset X$, and similarly define $\widehat{\mathrm{Ind}}$ and $\widehat{\dim}$. Restating our main theorem, we have:

THEOREM 4.1. $\widehat{\mathrm{ind}}(\mathcal{M}_\omega(S)) = \widehat{\mathrm{Ind}}(\mathcal{M}_\omega(S)) = \widehat{\dim}(\mathcal{M}_\omega(S)) = r(S)$.

The Rank Conjecture follows immediately as a corollary, since $\mathbb{R}^n$ is locally compact and $\mathrm{ind}(\mathbb{R}^n) = n$.

4.1. *Separation and dimension.* We will work with inductive dimension, which we define below. Equivalence of the different dimensions in our setting is provided by

LEMMA 4.2. *For a metric space* X, $\widehat{\dim}(X) = \widehat{\mathrm{ind}}(X) = \widehat{\mathrm{Ind}}(X)$.

Proof. This is essentially an appeal to the literature. First note the following standard topological facts:

(1) every metric space is paracompact;

(2) a locally compact space is paracompact if and only if it is strongly paracompact [Eng1, p. 329].

Engelking shows [Eng2, p. 220] that if Y is a strongly paracompact metrizable space, then $\mathrm{ind}(Y) = \mathrm{Ind}(Y) = \dim(Y)$. Thus, if $X' \subset X$ is a locally compact subset, then $\mathrm{ind}(X') = \mathrm{Ind}(X') = \dim(X')$. Taking the supremum over locally compact subsets finishes the proof. □

To prove Theorem 4.1 we provide a lemma reducing this result to Theorem 3.6. First we recall the definition of the *small inductive dimension*: $\mathrm{ind}(\emptyset) = -1$ and for any X, $\mathrm{ind}(X) = n$ if n is the smallest number such that for all $x \in X$ and neighborhood V of x, there exists a neighborhood $x \in U \subset V$ such that $\mathrm{ind}(\partial U) \leq n-1$. Here ∂U is the topological frontier of U in Y. (See [Eng2] for further details.)

LEMMA 4.3. *If* X *is a metric space for which every pair of points can be separated by a closed subset* $L \subset X$ *with* $\widehat{\mathrm{ind}}(L) \leq D-1$, *then* $\widehat{\mathrm{ind}}(X) = \widehat{\mathrm{Ind}}(X) = \widehat{\dim}(X) \leq D$.

Proof. By Lemma 4.2, we may henceforth restrict our attention to the small inductive dimension.

Let X' be a locally compact subset of X. Fixing $x \in X'$, consider any ϵ-ball B about x in the induced metric on X', where ϵ is assumed to be sufficiently small so that local compactness of X' implies ∂B is compact. For any $y \in \partial B$, let L be a closed separator of x and y, with $\widehat{\operatorname{ind}}(L) \leq D-1$, as provided by hypothesis. Since X' is locally compact, $L' = X' \cap L$ has $\operatorname{ind}(L') \leq D-1$. The separation property means that $X' \setminus L'$ is the union of a pair of disjoint open subsets of X', W_y and V_y, such that $x \in W_y$ and $y \in V_y$. Since ∂B is compact, we may extract a finite subcover of the covering $\{V_y\}$ of ∂B, which we relabel $V_1, \ldots, V_n$, with corresponding separators $L_1, \ldots, L_n$ and complementary $W_1, \ldots, W_n$. Then $\cup L_i'$ separates x from ∂B. More precisely, let $\mathcal{W} = \cap W_i$ and $\mathcal{V} = \cup V_i$. (In case $\partial B = \emptyset$, let $\mathcal{W} = X'$ and $\mathcal{V} = \emptyset$.) These are disjoint open sets with $x \in \mathcal{W}$, $\partial B \subset \mathcal{V}$, and $\partial \mathcal{W} \subset \cup L_i'$.

Now let $U = \mathcal{W} \cap B$. This is an open set, contained in B, whose boundary is contained in $\cup L_i'$ (since it cannot meet ∂B which lies in $\mathcal{V}$). Since ind is preserved by finite unions and monotonic with respect to inclusion, we have $\operatorname{ind}(\partial U) \leq D-1$, which is what we wanted to prove. □

4.2. *Proof of the dimension theorem.* We can now complete the proof of Theorem 4.1, by induction on $r(S)$.

Note that the lower bound $\widehat{\operatorname{ind}}(\mathcal{M}_\omega(S)) \geq r(S)$ is immediate since maximal abelian subgroups give quasi-isometrically embedded $r(S)$-flats [FLM]. We now prove the upper bound.

When $r(S) = 1$, S is $S_{1,1}$, $S_{0,4}$ or $S_{0,2}$. The asymptotic cones for the first two are the asymptotic cone for $\mathrm{SL}(2,\mathbb{Z})$ which is known to be an $\mathbb{R}$-tree. In the third case we really have in mind the annulus complex of an essential annulus, for which the asymptotic cone is just $\mathbb{R}$. Since $\widehat{\operatorname{ind}} = 1$ is well known for $\mathbb{R}$-trees, the theorem holds in this case.

Theorem 3.6 provides for each $x, y \in \mathcal{M}_\omega(S)$ a separator, L, which is homeomorphic to $\mathcal{M}_\omega(W^c)$, where W is an essential subsurface of S. Since r is additive over disjoint unions and $r(W) \geq 1$, we have $r(W^c) \leq r(S) - 1$. Thus by induction $\widehat{\operatorname{ind}}(L) \leq r(S) - 1$. (We can apply the inductive hypothesis to each component of W^c, and use subadditivity of ind over finite products, see [Eng2], and additivity of r over disjoint unions.)

Thus we have satisfied the hypotheses of Lemma 4.3 for $\mathcal{M}_\omega(S)$, and Theorem 4.1 follows.

5. Teichmüller space

In this section we deduce analogues of the results in the earlier sections for Teichmüller space with the Weil-Petersson metric. As shown in Brock [Bro], there is a combinatorial model for the Weil-Petersson metric on Teichmüller space provided by the *pants graph.* The combinatorial analysis as carried out above for the mapping class group can be done similarly in the pants graph, (cf. [MM2, §8]). Using Brock's result, we deduce the results below about Teichmüller space, while working only with the pants graph.

The rank statement we obtain below is also obtained, for $S_{2,0}$, by Brock-Masur [BM], as a consequence of an analysis of the special properties of quasi-geodesics in the pants graph for the genus 2 case.

Recall that the *Teichmüller space* of a topological surface is the deformation space of finite area hyperbolic structures which can be realized on that surface. Teichmüller space has many natural metrics, here we consider the *Weil-Petersson metric* which is a Kähler metric with negative sectional curvature.

Definition 5.1. The *pants graph* of S is a simplicial complex, $\mathcal{P}(S)$, with the following simplices:

(1) **Vertices**: one vertex for each pants decomposition of S, i.e., a top dimensional simplex in $\mathcal{C}(S)$.

(2) **Edges**: connect two pants decompositions by an edge if they agree on all but one curve, and those curves differ by an edge in the curve complex of the complexity one subsurface (complementary to the rest of the curves) in which they lie.

The following result of Brock [Bro] allows us to work with the pants graph in our study of Teichmüller space.

THEOREM 5.2. *$\mathcal{P}(S)$ is quasi-isometric to the Teichmüller space of S with the Weil-Petersson metric.*

An important remark recorded in [MM2] is that the pants graph is exactly what remains of the marking complex when annuli (and hence transverse curves) are ignored. Hence, one obtains the following version of Theorem 1.5:

THEOREM 5.3. *If $\mu, \nu \in \mathcal{P}(S)$, then there exists a constant $K(S)$, depending only on the topological type of S, such that for each $K > K(S)$ there exists $a \geq 1$ and $b \geq 0$ for which*:

$$d_{\mathcal{P}(S)}(\mu, \nu) \approx_{a,b} \sum_{\text{nonannular } Y \subseteq S} \left\{\!\!\left\{ d_{\mathcal{C}(Y)}(\pi_Y(\mu), \pi_Y(\nu)) \right\}\!\!\right\}_K .$$

We note that in [Be], analogues of both Theorems 1.4 and 3.1 are proved to hold for the pants graph of any surface of finite type. Further, by the above heuristic argument about ignoring annuli, one obtains product regions as produced for the mapping class group in Section 2. Again these product regions are quasi-isometrically embedded with uniform constants; in the pants graph the identification is:

$$Q_{\mathcal{P}(S)}(\Delta) \cong \prod_{\text{nonannular } U \in \sigma(\Delta)} \mathcal{P}(U). \tag{5.1}$$

This identification leads to the main difference between the case of the pants graph and the mapping class group; namely, one obtains different counts of how many distinct factors occur on the right-hand side of the above equation. In the mapping class group, this number is $3g+p-3$, whereas in the case of the pants graph, the count is easily verified to be $\lfloor \frac{3g+p-2}{2} \rfloor$.

As in the case of the mapping class group, one obtains:

LEMMA 5.4. *If $\mu \in \mathcal{P}(S)$ then*

$$d(\mu, Q_{\mathcal{P}(S)}(\Delta)) \approx \sum_{\substack{W \pitchfork \Delta \neq \emptyset \\ W \text{ nonannular}}} \left\{\!\!\left\{ d_{\mathcal{C}(W)}(\mu, \Delta) \right\}\!\!\right\}_K .$$

The remainder of the argument is completed as for the mapping class group, except for the count on the dimension of the separators. In the pants graph one obtains:

LEMMA 5.5. *For any two points $x, y \in \mathcal{P}_\omega$ there exists a closed set $L \subset \mathcal{P}_\omega$ which separates x from y, and such that $\widehat{\mathrm{ind}}(L) \leq \lfloor \frac{3g+p-2}{2} \rfloor - 1$.*

Thus, we have shown:

DIMENSION THEOREM FOR TEICHMÜLLER SPACE. *Every locally compact subset of an asymptotic cone of Teichmüller space with the Weil-Petersson metric has topological dimension at most $\lfloor \frac{3g+p-2}{2} \rfloor$.*

The Rank Theorem for Teichmüller space now follows just as for the mapping class group.

UNIVERSITY OF UTAH, SALT LAKE CITY, UT
E-mail address: jason@math.utah.edu

YALE UNIVERSITY, NEW HAVEN, CT
E-mail address: yair.minsky@yale.edu

References

[Ara] J. ARAMAYONA, The Weil-Petersson geometry of the five-times punctured sphere, in *Spaces of Kleinian Groups*, 219–231, *London Math. Soc. Lecture Notes Ser*. **329**, Cambridge Univ. Press, Cambridge, 2006, Ph.D. thesis, University of Southampton, U.K., May 2004.

[Be] J. BEHRSTOCK, Asymptotic geometry of the mapping class group and Teichmüller space, *Geometry & Topology* **10** (2006) 1523–1578.

[Bir] J. Birman, *Braids, Links, and Mapping Class Groups*, *Annals of Math. Studies* **82**, Princeton Univ. Press, Princeton, NJ, 1974.

[BLM] J. Birman, A. Lubotzky, and J. McCarthy, Abelian and solvable subgroups of the mapping class groups, *Duke Math. J.* **50** (1983), 1107–1120.

[Bow] B. Bowditch, Intersection numbers and the hyperbolicity of the curve complex, *J. Reine Angew. Math.* **598** (2006), 105–129.

[Bro] J. Brock, The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores, *J. Amer. Math. Soc.* **16** (2003), 495–535 (electronic).

[BF] J. Brock and B. Farb, Curvature and rank of Teichmüller space, *American J. Math.* **128** (2006), 1–22.

[BM] J. Brock and H. Masur, Quasi-geodesics and relative hyperbolicity in the pants complex, in preparation.

[Bur] J. Burillo, Dimension and fundamental groups of asymptotic cones, *J. London Math. Soc.* **59** (1999), 557–572.

[Deh] M. Dehn, *Papers on Group Theory and Topology*, Springer-Verlag, New York, 1987, translated from the German and with introductions and an appendix by John Stillwell, with an appendix by Otto Schreier.

[dDW] L. van den Dries and A. Wilkie, Gromov's theorem on groups of polynomial growth and elementary logic, *J. Algebra* **89** (1984), 349–374.

[DS] C. Druţu and M. Sapir, Tree-graded spaces and asymptotic cones of groups, *Topology* **44** (2005), 959–1058, with an appendix by Denis Osin and Mark Sapir.

[Eng1] R. Engelking, *General Topology*, *Sigma Series in Pure Math.* **6**, second edition, Heldermann Verlag, Berlin, 1989, translated from the Polish by the author.

[Eng2] ———, *Theory of Dimensions Finite and Infinite*, *Sigma Series in Pure Math.* **10**, Heldermann Verlag, Lemgo, 1995.

[FLM] B. Farb, A. Lubotzky, and Y. Minsky, Rank-1 phenomena for mapping class groups, *Duke Math. J.* **106** (2001), 581–597.

[Gro1] M. Gromov, Groups of polynomial growth and expanding maps, *IHES Sci. Publ. Math.* **53** (1981), 53–73.

[Gro2] ———, Infinite groups as geometric objects, in *Proc. of the International Congress of Mathematicians, Warsaw*, 385–392, Amer. Math. Soc., Providence, RI, 1983.

[Ham] U. Hamenstädt, Geometry of the mapping class groups III: Quasi-isometric rigidity, Preprint; arXiv:math.GT/0512429.

[KaL] M. Kapovich and B. Leeb, Quasi-isometries preserve the geometric decomposition of Haken manifolds, *Invent. Math.* **128** (1997), 393–416.

[KlL] B. Kleiner and B. Leeb, Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings, *IHES Publ. Math.* **86** (1997), 115–197.

[MM1] H. Masur and Y. Minsky, Geometry of the complex of curves, I. Hyperbolicity, *Invent. Math.* **138** (1999), 103–149.

[MM2] ———, Geometry of the complex of curves II: Hierarchical structure, *Geom. Funct. Anal.* **10** (2000), 902–974.

[Mos] L. Mosher, Mapping class groups are automatic, *Ann. of Math.* **142** (1995), 303–384.

(Received January 13, 2006)
(Revised January 11, 2007)