

Original citation:

Bowditch, B. H.. (2016) Large-scale rank and rigidity of the Teichmuller metric. Journal of Topology, 9 . pp. 985-1020.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/79124>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

"This is the peer reviewed version of the following article: Bowditch, B. H.. (2016) Large-scale rank and rigidity of the Teichmuller metric. Journal of Topology, 9 . pp. 985-1020. doi: 10.1112/jtopol/jtw017 has been published in final form at

<http://dx.doi.org/10.1112/jtopol/jtw017>

This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for SelfArchiving."

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

LARGE-SCALE RANK AND RIGIDITY OF THE TEICHMÜLLER METRIC

BRIAN H. BOWDITCH

ABSTRACT. We study the coarse geometry of the Teichmüller space of a compact orientable surface in the Teichmüller metric. We describe when this admits a quasi-isometric embedding of a euclidean space, or a euclidean half-space. We prove quasi-isometric rigidity for Teichmüller space of a surface of complexity at least 2: a result proven independently by Eskin, Masur and Rafi. We deduce that, apart from some well known coincidences, the Teichmüller spaces are quasi-isometrically distinct. (See also Lemma 2.5 for further discussion.) We also show that Teichmüller space satisfies a quadratic isoperimetric inequality. A key ingredient for proving these results is the fact that Teichmüller space admits a ternary operation, natural up to bounded distance, which endows the space with the structure of a coarse median space whose rank is equal to the complexity of the surface. From this, one can also deduce that any asymptotic cone is bilipschitz equivalent to a CAT(0) space, and so in particular, is contractible.

1. INTRODUCTION

In this paper we explore various properties of the large scale geometry of the Teichmüller space of a compact orientable surface in the Teichmüller metric. In particular, we prove a number of results relating to the coarse rank of Teichmüller space, as well as quasi-isometric rigidity. Our starting point is a combinatorial model of Teichmüller space [Ra, D1], which we use to show that Teichmüller space admits a coarse median structure in the sense of [Bo1]. From this, a number of facts follow immediately, though others require additional work. As we will note, some of these results have been obtained in some form before, while others seem to be new. The paper makes use of constructions in [Bo4], which studies the large-scale geometry of the mapping class group from a similar perspective. The idea of using medians (or “centroids”) in the mapping class group originates in [BeM]. We remark that Teichmüller space in the Weil-Petersson metric also admits a coarse median [Bo5]; which also has consequences for the large-scale geometry of that space.

Let Σ be a compact orientable surface of genus g with p boundary components. Let $\xi = \xi(\Sigma) = 3g + p - 3$ be the *complexity* of Σ . Unless otherwise stated, we will assume in this introduction that $\xi(\Sigma) \geq 2$. We will sometimes use $S_{g,p}$ to denote a generic surface of this type.

Date: First draft: 20th March 2015. Revised: 30th March 2016.

We write $\mathbb{T}(\Sigma)$ for the Teichmüller space of Σ , that is, the space of marked finite-area hyperbolic structures on the interior of Σ . We give $\mathbb{T}(\Sigma)$ the Teichmüller metric, ρ . This endows it with the structure of a complete Finsler manifold, diffeomorphic to $\mathbb{R}^{2\xi(\Sigma)}$. Note that the mapping class group, $\text{Map}(\Sigma)$, acts properly discontinuously on $\mathbb{T}(\Sigma)$. The properties we are mainly interested in here are quasi-isometry invariant, so for most of the paper we will be referring instead to the “decorated marking graph”, $\mathcal{R}(\Sigma)$, which is a slight variation on the “augmented marking graph” of [D1]. Both these spaces are equivariantly quasi-isometric to $\mathbb{T}(\Sigma)$.

The central observation of this paper is that Teichmüller space admits an equivariant median, with similar properties to that of the mapping class group [BeM, Bo1]. More specifically, we show:

Theorem 1.1. *There is a ternary operation, $\mu : \mathbb{T}(\Sigma)^3 \rightarrow \mathbb{T}(\Sigma)$, which is canonical up to bounded distance, and which endows $\mathbb{T}(\Sigma)$ with the structure of a coarse median space of rank $\xi(\Sigma)$.*

A more precise formulation of this (for $\mathcal{R}(\Sigma)$) is given as Theorem 4.1. Roughly speaking, this means that $\mathbb{T}(\Sigma)$ behaves like a median algebra of rank $\xi(\Sigma)$ up to bounded distance. In fact, any finite subset of $\mathbb{T}(\Sigma)$ lies inside a larger finite subset of $\mathbb{T}(\Sigma)$ which can be identified with the vertex set of a finite CAT(0) cube complex of dimension at most $\xi(\Sigma)$, in such a way that the median operation in $\mathbb{T}(\Sigma)$ agrees up to bounded distance with the usual median operation in a cube complex (see Section 2.5). As with the mapping class group, or the Weil-Petersson metric, the median can be characterised in terms of subsurface projection maps. Moreover, it is $\text{Map}(\Sigma)$ -equivariant up to bounded distance.

The proof of Theorem 1.1 depends on various properties of subsurface projection, as laid out in [Bo4]. It has also been recently noted in [BeHS2], that these properties are satisfied in any “hierarchically hyperbolic space”, which includes the mapping class group and Teichmüller space.

Theorem 1.1 is used here to prove various facts about the quasi-isometry type of $\mathbb{T}(\Sigma)$.

For example, the following is an immediate consequence:

Theorem 1.2. *For any compact surface Σ , $\mathbb{T}(\Sigma)$ satisfies a coarse quadratic isoperimetric inequality.*

Here the term “coarse quadratic isoperimetric inequality” refers to the standard quasi-isometric invariant formulation of the quadratic isoperimetric inequality, and will be made more precise in Section 5.

We note that Theorem 1.2 is also a consequence of the fact that Teichmüller space admits a bicombing, see [KR].

In fact, using some results about the local geometry of $\mathbb{T}(\Sigma)$, one can show that it satisfies a quadratic isoperimetric inequality in a more traditional sense (see Proposition 5.2).

We also have the following result, which has been proven by different methods in [EMR1]. An argument using asymptotic cones has been found independently by Durham [D2] (see the remark after Theorem 1.8). Another proof has been given in [BeHS1].

Theorem 1.3. *There is a quasi-isometric embedding of a euclidean n -dimensional half-space into $\mathbb{T}(\Sigma)$ if and only if $n \leq \xi(\Sigma)$.*

Of course, this implies that \mathbb{R}^n quasi-isometrically embeds in $\mathbb{T}(\Sigma)$ for any $n < \xi(\Sigma)$. The question of quasi-isometric embeddings of $\mathbb{R}^{\xi(\Sigma)}$ requires additional work. We will show:

Theorem 1.4. *There is a quasi-isometric embedding of the euclidean space, $\mathbb{R}^{\xi(\Sigma)}$, into $\mathbb{T}(\Sigma)$ if and only if Σ has genus at most 1, or is a closed surface of genus 2.*

Theorem 1.4 turns out to be equivalent to finding an embedded $(\xi(\Sigma) - 1)$ -sphere in the curve complex of Σ , which we will see happens precisely in the above cases (see Proposition 5.7).

Note that Theorem 1.3 immediately implies that if Σ and Σ' are both compact orientable surfaces and $\mathbb{T}(\Sigma)$ is quasi-isometric to $\mathbb{T}(\Sigma')$, then $\xi(\Sigma) = \xi(\Sigma')$. (One can distinguish some further cases by bringing Theorem 1.4 into play.) In fact we have:

Theorem 1.5. *If $\mathbb{T}(\Sigma)$ is quasi-isometric to $\mathbb{T}(\Sigma')$ then $\xi(\Sigma) = \xi(\Sigma')$, and if $\xi(\Sigma) = \xi(\Sigma') \geq 4$, then Σ and Σ' are homeomorphic. Moreover, if one of Σ or Σ' is homeomorphic to $S_{1,3}$, then they both are.*

This statement can be paraphrased by saying that if two Teichmüller spaces are quasi-isometric then they are isometric, since it is well known each of the pairs $\{S_{2,0}, S_{0,6}\}$, $\{S_{1,2}, S_{0,5}\}$ and $\{S_{1,1}, S_{0,4}\}$ have isometric Teichmüller spaces. They are all different apart from the above coincidences. Theorem 1.5 is a simple consequence of Theorem 1.6 below, together with the corresponding statement for the mapping class group.

Recall that we have a *thick part*, $\mathbb{T}_T(\Sigma)$, of $\mathbb{T}(\Sigma)$. Up to bounded Hausdorff distance, it can be defined in a number of equivalent ways. For example, given $\epsilon > 0$, it can be defined as the set those finite area hyperbolic structures with systole (i.e. length of the shortest closed geodesic) at least ϵ . For this, we need to chose ϵ sufficiently small so that $\mathbb{T}_T(\Sigma)$ is connected. Note that $\text{Map}(\Sigma)$ acts cocompactly on $\mathbb{T}_T(\Sigma)$. So, for example, $\mathbb{T}_T(\Sigma)$ is a bounded Hausdorff distance from any $\text{Map}(\Sigma)$ -orbit.

Theorem 1.6. *If $\phi : \mathbb{T}(\Sigma) \rightarrow \mathbb{T}(\Sigma)$ is any quasi-isometry, then the Hausdorff distance from $\phi(\mathbb{T}_T(\Sigma))$ and $\mathbb{T}_T(\Sigma)$ is finite, and bounded above in terms of $\xi(\Sigma)$ and the quasi-isometric parameters of ϕ .*

Our proof will use asymptotic cones. However, it has been pointed out to me by Kasra Rafi that it is also a consequence of the results of Mosher [Mo]

and Minsky [Mi1, Mi2] which imply that a quasigeodesic in Teichmüller space is stable if and only if it lies a bounded distance from the thick part. Thus, up to bounded distance, the thick part can be characterised as the union of all stable quasigeodesics.

Now, up to bounded Hausdorff distance, $\mathbb{T}_T(\Sigma)$ can be viewed as a uniformly embedded copy of (any Cayley graph of) $\text{Map}(\Sigma)$. It follows that ϕ gives rise to a quasi-isometry of $\text{Map}(\Sigma)$ to $\mathbb{T}_T(\Sigma)$. By quasi-isometric rigidity of $\text{Map}(\Sigma)$, [BeKMM, Ham] (see also [Bo4]), it follows that $\phi|_{\mathbb{T}_T(\Sigma)}$ agrees up to bounded distance with the map induced by an element of $\text{Map}(\Sigma)$.

Building on this, we get one of the main results of this paper, namely the quasi-isometric rigidity of the Teichmüller metric:

Theorem 1.7. *If $\xi(\Sigma) \geq 2$, there is some $m \geq 0$, depending only on $\xi(\Sigma)$ and quasi-isometry parameters, such that if $\phi : \mathbb{T}(\Sigma) \rightarrow \mathbb{T}(\Sigma)$ is a quasi-isometry, then there is some $g \in \text{Map}(\Sigma)$ such that $\rho(\phi x, gx) \leq m$ for all $x \in \mathbb{T}(\Sigma)$.*

This result has been obtained independently by Eskin, Masur and Rafi [EMR2], using different methods.

Retrospectively, of course, Theorem 1.6 is an immediate consequence of Theorem 1.7.

Theorem 1.7 is a coarse version of the well known result of Royden [Ro] that any isometry of $\mathbb{T}(\Sigma)$ is induced by an element of $\text{Map}(\Sigma)$.

The proofs of Theorems 1.3 to 1.7 involve studying an asymptotic cone, $\mathbb{T}^\infty(\Sigma)$, of $\mathbb{T}(\Sigma)$ (see [G, VaW]). As a consequence of Theorem 1.1, we know that $\mathbb{T}^\infty(\Sigma)$ is a topological median algebra of rank $\xi(\Sigma)$, and can be bilipschitz embedded in a finite product of \mathbb{R} -trees [Bo1, Bo2]. In particular, it is bilipschitz equivalent to a median metric space. We also derive the following facts:

Theorem 1.8. *Any asymptotic cone, $\mathbb{T}^\infty(\Sigma)$, of $\mathbb{T}(\Sigma)$ has locally compact dimension $\xi(\Sigma)$. It is bilipschitz equivalent to a $CAT(0)$ space (and so, in particular, is contractible).*

Here the *locally compact dimension* of a topological space is the maximal dimension of any locally compact subset.

Another proof that the asymptotic cone has locally compact dimension at most $\xi(\Sigma)$ has been found independently by Durham [D2]. From this, the “only if” part of Theorem 1.3 follows. It also follows from the discussion in [BeHS1].

In fact, analysing the structure of $\mathbb{T}^\infty(\Sigma)$ will be a significant part of the work of this paper.

We remark that there are various other spaces naturally associated to a compact surface. Of particular note are the Weil-Petersson metric on Teichmüller space, the (Cayley graph of the) mapping class group, and the curve complex. Some discussion of the quasi-isometry types of these and other related spaces can be found in [Y]. Various results regarding rank and rigidity of such spaces can

be found, for example, in [BeKMM, Ham, EMR1, EMR2, Bo1, Bo4, Bo5], and references therein.

Theorem 1.1 of this paper is proven in Section 4 (modulo the lower bound on rank, which will be a consequence of Theorem 1.3). Theorems 1.2, 1.3, 1.8 and half of 1.4 are proven in Section 5. Theorems 1.5 and 1.6 are proven in Section 7, and the proof of Theorem 1.7 is completed in Section 8. The remaining half of Theorem 1.4 is proven in Section 9.

Much of the work for this paper was carried out while visiting Tokyo Institute of Technology in 2014. I am grateful to that institution for its support and hospitality, and to Sadayoshi Kojima for his invitation. The material about quasi-isometric embeddings was added in 2015, together with a different proof of the existence of median. I also thank Kasra Rafi for his interest and comments.

2. BACKGROUND

In this section, we review a few items of background material.

2.1. Conventions and terminology.

If $x, y \in \mathbb{R}$, we often use the notation $x \sim y$ to mean that $|x - y|$ is bounded above by some additive constant. The factors determining the relevant constant at any given moment, if not specified, should be clear from context. Ultimately, it should only depend on the parameters of the hypotheses, namely the complexity of the surface, or quasi-isometry constants etc. If a, b are points in a metric space, we write $a \sim b$ to mean that they are a bounded distance apart. Again, the relevant bounds should be clear from context, and ultimately depend only on the parameters of the hypotheses.

If $x, y > 0$, we will similarly write $x \asymp y$ to mean that y is bounded above and below by increasing linear functions of x . Again, the factors determining these functions should be clear from context.

We will generally behave as though the relations \sim and \asymp were transitive, though clearly each application of the transitive law will implicitly entail a change in the defining constants.

Except when we are working in the asymptotic cone, maps between our various metric spaces are generally defined up to bounded distance. It will generally be assumed that maps between graphs send vertices to vertices.

In the context of asymptotic cones, we will be using ultraproducts of various sets or graphs associated to a surface, for example curves, multicurves and markings. In this case, we will understand a “curve” to mean an element of the ultraproduct of the set of curves, and a “standard curve” to mean an element of the original set, that is, a curve in the traditional sense. We apply similar terminology to multicurves and subsurfaces etc. We will elaborate on this later.

Let $\phi : (M, \rho) \rightarrow (M', \rho')$ be a map (not necessarily continuous) between metric spaces. We say that ϕ is *coarsely lipschitz* if for all $x, y \in M$, $\rho'(\phi x, \phi y)$ is

bounded above by a linear function of $\rho(x, y)$. We say it is a *quasi-isometric embedding* if, in addition, $\rho(x, y)$ is bounded above by a linear function of $\rho'(\phi x, \phi y)$. A *quasi-isometry* is a quasi-isometric map with cobounded image (i.e., every point of M' is at a bounded distance from $\phi(M)$).

2.2. Marking graphs.

Let Σ be a compact orientable surface of genus g with p boundary components. Let $\xi = \xi(\Sigma) = 3g + p - 3$ be the *complexity* of Σ . We will write $S_{g,p}$ to denote the topological type of Σ .

As usual, a *curve* in Σ will mean a homotopy class of essential non-peripheral simple closed curves (except when it refers to an element of the ultraproduct of such, as mentioned above). We write $\iota(\alpha, \beta)$ for the geometric intersection number of two curves, α, β . We write $\text{Map}(\Sigma)$ for the mapping class group of Σ . A *multicurve* in Σ is a set of pairwise disjoint curves. Here we will generally allow empty multicurves. We write $\mathcal{S} = \mathcal{S}(\Sigma)$ for the set of multicurves on Σ . A multicurve is *complete* if it cuts Σ into $S_{0,3}$'s. A complete multicurve has exactly $\xi(\Sigma)$ components.

Central to the discussion are the notions of “markings” and “marking graphs”. A specific construction of a marking graph is described in [MaM2]; though the notion is quite robust, and there are many variations which could serve for our purposes. We summarise below the essential properties we need.

Given two finite sets, a, b , of curves on Σ , we write $\iota(a, b) = \max\{\iota(\alpha, \beta) \mid \alpha \in a, \beta \in b\}$. We abbreviate $\iota(a, \beta) = \iota(a, \{\beta\})$. Note that $\iota(a, a)$ measures the “self intersection” of a . (In fact, this quantity being finite is equivalent to a being finite. Indeed, $|a|$ is bounded in terms of $\iota(a, a)$.) We say that a collection, a , of curves *fills* Σ if $\iota(a, \gamma) > 0$ for all curves γ . (Less formally, this means that $\bigcup a$, realised so as to minimise total intersection, cuts Σ into discs and peripheral annuli.) By a *p-marking* on Σ , we mean a finite set, a , of curves which fill Σ and with $\iota(a, a) \leq p$. Note that, for any $p \in \mathbb{N}$, $\text{Map}(\Sigma)$ acts naturally on the set of p -markings, with finite quotient.

Definition. By a *marking graph* we mean a connected $\text{Map}(\Sigma)$ -invariant graph, \mathcal{M} , whose vertex set, \mathcal{M}^0 , consists of a set of p -markings for some $p < \infty$, and such that $\iota(a, b) \leq q$ for all adjacent $a, b \in \mathcal{M}^0$, for some $q < \infty$.

Here, of course, “ $\text{Map}(\Sigma)$ -invariant” means that \mathcal{M}^0 is closed under the natural action of $\text{Map}(\Sigma)$ induced by its action on the set of curves, and that adjacency is also respected by this action. Note that the conditions imply that the action is cofinite, i.e. the quotient is a finite graph.

It will be convenient to impose some more conditions on our marking graph. First, we suppose that some marking should contain a complete multicurve. (Note that it follows that if $a \in \mathcal{M}^0$ contains a multicurve, τ , then a is a bounded distance from some $b \in \mathcal{M}^0$, where b contains a complete multicurve containing τ .) We also require the following. If $a, b \in \mathcal{M}^0$ both contain a multicurve τ , then

they can be connected by a path in \mathcal{M} , whose vertices all contain τ , and whose length is bounded above in terms of the distance between a and b in \mathcal{M} .

The above properties are easy to arrange. In fact, they hold for the marking graph described in [MaM2]. Alternatively one can construct a graph, $\mathcal{M}(p, q)$, as follows. Given $q \geq p > 0$ and the set of vertices is the set of all p -markings, and we deem markings a, b to be adjacent if $\iota(a, b) \leq q$. One can check that if p, q are large enough, then the above conditions are satisfied for $\mathcal{M}(p, q)$. (In fact, $p = q = 4$ is sufficient for it to be a marking graph as defined.) This was the definition used in [Bo1]. We note that whatever marking graph we choose, it will embed into $\mathcal{M}(p, q)$ for large enough p, q . Moreover, this embedding will be a quasi-isometry.

We now fix $\mathcal{M} = \mathcal{M}(\Sigma)$ to be any marking graph satisfying the above conditions. Henceforth by a *marking* of Σ we will mean an element of \mathcal{M}^0 for this marking graph. We write ρ^\wedge for the combinatorial metric on \mathcal{M} , where each edge is assigned unit length. (This notation will be explained in Section 2.3 below.)

2.3. Subsurfaces.

By a *subsurface* in Σ , we mean a subsurface X of Σ , defined up to homotopy, such that the intrinsic boundary, ∂X , of X is essential in Σ , and such that X is not a three-holed sphere. We write \mathcal{X} for the set of subsurfaces. We can partition \mathcal{X} as $\mathcal{X}_A \sqcup \mathcal{X}_N$ into annular and non-annular subsurfaces. Given a curve, γ , in Σ , we write $X(\gamma) \in \mathcal{X}_A$, for the regular neighbourhood of γ . Given $X \in \mathcal{X}$, write $\partial_\Sigma X$ for the relative boundary of X in Σ , thought of as a multicurve in Σ .

Given $X, Y \in \mathcal{X}$, we have the following pentachotomy:

$X = Y$.

$X \prec Y$: $X \neq Y$, and X can be homotoped into Y but not into ∂Y .

$Y \prec X$: $Y \neq X$, and Y can be homotoped into X but not into ∂X .

$X \wedge Y$: $X \neq Y$ and X, Y can be homotoped to be disjoint.

$X \pitchfork Y$: none of the above.

We will be using subsurface projections to curve graphs and marking complexes.

We can associate to each $X \in \mathcal{X}$ the *curve graph*, $\mathcal{G}(X)$, in the usual way. Thus, if $X \in \mathcal{X}_N$, then the vertex set, $\mathcal{G}^0(X)$, is the set of curves in X , where two curves are adjacent if they have minimal possible intersection number. If $\xi(\Sigma) \geq 2$, this is 0. We write σ_X^\wedge for the combinatorial metric on $\mathcal{G}(X)$. (The caret superscript will be explained later.) It is shown in [MaM1] that $\mathcal{G}(X)$ is Gromov hyperbolic.

We will need to deal with annular surfaces differently. We begin with a general discussion.

Suppose that A is a compact topological annulus. Let $\mathcal{G}(A)$ be the graph whose vertex, $\mathcal{G}^0(A)$, consists of arcs connecting the two boundary components of A , defined up to homotopy fixing their endpoints, and where two such arcs are deemed adjacent in $\mathcal{G}(A)$ if they can be realised so that they meet at most at

their endpoints. We write σ_A^\wedge for the combinatorial metric on $\mathcal{G}(A)$. It is easily seen that $\mathcal{G}(A)$ is quasi-isometric to the real line. In fact, we will choose points x, y in the different boundary components, and let $\mathcal{G}_0(A)$ be the full subgraph of $\mathcal{G}(A)$, whose vertex set, $\mathcal{G}_0^0(A)$, consists of those arcs with endpoints at x and y . Now $\mathcal{G}_0(A)$ can be identified with real line, \mathbb{R} , with vertex set \mathbb{Z} . In fact, if t is the Dehn twist in A , and $\delta \in \mathcal{G}_0^0(A)$ is any fixed arc, then the map $[r \mapsto t^r \delta]$ gives an identification of \mathbb{Z} with the vertex set. One can also check that the inclusion of $\mathcal{G}_0(A)$ into $\mathcal{G}(A)$ is an isometric embedding. Moreover, $\mathcal{G}(A)$ is the 1-neighbourhood of its image. It follows that $|\sigma_A^\wedge(\delta, t^r \delta) - |r|| \leq 1$, for all $r \in \mathbb{Z}$. In fact, since x, y and δ can be chosen arbitrarily, this holds for all $\delta \in \mathcal{G}^0(A)$. (The fact that we have only an additive error here is important for later discussion.)

Now given $X \in \mathcal{X}$, we have a well defined ‘‘subsurface projection’’ map: $\theta_X^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(X)$, well defined up to bounded distance (see [MaM2]). (Here we are using the notation ‘‘ θ_X^\wedge ’’ to remind us that the we are dealing with marking graphs and curve graphs. In [Bo1, Bo4] the notation θ_X was used, first in a general setting, and then specialised to the marking graph. Here we will use the notation θ_X for projection between ‘‘decorated’’ graphs. which will play an equivalent role in this paper. This also explains the notation ρ^\wedge and σ_X^\wedge introduced above.)

We also have a projection map $\psi_X^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(X)$. In fact, these maps can be defined intrinsically to subsurfaces. In this way, if $Y \preceq X$, then $\theta_Y^\wedge \circ \psi_X^\wedge = \theta_Y^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(Y)$ and $\psi_Y^\wedge \circ \psi_X^\wedge = \psi_Y^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(Y)$. Moreover, if $\gamma \in a \in \mathcal{M}(\Sigma)$ with $\gamma \prec X$, we may always assume that $\gamma \in \psi_X^\wedge a$.

If $X \in \mathcal{X}_N$, we have a map $\chi_X^\wedge : \mathcal{M}(X) \rightarrow \mathcal{G}(X)$ where χ_X^\wedge simply selects one curve from the marking. If $X \in \mathcal{X}_A$, we set $\mathcal{M}(X)$ equal to \mathcal{G} , and take $\chi_X^\wedge : \mathcal{M}(X) \rightarrow \mathcal{H}(X)$ to be the identity map. In all cases, we have $\theta_X^\wedge = \chi_X^\wedge \circ \psi_X^\wedge$ (at least up to bounded distance).

If $Y \preceq X$ or $Y \pitchfork X$, then we also have a projection, $\theta_X^\wedge Y \in \mathcal{G}(X)$.

The distance formula of [MaM2] relates distances in $\mathcal{M}(\Sigma)$ to subsurface projection distances. In particular, they show:

Lemma 2.1. *There is some $r_0 \geq 0$ depending only on $\xi(\Sigma)$, such that if $r \geq r_0$ $a, b \in \mathcal{M}(\Sigma)$ then the set $\mathcal{A}^\wedge(a, b; r) = \{X \in \mathcal{X} \mid \sigma_X^\wedge(\theta_X^\wedge a, \theta_X^\wedge b) \geq r\}$ is finite. Moreover, $\rho(a, b) \asymp \sum_{X \in \mathcal{A}^\wedge(a, b)} \sigma_X^\wedge(\theta_X^\wedge a, \theta_X^\wedge b)$.*

Here the linear bounds implicit in \asymp depend only on $\xi(\Sigma)$ and r . (A similar formula for Teichmüller space was given in [Ra]. It is given as Proposition 4.8 here.)

One immediate consequence is the following:

Lemma 2.2. *Given $r \geq 0$, there is some $r' \geq 0$, depending only on $\xi(\Sigma)$ and r such that if $a, b \in \mathcal{M}(\Sigma)$ and $\sigma_X^\wedge(\theta_X^\wedge a, \theta_X^\wedge b) \leq r$ for all $X \in \mathcal{X}$, then $\rho^\wedge(a, b) \leq r$.*

Another important ingredient is the following lemma of Behrstock [Be] (given as Lemma 11.3 of [Bo1]).

Lemma 2.3. *There is a constant, l , depending only on $\xi(\Sigma)$ with the following property. Suppose that $X, Y \in \mathcal{X}$ with $X \pitchfork Y$, and that $a \in \mathcal{M}^0$. Then $\min\{\sigma_X^\wedge(\theta_X^\wedge a, \theta_X^\wedge Y), \sigma_Y^\wedge(\theta_Y^\wedge a, \theta_Y^\wedge X)\} \leq l$.*

2.4. Median algebras.

Let (M, μ) be a median algebra; that is, a set, M , equipped with a ternary operation, $\mu : M^3 \rightarrow M$, such that $\mu(a, b, c) = \mu(b, c, a) = \mu(c, a, b)$, $\mu(a, a, b) = a$ and $\mu(a, b, \mu(c, d, e)) = \mu(\mu(a, b, c), \mu(a, b, d), e)$, for all $a, b, c, d, e \in M$. (For an exposition, see for example [BaH].) Given $a, b \in M$, write $[a, b] = \{x \in M \mid \mu(a, b, x) = x\}$, for the *median interval* from a to b . A subset, A , of M is a *subalgebra* if $\mu(a, b, c) \in A$ for all $a, b, c \in A$. It is *convex* if $[a, b] \subseteq A$ for all $a, b \in A$. One checks easily that $[a, b]$ is itself convex. One defines homomorphisms and isomorphisms between median algebras in the obvious way. One also checks that the map $[x \mapsto \mu(a, b, x)]$ is a median retraction of M to $[a, b]$.

If $a, b \in M$, then $[a, b]$ admits a partial order, \leq , defined by $x \leq y$ if $x \in [a, y]$ (or equivalently $y \in [b, x]$). If $[a, b]$ has intrinsic rank 1 (as defined below), then this is a total order.

A (*directed*) *wall* in M is (equivalent to) an epimorphism of M to the two-point median algebra. One can show that any two points of M are separated by some wall (i.e., they have different images under the epimorphism).

An *n-cube* in M is a subalgebra isomorphic to the direct product on n two-point median algebras: $\{-1, 1\}^n$. The *rank* of M is the maximal n such that M contains an n -cube (deemed infinite, if there is no such bound).

There is a stronger notion of “ n -colourability” defined in [Bo1]. We say that two walls, $\phi : M \rightarrow [0, 1]$ and $\phi' : M \rightarrow [0, 1]$, *cross* if the product homomorphism, $(\phi, \phi') : M \rightarrow [0, 1]^2$ is surjective. We say that M is *n-colourable* if we can partition the walls into n “colours” so that no two walls of the same colour cross. It is not hard to see that this implies that M has rank at most n .

We will refer to a 2-cube as a *square*. We will generally denote a square by cyclically listing its points as a_1, a_2, a_3, a_4 , so that $\{a_i, a_{i+1}\}$ is a side for all i . Note that $a_i \in [a_{i-1}, a_{i+1}]$ for all i (which one can check is, in fact, an equivalent way of characterising a square, provided we assume the a_i to be pairwise distinct).

Two ordered pairs, a, b and a', b' of elements of M are said to be *parallel* if $[a, b'] = [a', b]$. This is equivalent to saying that $a, b' \in [a', b]$ and $a', b \in [a, b']$. It is also equivalent to saying that $(a = b \text{ and } a' = b')$ or $(a = a' \text{ and } b = b')$ or a, b, b', a' form a square. Note that this is an equivalence relation on the set of ordered pairs in M .

If a, b is parallel to a', b' , and $x \in [a, b]$, then $\mu(a, b, \mu(a', b', x)) = \mu(\mu(a, b, a'), \mu(a, b, b'), x) = \mu(a, b, x) = x$. Similarly, swapping a, b with a', b' , one sees that the maps $[x \mapsto \mu(a, b, x)]$ and $[x \mapsto \mu(a', b', x)]$ are inverse median isomorphisms between $[a, b]$ and $[a', b']$.

We need to give some discussion to “gate maps”. Suppose $C \subseteq M$ is (a-priori) any subset. A map $\omega : M \rightarrow C$ is a *gate map* if $\omega x \in [c, x]$ for all $x \in M$ and

$c \in C$. If such a map exists, then it is unique, and $\omega|_C$ is the identity. Moreover, C must be convex (since if $a, b \in C$ and $d \in [a, b]$, then $\omega d \in [a, d] \cap [d, b] = \{d\}$, so $d \in C$). In fact, given that C is convex, ωx is characterised by the fact that $[x, \omega x] \cap C = \{\omega x\}$. In particular, we see that if $y \in [x, \omega x]$ then $\omega y = \omega x$. We also claim that ω is a median homomorphism. For this, it is enough to show that if $c \in [a, b]$, then $\omega c \in [\omega a, \omega b]$. But now the statements $c \in [a, b]$, $\omega c \in [c, \omega b]$ and $\omega b \in [b, \omega c]$ together imply $\omega c \in [a, \omega b]$. Thus (by the same observation, with a, b, c replaced by $\omega b, a, \omega c$) we get $\omega c = \omega \omega c \in [\omega a, \omega b]$ as required. In other words, since $\omega|_C$ is the identity, ω is a median retraction of M onto C . In fact, if C is convex, then any median retraction to C is a gate map. (For if $x \in M$ and $y \in [x, \omega x] \cap C$, then $y = \omega y \in [\omega x, \omega \omega x] = [\omega x, \omega x] = \{\omega x\}$, so $y = \omega x$.)

In general, a gate map to a closed convex set might not exist. However, it will exist, for example, if C is compact, or if all intervals in M are compact. (The latter will hold in the cases of interest to us here.) Also, if $a, b \in M$, then it is easily checked that the map $[x \mapsto \mu(a, b, x)]$ is a gate map to $[a, b]$.

If $D \subseteq M$ is convex, and $\omega : M \rightarrow C$ is a gate map, then ωD is convex. To see this, suppose that $a, b \in D$, and that $x \in [\omega a, \omega b] \subseteq C$. Then, $\omega x = x$. Let $c = \mu(a, b, x) \in D$. Since ω is a median homomorphism, we have $x = \mu(\omega a, \omega b, x) = \mu(\omega a, \omega b, \omega x) = \omega c$, so $x \in \omega D$, as required. (Note that, if $D \cap C \neq \emptyset$, then $\omega(D) = D \cap C$.)

If $C, C' \subseteq M$ are closed convex, with respective gate maps $\omega_C, \omega_{C'}$, we say that C, C' are *parallel* if $\omega_{C'}|_C$ and $\omega_C|_{C'}$ are inverse bijections (necessarily median isomorphisms). In this case, C, C' are either disjoint or equal. As an example, we have seen that if a, b and a', b' are parallel pairs, then the intervals $[a, b]$ and $[a', b']$ are parallel. (Again, note that if $B \cap B' \neq \emptyset$, then $C = C' = B \cap B'$.)

More generally, suppose that we have closed convex sets, $B, B' \subseteq M$, with gate maps $\omega : M \rightarrow B$ and $\omega' : M \rightarrow B'$. Let $C = \omega B' \subseteq B$ and $C' = \omega' B \subseteq B'$. By an earlier observation, we know that C, C' are convex.

If $a \in B'$, then since $\omega' \omega a \in [a, \omega a]$ we see that $\omega \omega' \omega a \in [\omega a, \omega \omega a] = \{\omega a\}$, so $\omega \omega' \omega a = \omega a$. It follows that $\omega \omega' x = x$ for all $x \in C$. Similarly, $\omega' \omega|_{C'}$ is the identity. We see that $\omega'|_C$ and $\omega|_{C'}$ are inverse bijections between C and C' . We refer to them as *parallel maps*.

In summary, we have shown:

Lemma 2.4. *Suppose that $B, B' \subseteq M$ are closed convex subsets with gate maps $\omega : M \rightarrow B$ and $\omega' : M \rightarrow B'$. Let $C = \omega B' \subseteq B$ and $C' = \omega' B \subseteq B'$. Then C, C' are parallel convex sets, with $\omega'|_C$ and $\omega|_{C'}$ the inverse parallel isomorphisms.*

Now let $\lambda = \omega \omega' : M \rightarrow C$ and $\lambda' = \omega' \omega : M \rightarrow C'$. These are median retractions, hence gate maps to C and C' . Note that $\omega' \lambda = \omega'$ and $\omega \lambda' = \omega$. In other words, ω is the gate map to C' composed with a parallel map to C . Similarly for ω' .

Suppose that $a \in A$ and $b \in B'$. Then $\lambda a = \omega \omega' a \in [a, \omega' a]$. Also $\omega' a \in [a, \lambda' b]$, so $\lambda a \in [a, \lambda b]$. Similarly $\lambda' b \in [b, \lambda a]$. We also have $\lambda a, \lambda' b \in [a, b]$.

We will also use the notion of a topological median algebra. This consists of a hausdorff topological space, M , and a continuous ternary operation $\mu : M^3 \rightarrow M$ such that (M, μ) is a median algebra. We say that M is *locally convex* if every point has a base of convex neighbourhoods. We say that M is *weakly locally convex* if, given any open set $U \subseteq M$ and any $x \in U$, there is another open set, $V \subseteq U$, containing x such that if $y \in V$, then $[x, y] \subseteq U$. (In fact, finite rank together with weakly locally convex implies locally convex, see Lemma 7.1 of [Bo1].) All the topological median algebras that arise in this paper will be locally convex.

Examples of topological median algebras are median metric spaces. A *median metric space* is (equivalent to) a median algebra with a metric ρ^M , such that for all $a, b \in M$, $[a, b] = \{x \in M \mid \rho^M(a, b) = \rho^M(a, x) + \rho^M(x, b)\}$. Discussion of this can be found, for example, in [Ve, Bo3].

The following gives a means of obtaining a median metric space in the context of interest to us.

Lemma 2.5. *Suppose that (M, ρ) is a geodesic metric space equipped with a ternary operation, $\mu : M^3 \rightarrow M$ such that (M, μ) a finite-rank median algebra. Suppose that there is some $\kappa \geq 1$ such that for all $a, b, c, d \in M$ we have $\rho(\mu(a, b, c), \mu(a, b, d)) \leq \kappa \rho(c, d)$. Then, ρ is bilipschitz equivalent to a median metric, ρ^M , which induces the given median μ .*

In fact, the bilipschitz constant can be chosen explicitly to depend only on κ and $\text{rank}(M)$. The metric ρ^M is not canonically determined by ρ . However, if $C \subseteq M$ is convex and such that μ restricted to C a median metric, then we can assume that ρ^M equals ρ on C .

Under the stronger assumption that μ is finitely colourable, Lemma 2.5 is proven in [Bo3]. This is achieved by embedding M into a finite product of \mathbb{R} -trees by a bilipschitz median homomorphism and pulling back the l^1 metric. (This is sufficient for the applications of this paper.) In fact, as observed in [Bo4], if we just assume finite rank, then essentially the same argument suffices to give us a median metric. We remark that one can weaken the geodesic hypothesis to assume only that M is lipschitz path-connected (though we don't need that here).

If we assume in addition that M is complete, then it also follows from [Bo2] that every interval in M is compact.

Note that, in the situation described above, where $B, B' \subseteq M$ are closed convex subsets, then the gate maps to B and B' are both 1-lipschitz. If $b \in B$ and $b' \in B'$, we have noted that $\lambda b, \lambda' b' \in [b, b']$ and $\lambda' b' \in [\lambda b, b']$. It follows that $\rho^M(b, b') = \rho^M(b, \lambda b) + \rho^M(\lambda b, \lambda' b') + \rho^M(\lambda' b', b')$. Also, $\rho^M(b, C) = \rho^M(b, \lambda b)$ and $\rho^M(b', C') = \rho^M(b', \lambda' b')$. In particular, we get $\rho^M(b, b') \geq \rho^M(b, C) + \rho^M(b', C')$, where C, C' are the images of B' and B under the respective gate maps, as above. (This observation will be used in the proof of Lemma 8.2.)

Note that any \mathbb{R} -tree is a median metric space of rank 1. Hence, a direct product of n \mathbb{R} -trees with the l^1 metric is median metric space of rank n . We recall the following definition from [Bo4]:

Definition. An \mathbb{R} -tree is *almost furry* if it has no point of valence 2.

In other words, the complement of any point is either connected or has at least 3 components.

We note the following from [Bo4]:

Proposition 2.6. *Suppose that M is a median metric space of rank n , that D is a direct product of n almost furry \mathbb{R} -trees, and that $f : D \rightarrow M$ is a continuous injective map with closed image. Then f is a median homomorphism, and $f(D) \subseteq M$ is convex.*

2.5. Coarse median spaces.

Let (Λ, ρ) be a geodesic metric space. The following definition was given in [Bo1]:

Definition. We say that (Λ, ρ, μ) is a *coarse median space* if it satisfies:

(C1): There are constants, $k, h(0)$, such that for all $a, b, c, a', b', c' \in \Lambda$,

$$\rho(\mu(a, b, c), \mu(a', b', c')) \leq k(\rho(a, a') + \rho(b, b') + \rho(c, c')) + h(0).$$

(C2): There is a function $h : \mathbb{N} \rightarrow [0, \infty)$ such that $1 \leq |A| \leq p < \infty$, then there is a finite median algebra (Π, μ_Π) and a map $\lambda : \Pi \rightarrow \Lambda$ such that

$$\rho(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z)) \leq h(p)$$

for all $x, y, z \in \Pi$ and such that $\rho(a, \lambda\pi a) \leq h(p)$ for all $a \in A$.

We say that (Λ, ρ, μ) has *rank* at most n if Π can always be chosen to have rank at most n as a median algebra.

We say that (Λ, ρ, μ) is *n-colourable* if we can always choose Π to be n -colourable as a median algebra.

Given $a, b \in \Lambda$, write $[a, b] = \{\mu(a, b, x) \mid x \in \Lambda\}$, for the *coarse interval* from a to b . If $c \in [a, b]$, then one can check that $\mu(a, b, c) \sim c$.

Definition. We say that $C \subseteq \Lambda$ is *r-quasiconvex* if $[a, b] \subseteq N(C; r)$ for all $a, b \in C$. We say that C is *quasiconvex* if it is r -quasiconvex for some $r \geq 0$.

Note that a quasiconvex set, C , is always quasi-isometrically embedded, or more precisely, there is some s depending only on r and the parameters of Λ , such that the inclusion of $N(C; s)$ into Λ is a quasi-isometric embedding.

We can define a *coarse gate map* to be a map $\omega : \Lambda \rightarrow C$ such that $\mu(x, \omega x, c) \sim \omega x$ for all $x \in \Lambda$ and $c \in C$. Similarly as with gate maps in a median algebra, the existence of such a map implies that C is quasiconvex.

We note the following:

Lemma 2.7. *Suppose that $a, b, c \in \Lambda$ and $r \geq 0$, with $\rho(\mu(a, b, c), c) \leq r$. Then $\rho(a, c) + \rho(c, b) \leq k_1\rho(a, b) + k_2$, where k_1, k_2 are constants depending only on r and the parameters of Λ .*

Proof. This is equivalent to saying that $\rho(a, c)$ is linearly bounded above in terms of $\rho(a, b)$. By (C1), $\rho(\mu(a, c, a), \mu(a, c, b)) \leq k\rho(a, b) + h(0)$. Also, $\mu(a, c, a)$ and $\mu(a, c, b)$ are respectively at a bounded distance from a and c , so the statement follows. \square

In other words, if c lies “between” a and b in the coarse median sense, then $\rho(a, c) + \rho(c, b)$ agrees with $\rho(a, b)$ up to linear bounds.

A map, $\phi : (\Lambda, \rho, \mu) \rightarrow (\Lambda', \rho', \mu')$, between two coarse median spaces is said to be a *h-quasimorphism* if $\rho'(\phi\mu(x, y, z), \mu(\phi x, \phi y, \phi z)) \leq h$ for all $x, y, z \in \Lambda$. We abbreviate this to *quasimorphism* if the constant, h , is understood from context.

A particular example of a coarse median space is a Gromov hyperbolic space. In this case, the median is the usual centroid of three points (a bounded distance from any geodesic connecting any two of these points). Such a coarse median space has rank at most 1. (In fact, any rank-1 coarse median space arises in this way.)

2.6. Asymptotic cones.

Let \mathcal{Z} be a countable set. By Zorn’s lemma, \mathcal{Z} admits a non-principal ultrafilter, and we assume \mathcal{Z} to be equipped with some such. In general our constructions might depend on the choice of ultrafilter. It is unclear to what extent they do here; but in any case, all the properties we describe are valid for any such choice (which we fix, once and for all).

Given a \mathcal{Z} -sequence of sets, $\vec{A} = (A_\zeta)_\zeta$, we write $\mathcal{U}\vec{A}$ for its ultraproduct, that is, $\mathcal{U}\vec{A} = \left(\prod_{\zeta \in \mathcal{Z}} A_\zeta\right) / \approx$, where $(a_\zeta)_\zeta \approx (b_\zeta)_\zeta$ if $a_\zeta = b_\zeta$ almost always. If $A_\zeta = A$ is constant, we write $\mathcal{U}A = \mathcal{U}\vec{A}$. We can identify A as a subset of $\mathcal{U}A$ via constant sequences. We then refer to an element of A in $\mathcal{U}A$ as being *standard*.

For example, in this context, we will refer to an element of $\mathcal{U}\mathcal{G}^0(\Sigma)$ as a “curve”, and an element of $\mathcal{G}^0(\Sigma) \subseteq \mathcal{U}\mathcal{G}^0(\Sigma)$ as a “standard” curve. We can also talk about “subsurfaces” (in $\mathcal{U}\mathcal{X}$) and “standard subsurfaces” in \mathcal{X} . We will sometimes abuse terminology and view a (non-standard) multicurve as a set of disjoint (non-standard) curves.

Note that the ultraproduct of the reals, $\mathcal{U}\mathbb{R}$, is an ordered field, where the order is given by $x < y$ if and only if $x_\zeta < y_\zeta$ for almost all ζ . We say that $x \in \mathcal{U}\mathbb{R}$ is *infinitesimal* if $|x| < y$ for all $y \in \mathbb{R}$ with $y > 0$. We say that $x \in \mathcal{U}\mathbb{R}$ is *limited* if $|x| < y$ for some $y \in \mathbb{R}$. If we quotient $\mathcal{U}\mathbb{R}$ by the infinitesimals (i.e. two numbers are equivalent if they differ by an infinitesimal), we get the “extended reals”, \mathbb{R}^* , which is an ordered abelian group containing \mathbb{R} as the convex subgroup of *limited* extended reals.

Suppose that $((\Lambda_\zeta, \rho_\zeta))_\zeta$ is a \mathcal{Z} -sequence of metric spaces. Then $(\mathcal{U}\vec{\Lambda}, \mathcal{U}\rho)$ is a $(\mathcal{U}\mathbb{R})$ -metric space. After identifying points an infinitesimal distance apart, we get a quotient (Λ^*, ρ^*) , which is an \mathbb{R}^* -metric space. We say that two points of Λ^* lie in the same *component* if they are a limited distance apart. Thus, each

component of Λ is a metric space in the usual sense (that is, the metric takes real values). In fact, one can show that such a component is a complete metric space.

In particular, suppose that (Λ, ρ) is a fixed metric space, and that $t \in \mathcal{U}\mathbb{R}$ is a positive infinitesimal. Let $(\Lambda_\zeta, \rho_\zeta) = (\Lambda, t_\zeta \rho)$. In this case, we write (Λ^*, ρ^*) for the resulting \mathbb{R}^* -metric space (where the scaling factors, t_ζ , are implicitly understood). We will use Λ^∞ to denote an arbitrary component of Λ^* , and ρ^∞ for the restriction of ρ^* . Thus, $(\Lambda^\infty, \rho^\infty)$ is a complete metric space in the usual sense. This is called an *asymptotic cone* of Λ and we refer to Λ^* as the *extended asymptotic cone*. Given a \mathcal{Z} -sequence of points $x_\zeta \in \Lambda_\zeta$, and $x \in \Lambda^*$, write $x_\zeta \rightarrow x$ to mean that x is the class corresponding to $(x_\zeta)_\zeta$. If Λ is a geodesic space, so is Λ^∞ . We put the metric topology on each component of Λ^* , and topologise Λ^* as the disjoint union of its components. Note that Λ^* has a preferred basepoint, namely that corresponding to any constant sequence in Λ . The component of Λ^* containing this basepoint is sometimes referred to as *the* asymptotic cone of Λ . (Again, the choice of scaling factors is implicitly assumed.)

If a group Γ acts by isometry on Λ , we get an action of its ultraproduct, $\mathcal{U}\Gamma$, on Λ^* .

If (Λ, ρ) and (Λ', ρ') are metric spaces, and $\phi : \Lambda \rightarrow \Lambda'$ is a coarsely lipschitz map, we get an induced map $\phi^* : \Lambda^* \rightarrow (\Lambda')^*$, which is lipschitz (in the sense that the multiplicative bound is real). This restricts to a lipschitz map $\phi : \Lambda^\infty \rightarrow (\Lambda')^\infty$ (where “lipschitz” here has its usual meaning). If ϕ is a quasi-isometric embedding, then ϕ^∞ is bilipschitz onto its range. In particular, if ϕ is a quasi-isometry then, ϕ^∞ is a bilipschitz homeomorphism.

A similar discussion applies if we have a \mathcal{Z} -sequence of uniformly coarsely lipschitz maps, $\phi_\zeta : \Lambda \rightarrow \Lambda'$.

If (Λ, ρ, μ) is a coarse median space, then we get a ternary operation $\mu^* : (\Lambda^*)^3 \rightarrow \Lambda^*$ which restricts to $\mu^\infty : (\Lambda^\infty)^3 \rightarrow \Lambda^\infty$. One can check that (Λ^*, μ^*) is a median algebra, with $(\Lambda^\infty, \mu^\infty)$ as a subalgebra. Note that $(\Lambda^\infty, \mu^\infty)$ is a topological median algebra, in the sense that the median is continuous with respect to the topology induced by ρ^∞ .

If Λ has finite rank, ν , then Λ^* has rank at most ν as a median algebra. It is easy to see that Λ^∞ satisfies the hypotheses of Lemma 2.5 and so Λ^∞ is bilipschitz equivalent to a median metric space. (See [Bo2] for more discussion.) Also, intervals in Λ^∞ are compact. (In particular, we have a gate map for any closed convex subset.)

A standard example is that of a Gromov hyperbolic space, Λ , in which case, Λ^* is an \mathbb{R}^* -tree. Each component, Λ^∞ , is an \mathbb{R} -tree. For example, if Λ is the hyperbolic plane, then Λ^∞ is the unique complete 2^{\aleph_0} -regular tree. If Λ is a horodisc, then Λ^∞ is a closed subtree thereof. In both cases, Λ^∞ is almost furry (in the sense of Proposition 2.6).

3. A COMBINATORIAL MODEL

In this section, we describe a combinatorial model, $\mathcal{R} = \mathcal{R}(\Sigma)$, for $\mathbb{T}(\Sigma)$. It is a slight variation on the “augmented marking complex” described in [Ra] and in [D1]. In order to distinguish it, we will refer to the model described here as the “decorated marking complex”. The model $\mathcal{R}(\Sigma)$ contains $\mathcal{M}(\Sigma)$ as a subgraph. We define $\mathcal{R}(\Sigma)$ as follows.

A vertex, a , of \mathcal{R} consists of a marking, $\bar{a} \in \mathcal{M}^0$, together with a map $\eta_a : \bar{a} \rightarrow \mathbb{N}$ such that if $\alpha, \beta \in \bar{a}$, with $\eta_a(\alpha) > 0$ and $\eta_a(\beta) > 0$, then $\iota(\alpha, \beta) = 0$. Thus, $\hat{a} = \{\alpha \in \bar{a} \mid \eta_a(\alpha) > 0\}$ is a (possibly empty) multicurve in Σ . We refer to such an a as a *decorated marking*, and to $\eta_a(\alpha)$ as the *decoration* on α . Two decorated markings, $a, b \in \mathcal{R}^0$ are deemed adjacent in \mathcal{R} if one of the following three conditions hold:

(E1): $\bar{a} = \bar{b}$ and $|\eta_a(\alpha) - \eta_b(\alpha)| \leq 1$ for all $\alpha \in \bar{a}$.

(E2a): $\hat{a} = \hat{b}$, $\eta_a|_{\hat{a}} = \eta_b|_{\hat{b}}$ and \bar{a}, \bar{b} are adjacent in \mathcal{M} .

(E2b): $\hat{a} = \hat{b}$, $\eta_a|_{\hat{a}} = \eta_b|_{\hat{b}}$ and $b = t_\alpha^r a$, where $\alpha \in \bar{a}$, t_α is the Dehn twist about α , and $|r| \leq 2^{\eta_a(\alpha)}$.

We refer to condition (E1) as “vertical adjacency” and to (E2) (that is (E2a) or (E2b)) as “horizontal adjacency”. Given that \mathcal{M} is connected, it is easily seen that \mathcal{R} is connected also. We write ρ for the combinatorial metric on \mathcal{R} (assigning each edge unit length).

We say that an element $a \in \mathcal{R}^0$ is *thick* if $\hat{a} = \emptyset$. The *thick part*, \mathcal{R}_T , of \mathcal{R} is the full subgraph of \mathcal{R} whose vertex set consists of thick decorated markings. Note that there is a natural embedding, $v : \mathcal{M} \rightarrow \mathcal{R}_T \subseteq \mathcal{R}$, extending this inclusion. It is easily seen that this is a quasi-isometry, with respect to the intrinsic path metric induced on \mathcal{R}_T . Note that this is again robust — if we were to take a different marking complex satisfying the conditions laid out in Section 2, we would get a quasi-isometric space.

Given $a \in \mathcal{R}$, define the map $h : \mathcal{R} \rightarrow [0, \infty)$ by setting $h(a) = \rho(a, \mathcal{R}_T)$. It is easily checked that if $a \in \mathcal{R}^0$, then $h(a) = \sum_{\gamma \in \bar{a}} h_\gamma(a) = \sum_{\gamma \in \hat{a}} h_\gamma(a)$, where $h_\gamma(a) = \eta_a(\gamma)$. This will be used in Section 7.

In Section 2 above we described subsurface projections, $\theta_X^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(X)$. Here we need to modify that in the case where X is an annulus. So suppose that $X \in \mathcal{X}_A$. The open annular cover of Σ corresponding to X has a natural compactification to a compact annulus, $A(X)$ (cf. [MaM2]).

Now, to any annulus, A , we have associated a graph $(\mathcal{G}(A), \sigma_A^\wedge)$, as described in Section 2.3. From this, we can define the *decorated arc graph*, $\mathcal{H}(A)$. Its vertex set, $\mathcal{H}^0(A)$ is $\mathcal{G}^0(A) \times \mathbb{N}$, where (δ, i) and (ϵ, j) are deemed adjacent if either $\delta = \epsilon$ and $|i - j| = 1$, or if $i = j$ and $\sigma_A^\wedge(\delta, \epsilon) \leq 2^i$. We write σ_A for the induced combinatorial metric. One can see that $\mathcal{H}(A)$ is quasi-isometric to a horodisc in the hyperbolic plane. In fact, if $\mathcal{H}_T(A)$ is the full subgraph of $\mathcal{H}(A)$ with vertex

set $\mathcal{G}^0(A) \times \mathbb{N}$, then the inclusion of $\mathcal{H}_T(A)$ in $\mathcal{G}(A)$ is an isometric embedding with 1-cobounded image. Moreover the map sending $(t^r \delta, i)$ to the point $(r, 1 + i \log 2)$ in the upper-half-plane model gives a quasi-isometry of $\mathcal{H}_T(A)$ to the horodisc $\mathbb{R} \times [1, \infty)$.

Returning to $X \in \mathcal{X}_A$, we write $(\mathcal{G}(X), \sigma_X^\wedge) = (\mathcal{G}(A(X)), \sigma_{A(X)}^\wedge)$, and $(\mathcal{H}(X), \sigma_X) = (\mathcal{H}(A(X)), \sigma_{A(X)})$. Write $\mathcal{H}_T(X) = \mathcal{G}(X) \subseteq \mathcal{H}(X)$.

Let $\theta_X^\wedge : \mathcal{M} \rightarrow \mathcal{G}_0(X)$ be the usual subsurface projection map (which we assume sends vertices to vertices). This commutes with the Dehn twist, t , about the core curve of X . In particular, it follows from the above discussion that $|\sigma_X^\wedge(\theta_X^\wedge m, \theta_X^\wedge t^r m) - |r|| \leq 1$ for all $r \in \mathbb{Z}$ and $m \in \mathcal{M}$. If $a \in \mathcal{R}^0$, set $\theta_X a = (\theta_X^\wedge \bar{a}, i)$, where $i = \eta_a(\alpha)$ if $\alpha \in \bar{a}$, and $i = 0$ if $\alpha \notin \bar{a}$.

If $X \in \mathcal{X}_N$, we simply set $(\mathcal{H}(X), \sigma_X) = (\mathcal{G}(X), \sigma_X^\wedge)$. We define $\theta_X : \mathcal{R}^0 \rightarrow \mathcal{H}^0(X)$ just by setting $\theta_X a = \theta_X^\wedge(\bar{a})$.

Lemma 3.1. *If $X \in \mathcal{X}$, then the map $\theta_X : \mathcal{R}^0 \rightarrow \mathcal{H}^0(X)$ extends to a coarsely lipschitz map $\theta_X : \mathcal{R} \rightarrow \mathcal{H}(X)$.*

Proof. In other words, we claim that if $a, b \in \mathcal{R}^0$ are adjacent in \mathcal{R} , then $\sigma_X(\theta_X a, \theta_X b)$ is bounded above (in terms of $\xi(\Sigma)$). We deal with the three types of edges in turn.

(E1): We have $\bar{a} = \bar{b}$. If $X \in \mathcal{X}_N$ then $\theta_X a = \theta_X b$. If $X \in \mathcal{X}_A$, then the first coordinates of $\theta_X a$ and $\theta_X b$ are equal, and their second coordinates are differ by at most 1. Thus, $\sigma_X(\theta_X a, \theta_X b) \leq 1$.

(E2a): We have $\hat{a} = \hat{b}$ and $\eta_a|\hat{a} = \eta_b|\hat{b}$, and that $\sigma_X^\wedge(\theta_X^\wedge a, \theta_X^\wedge b)$ is bounded. If $X \in \mathcal{X}_N$, we are done. If $X \in \mathcal{X}_A$, the first coordinates are a bounded distance apart in $\mathcal{G}(X)$, and the second coordinates are equal (to $\eta_a \alpha = \eta_b \alpha$ or to 0, depending on whether or not the core curve, α , of X lies in \hat{a}).

(E2b): We have $b = t_\alpha^r a$ (so that $\hat{a} = \hat{b}$). Suppose first that X is not a regular neighbourhood of α . In this case, we have $\sigma_X^\wedge(\theta_X^\wedge \bar{a}, \theta_X^\wedge \bar{b})$ bounded. (To see this, let γ be any curve not homotopic into $\Sigma \setminus X$, with $\iota(\gamma, \alpha) = 0$, and with $\iota(\gamma, \delta)$ bounded for all $\delta \in \bar{a}$. Such a curve is easy to construct — note that we allow $\gamma = \alpha$. Now, $t_\alpha \gamma = \gamma$, and we see that $\iota(\gamma, \epsilon) = \iota(\gamma, t_\alpha^{-r} \epsilon)$ is bounded for all $\epsilon \in \bar{b}$. It follows that $\theta_X^\wedge \bar{a}$ and $\theta_X^\wedge \bar{b}$ are both a bounded distance from the projection of the curve γ to X in $\mathcal{G}(X)$.) Since $\mathcal{H}(X) = \mathcal{G}(X)$ in this case, we see that $\sigma_X(\theta_X a, \theta_X b)$ is bounded. We are therefore reduced to considering the case where X is an annulus with core curve α . From the earlier discussion, we know that $\sigma_X^\wedge(\theta_X^\wedge \bar{a}, \theta_X^\wedge \bar{b}) = \sigma_X^\wedge(\theta_X^\wedge \bar{a}, \theta_X^\wedge t_\alpha^r \bar{a})$ differs by at most 1 from $|r|$. Moreover, $|r| \leq 2^i$, where $i = \eta_a \alpha = \eta_b \alpha$. Note that i is also the second coordinate of $\theta_X a$ and of $\theta_X b$. By construction of $\mathcal{H}(X)$, we see that $\sigma_X(\theta_X a, \theta_X b) \leq 2$ in this case. \square

If $X = \Sigma$, we write $\chi_\Sigma = \theta_\Sigma : \mathcal{R}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$. Up to bounded distance, this simply selects some curve from the marking.

Also, if $X, Y \in \mathcal{X}$ with $Y \pitchfork X$ or $Y \prec X$, we write $\theta_X Y = \theta_X^\wedge Y \in \mathcal{G}(X) = \mathcal{H}^0(X) \subseteq \mathcal{H}(X)$ for the usual subsurface projection (as in [MaM2]).

Given $a \in \mathcal{R}$, $\alpha \in \hat{a}$ and $i \in \mathbb{N}$, write $a_i = a_i(\alpha) \in \mathcal{R}$ for the decorated marking obtained by changing the decoration on α to i . (That is, $\bar{a}_i = \bar{a}$, $\eta_{a_i}(\beta) = \eta_a(\beta)$ for all $\beta \in \bar{a} \setminus \{\alpha\}$, and $\eta_{a_i}(\alpha) = i$.) Write $t = t_\alpha$ for Dehn twist about α . Write $H_a(\alpha) = \{t^r a_i \mid r \in \mathbb{Z}, i \in \mathbb{N}\}$.

Lemma 3.2. *Suppose $a \in \mathcal{R}$, $\alpha \in \hat{a}$ and $X = X(\alpha)$. There is a quasi-isometric embedding $\kappa : \mathcal{H}(X) \rightarrow \mathcal{R}$ with image a bounded Hausdorff distance from $H_a(\alpha)$, and with $\theta_X \circ \kappa$ a bounded distance from the identity on $\mathcal{H}(X)$.*

Proof. Let $\delta = \theta_X^\wedge \bar{a}$. We can assume that $\delta \in \mathcal{H}_T^0(X) \subseteq \mathcal{H}^0(X)$. Thus, $\mathcal{H}_T^0(X) = \{(t^r \delta, i) \mid r \in \mathbb{Z}, i \in \mathbb{N}\}$. Define $\kappa|_{\mathcal{H}_T^0(X)}$ by $\kappa((t^r \delta, i)) = t^r a_i$. Thus, by construction, $\kappa(\mathcal{H}_T(X)) = H_a(\alpha)$ and $\theta_X \circ \kappa|_{(\mathcal{G}_0^0(X))}$ is the identity. If $b, c \in \mathcal{H}_T^0(X)$, the $\kappa b, \kappa c$ are connected by an edge of \mathcal{R} (of type (E1) or (E2b)). Thus, we can extend this to an embedding of $\mathcal{H}_T(X)$ in \mathcal{R} , and hence to a coarsely lipschitz map, $\kappa : \mathcal{H}(X) \rightarrow \mathcal{R}$. Given that it has a left quasi-inverse, this must be a quasi-isometric embedding. \square

Note that, in fact, we can see that the multiplicative constant of the quasi-isometry is 1 in this case, i.e. distances agree up to an additive constant. In other words, we see that if $b, c \in H_a(\alpha)$, then $|\rho(b, c) - \theta_X(b, c)|$ is bounded.

We can extend this to a statement about twists on multicurves. Given $a, b \in \mathcal{R}$, suppose that there is some $\tau \subseteq \hat{a} \cap \hat{b}$ such that b is obtained from a by applying powers of Dehn twists about elements of τ , and changing the decorations on these curves. In this case, we get:

Lemma 3.3. *If $a, b \in \mathcal{R}$ are as above, then $|\rho(a, b) - \sum_{\alpha \in \tau} \sigma_{X(\alpha)}(a, b)|$ is bounded.*

(We will only really need that $\rho(a, b) \asymp \sum_{\alpha \in \tau} \sigma_{X(\alpha)}(a, b)$.)

Note that for all $X \in \mathcal{X}$, $\mathcal{H}(X)$ is uniformly hyperbolic (in the sense of Gromov). In particular, they each admit a median operation $\mu_X : \mathcal{H}(X)^3 \rightarrow \mathcal{H}(X)$, well defined up to bounded distance, and such that $(\mathcal{H}(X), \mu_X)$ is a coarse median space of rank 1.

We will need the following observation:

Lemma 3.4. *$\mathcal{R}_T(\Sigma)$ is uniformly embedded in $\mathcal{R}(\Sigma)$.*

Proof. In fact, $\mathcal{R}_T(\Sigma)$ is exponentially distorted in $\mathcal{R}(\Sigma)$. Given $a, b \in \mathcal{R}_T^0(\Sigma)$, let $n = \rho(a, b)$. Let $a = a_0, \dots, a_n = b$ be vertex path from a to b in $\mathcal{R}(\Sigma)$. Certainly, $\rho(a_i, \bar{a}_i) = \rho(a_i, \mathcal{R}_T^0(\Sigma)) \leq n$ for all i . So by construction of $\mathcal{R}(\Sigma)$, we have $\rho(a_i, a_{i+1}) \leq 2^n$, so $\rho(a, b) \leq 2^n n$. \square

There are many variations on the construction of \mathcal{R} which would give rise to quasi-isometrically equivalent graphs. As observed in Section 2.2, our marking graph, \mathcal{M} , quasi-isometrically embeds into $\mathcal{M}(p, q)$ for all sufficiently large p, q . This naturally induces an embedding of \mathcal{R} into the decorated marking graph, $\mathcal{R}(p, q)$, similarly constructed from $\mathcal{M}(p, q)$. Moreover, it is easily checked that this inclusion is a quasi-isometry. Therefore, for our purposes, it doesn't matter

which marking graph we choose. Similarly, in (E2b) we could replace the exponent base, 2, with any fixed number strictly bigger than 1.

The construction of the “augmented marking graph” in [D1] fits (more or less) into this picture. There, the marking graph, \mathcal{M} , is taken to be the marking graph as defined in [MaM2]. There is a restriction that \hat{a} is required to be a subset of the “base” of a marking a (but any multicurve with bounded intersection with a will be the base of some nearby marking, so this requirement does not change things on a large scale). Our edges of type (E1), (E2a) and (E2b) correspond to “vertical moves”, “flip moves” and “twist moves” in the terminology there. Note that the exponent e is used instead of 2 for the twist moves. In any case, it is easily seen from the above, that the augmented marking complex of [D1] is equivariantly quasi-isometric to the decorated marking complex as we have defined it.

It was shown in [D1] that the augmented marking graph is equivariantly quasi-isometric to $\mathbb{T}(\Sigma)$. We deduce:

Proposition 3.5. *There is a $\text{Map}(\Sigma)$ -equivariant quasi-isometry of $\mathbb{T}(\Sigma)$ to $\mathcal{R}(\Sigma)$.*

Note that this quasi-isometry necessarily maps $\mathbb{T}_T(\Sigma)$ to within a bounded Hausdorff distance of $\mathcal{R}_T(\Sigma)$. (This can be seen explicitly from the various constructions, but also follows from the fact that the constructions are equivariant and that $\text{Map}(\Sigma)$ acts coboundedly on both $\mathbb{T}_T(\Sigma)$ and $\mathcal{R}_T(\Sigma)$.)

For the proof of Theorem 1.7 we will need to give some consideration to the complexity-1 case (see the “thick case” of the proof of Lemma 8.2). We briefly describe that here. (See also the discussion at the end of Section 4.)

So suppose that Σ is an $S_{1,1}$ or an $S_{0,4}$. In this case, we can take $\mathcal{G}^0(\Sigma)$, as usual, to be the set of curves in Σ . We deem α, β to be adjacent in \mathcal{G}^0 if they have minimal intersection (that is, $\iota(\alpha, \beta) = 1$ for $S_{1,1}$ and $\iota(\alpha, \beta) = 2$ for $S_{0,4}$). Thus, $\mathcal{G}(\Sigma)$ is a Farey graph, which we can identify with the 1-skeleton of a regular ideal tessellation of the hyperbolic plane, \mathbb{H}^2 .

We can take $\mathcal{M}(\Sigma)$ to be the dual 3-regular tree. Its vertices are at the centres of the ideal triangles, and its edges are geodesic segments. In this way, an element of $\mathcal{R}(\Sigma)$ consists of a triple of curves $\{\alpha, \beta, \gamma\}$ corresponding to a triangle in $\mathcal{G}(\Sigma)$, with decorations assigned to these curves, at most one of which is non-zero. We define a map $f : \mathcal{R}^0 \rightarrow \mathbb{H}^2$ as follows. If all the decorations of $\{\alpha, \beta, \gamma\}$ are 0, then we map it to the centre, m , of the corresponding triangle. If the decoration on α , say, is $i > 0$, then we map it to $\lambda(i \log 2)$ where $\lambda : [0, \infty) \rightarrow [0, \infty)$ is the geodesic ray with $\gamma(0) = m$, and tending to the ideal point of \mathbb{H}^2 corresponding to α . Mapping edges to geodesic segments, we get a map $f : \mathcal{R} \rightarrow \mathbb{H}^2$, which extends the inclusion of \mathcal{M} into \mathbb{H}^2 . It is easily checked that f is a quasi-isometry.

Note that each curve $\alpha \in \mathcal{G}^0$ corresponds to a component, $C(\alpha)$, of the complement of \mathcal{M} in \mathbb{H}^2 . Now $f|_{\mathcal{H}(\alpha)}$ is a quasi-isometry of $\mathcal{H}(\alpha)$ to $C(\alpha)$, which is in turn a bounded Hausdorff distance from a horodisc in \mathbb{H}^2 .

Finally note that there is a natural $\text{Map}(\Sigma)$ -equivariant identification of \mathbb{H}^2 (modulo scaling the metric by a factor of 2) with the Teichmüller space of Σ with the Teichmüller metric.

4. THE MEDIAN CONSTRUCTION

The main aim of this section will be to prove the following:

Theorem 4.1. *There is a coarsely lipschitz ternary operation $\mu : \mathcal{R}(\Sigma)^3 \rightarrow \mathcal{R}(\Sigma)$, unique up to bounded distance, such that for all $a, b, c \in \mathcal{R}(\Sigma)$ and all $X \in \mathcal{X}$, $\theta_X \mu(a, b, c)$ agrees up to bounded distance with $\mu_X(\theta_X a, \theta_X b, \theta_X c)$. Moreover, $(\mathcal{R}(\Sigma), \mu)$ is a finitely colourable coarse median metric space of rank $\xi(\Sigma)$.*

As we will see, the median is characterised up to bounded distance by the fact that all the subsurface projection maps, $\theta_X : \mathcal{R}(\Sigma) \rightarrow \mathcal{H}(X)$, are uniform quasi-morphisms for all $X \in \mathcal{X}$. Since this condition is $\text{Map}(\Sigma)$ -equivariant, the median is necessarily $\text{Map}(\Sigma)$ -equivariant up to bounded distance. Here all bounds depend only on $\xi(\Sigma)$. In view of Proposition 3.5, we see that this will then imply Theorem 1.1.

It is not hard to deduce the existence of medians on $\mathcal{R}(\Sigma)$ from the existence of medians in $\mathcal{M}(\Sigma)$. However, the approach we take here is to verify the properties laid out in [Bo4]. The statement then follows directly from Theorem 1.4 of that paper. We begin with a general discussion of decorated markings.

To summarise so far, we have a graph, (\mathcal{R}, ρ) , and a collection of uniformly hyperbolic spaces, $(\mathcal{H}(X), \sigma_X)$, indexed by the set, \mathcal{X} , of subsurfaces of Σ , together with uniformly coarsely lipschitz maps $\theta_X : \mathcal{R}(\Sigma) \rightarrow \mathcal{H}(X)$. The maps θ_X were constructed out of the uniformly lipschitz maps $\theta_X^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(X)$, for the family of graphs, $(\mathcal{G}(X), \sigma_X^\wedge)$. Note that these constructions are all $\text{Map}(\Sigma)$ -equivariant up to bounded distance.

Given $a, b \in \mathcal{R}$, we will often abbreviate $\sigma_X(a, b) = \sigma_X(\theta_X a, \theta_X b)$ and $\sigma_X^\wedge(a, b) = \sigma_X^\wedge(\theta_X^\wedge \bar{a}, \theta_X^\wedge \bar{b})$. We also write $\theta_X a = \theta_X \bar{a}$. If γ is a curve in Σ , we will abbreviate $\sigma_\gamma = \sigma_{X(\gamma)}$, $\sigma_\gamma^\wedge = \sigma_{X(\gamma)}^\wedge$, $\theta_\gamma = \theta_{X(\gamma)}$, $\theta_\gamma^\wedge = \theta_{X(\gamma)}^\wedge$, etc. We write $\mathcal{G}(\gamma) = \mathcal{G}(X(\gamma))$, $\mathcal{H}(\gamma) = \mathcal{H}(X(\gamma))$, etc.

The following two statements are immediate consequences of Proposition 4.8 below, though we offer a more direct proofs which only use properties of the decorated marking complex.

Lemma 4.2. *There is some $l_0 \geq 0$ such that for all $a, b \in \mathcal{R}(\Sigma)$, $\{X \in \mathcal{X} \mid \sigma_X(a, b) \geq l_0\}$ is finite.*

Proof. This is a simple consequence of the corresponding statement for $\mathcal{M}(\Sigma)$ [MaM2], given in Lemma 2.1 here. This tells us immediately that there are only finitely many such $X \in \mathcal{X}_N$ for large enough l_0 . We therefore need only consider $X \in \mathcal{X}_A$. Now, $\hat{a} \cup \hat{b}$ is finite, so we can assume that the core curve of X does not lie in $\hat{a} \cup \hat{b}$. But in this case, the second coordinates of $\theta_X a$ and of $\theta_X b$ are

bounded, so $\sigma_X(a, b)$ is bounded above in terms of $\sigma_X^\wedge(a, b)$. Lemma 2.1 again tells that there are only finitely many such X provided the lower bound is sufficiently large. \square

Lemma 4.3. *For all $l \geq 0$, there is some $l' \geq 0$, depending only on l and $\xi(\Sigma)$ such that if $a, b \in \mathcal{R}$ satisfy $\sigma_X(a, b) \leq l$ for all $X \in \mathcal{X}$, then $\rho(a, b) \leq l'$.*

Proof. First note that we can assume that $\hat{a} = \hat{b}$. For suppose that $\alpha \in \hat{a} \setminus \hat{b}$. Let $X(\alpha)$ be the regular neighbourhood of α . Since $\eta_b(\alpha) = 0$, we have $\eta_a(\alpha) \leq \sigma_{X(\alpha)}(a, b) \leq l$. Let $a_0 = a_0(\alpha) \in \mathcal{R}^0$ (i.e. we reset the decoration on α to 0). Thus, $\rho(a, a_0) = \eta_a(\alpha) \leq l$, and $\hat{a}_0 = \hat{a} \setminus \{\alpha\}$. We now replace a with a_0 . We continue with this process for all curves in $\hat{a} \setminus \hat{b}$ and in $\hat{b} \setminus \hat{a}$, until \hat{a} and \hat{b} are both equal to some (possibly empty) multicurve, τ , say. Note that this process does not change \bar{a} or \bar{b} , and moves a and b each a bounded amount. The hypotheses are still satisfied (since the maps θ_X are coarsely lipschitz).

If X is not a regular neighbourhood of any element of τ , then $\sigma_X^\wedge(\bar{a}, \bar{b}) = \sigma_X(a, b)$ is bounded in terms of l . Moreover, if $\beta \in \tau$, we can find some $r(\beta) \in \mathbb{Z}$ such that $\sigma_\beta^\wedge(\bar{b}, t_\beta^{r(\beta)}\bar{a})$ is bounded (since $\mathcal{G}(X)$ is quasi-isometric to the real line). Let g be the composition of the twists $t_\beta^{r(\beta)}$ as β ranges over τ , and set $e = ga \in \mathcal{R}^0$. We now get that $\sigma_X^\wedge(\bar{b}, \bar{e})$ is bounded for all $X \in \mathcal{X}$. Therefore by Lemma 2.2, we get $\rho^\wedge(\bar{b}, \bar{e})$ is bounded in terms of l . We are therefore reduced to the case where $\bar{a} = \bar{b}$.

But now, $|\eta_a(\beta) - \eta_b(\beta)| \leq \sigma_\beta(a, b)$ is bounded for all $\beta \in \tau = \hat{a} = \hat{b}$. In other words, all decorations differ by a bounded amount, so $\sigma(a, b)$ is bounded (using edges of type (E2a)). \square

The following is an analogue of Behrstock's Lemma [Be], given as Lemma 2.3 here.

Lemma 4.4. *There is a constant, l_1 , depending only on $\xi(\Sigma)$ with the following property. Suppose that $X, Y \in \mathcal{X}$ with $X \pitchfork Y$, and that $a \in \mathcal{R}_T^0$. Then $\min\{\sigma_X(a, Y), \sigma_Y(a, X)\} \leq l_1$.*

Proof. Suppose, for contradiction, that $\sigma_X(a, Y)$ and $\sigma_Y(a, X)$ are both large. Now by Lemma 2.3, we can assume, without loss of generality that $\sigma_X^\wedge(a, X)$ is bounded. This can only happen if $X = X(\gamma)$ for some curve $\gamma \in \hat{a}$. By assumption, γ crosses Y , so we see that $\sigma_Y^\wedge(a, X)$ is also bounded. We must therefore have $Y = X(\beta)$ for some $\beta \in \hat{a}$. But $\beta \pitchfork \alpha$, giving a contradiction, since \hat{a} is assumed to be a multicurve. \square

We will write

$$\langle x, y : z \rangle_\sigma = \frac{1}{2} (\sigma(x, z) + \sigma(y, z) - \sigma(x, y))$$

for the Gromov product of x, y with basepoint z . We will write $\langle x, y : z \rangle_X = \langle \theta_X x, \theta_X y : \theta_X z \rangle_{\sigma_X}$, and $\langle x, y : z \rangle_X^\wedge = \langle \theta_X^\wedge x, \theta_X^\wedge y : \theta_X^\wedge z \rangle_{\sigma_X^\wedge}$.

The following is an analogue of Lemma 11.4 of [Bo1]:

Lemma 4.5. *There are some l_1, l_2 , depending only on $\xi(\Sigma)$, with the following property. Suppose that $X, Y \in \mathcal{X}$ with $Y \prec X$, and suppose that $a, b \in \mathcal{R}(\Sigma)$ with $\langle a, b:Y \rangle_X \geq l_1$. Then $\sigma_Y(a, b) \leq l_2$.*

Proof. Note that $X \in \mathcal{X}_N$, so we have that $\langle a, b:Y \rangle_X^\wedge = \langle a, b:Y \rangle_X$ is large. From the bounded geodesic image theorem of [MaM2] (see Lemma 11.4 of [Bo1]), it follows that $\sigma_Y^\wedge(a, b)$ is bounded. If $Y \in \mathcal{X}_N$, $\sigma_Y(a, b) = \sigma_Y^\wedge(a, b)$ and we are done. So suppose that $Y = X(\gamma)$ for some curve γ . If $\sigma_\gamma(a, b)$ is large, then we must have $\gamma \in \hat{a} \cup \hat{b}$, and so we suppose $\gamma \in \hat{a} \subseteq \bar{a}$. But then $\langle a, b:Y \rangle_X \leq \sigma_X(a, Y)$ is bounded, giving a contradiction. \square

If $X \in \mathcal{X}_N$, we set $\mathcal{R}(X)$ to be the decorated marking complex of X (defined intrinsically to X). We have the map $\chi_X : \mathcal{R}(X) \rightarrow \mathcal{H}(X)$, defined up to bounded distance, where $\chi_X(a)$ just selects a curve from the marking, \bar{a} . If $X \in \mathcal{X}_A$, we set $\mathcal{R}(X) = \mathcal{H}(X)$, and in this case, take $\chi_X : \mathcal{R}(X) \rightarrow \mathcal{H}(X)$ to be the identity map.

Next, we need to define maps, $\psi_X : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(X)$ for $X \in \mathcal{X}$. As noted in Section 2.2, we already have corresponding maps $\psi_X^\wedge : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(X)$. Suppose, first that $X \in \mathcal{X}_N$, and that $a \in \mathcal{R}^0(\Sigma)$. Write $\tau = \{\gamma \in \hat{a} \mid \gamma \prec X\}$. As noted in Section 2.2, we can suppose that $\tau \subseteq \psi_X^\wedge \bar{a}$. Now let $\psi_X a$ be the decorated marking with decorations determined by $\eta_a|_\tau$. It is easily seen that ψ_X is uniformly coarsely lipschitz, so we also get a map $\psi_X : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(X)$. Note also that $\chi_X \circ \psi_X \sim \theta_X : \mathcal{R}(\Sigma) \rightarrow \mathcal{H}(X)$. In the case where $X \in \mathcal{X}_A$, we set $\mathcal{R}(X) = \mathcal{H}(X)$ and $\psi_X = \theta_X$. In this case, we set χ_X to be the identity map, (so trivially, $\chi_X \circ \psi_X = \theta_X$). Note that we can perform the above constructions intrinsically to any $X \in \mathcal{X}$; so if $Y \preceq X$, we get maps $\psi_{YX} : \mathcal{R}(X) \rightarrow \mathcal{R}(Y)$, with $\chi_Y \circ \psi_{YX} \sim \theta_Y$. It is also immediate from the construction (and corresponding fact in $\mathcal{M}(\Sigma)$) that if $Z \preceq Y \preceq X$, then $\psi_{ZY} \circ \psi_{YX} \sim \psi_{ZX}$.

We are now in a position to apply the results of [Bo4], which give a set of conditions which imply the existence of a coarse median on a geodesic space. Let us begin by summarising the current set-up.

Our spaces are indexed by the set, \mathcal{X} , of subsurfaces of Σ .

Given $X \in \mathcal{X}$, we write $\mathcal{X}(X) = \{Y \in \mathcal{X} \mid Y \preceq X\}$. For each X we have geodesic metric spaces, $(\mathcal{R}(X), \rho_X)$ and $(\mathcal{H}(X), \sigma_X)$ and a map $\chi_X : \mathcal{R}(X) \rightarrow \mathcal{H}(X)$. If $Y \in \mathcal{X}(X)$, we have a map $\psi_{YX} : \mathcal{R}(X) \rightarrow \mathcal{R}(Y)$. Also, if $Z \in \mathcal{X}$ with $Z \pitchfork X$ or $Z \prec X$, we have an element $\theta_X Z \in \mathcal{H}(X)$.

We set $\theta_{YX} = \chi_X \circ \psi_{YX} : \mathcal{R}(X) \rightarrow \mathcal{H}(Y)$. Given $a, b \in \mathcal{R}(X)$ and $Y, Z, W \in \mathcal{X}(X)$, we abbreviate $\sigma_X(a, b) = \sigma_X(\chi_X a, \chi_X b)$, $\rho_Y(a, b) = \rho_Y(\psi_{YX} a, \psi_{YX} b)$, $\sigma_Y(a, b) = \sigma_Y(\theta_{YX} a, \theta_{YX} b)$, $\sigma_Y(a, Z) = \sigma_Y(\theta_{YX} a, \theta_Y Z)$ and $\sigma_Y(W, Z) = \sigma_Y(\theta_Y W, \theta_Y Z)$. We write $\langle a, b:Y \rangle_X$ for the Gromov product, $\langle \chi_X a, \chi_X b: \theta_X Y \rangle$ in $(\mathcal{H}(X), \sigma_X)$.

We list the following properties (most of which we have already verified in the context of the present paper).

- (A1): $(\exists k \geq 0)(\forall X \in \mathcal{X}) \mathcal{H}(X)$ is k -hyperbolic.
- (A2): $(\exists k_1, k_2 \geq 0)(\forall X \in \mathcal{X})(\forall a, b \in \mathcal{R}(X)) \sigma_X(a, b) \leq k_1 \rho_X(a, b) + k_2$.
- (A3): $(\exists k_1, k_2 \geq 0)(\forall X \in \mathcal{X})(\forall Y \in \mathcal{X}(X))(\forall a, b \in \mathcal{R}(X)) \rho_Y(a, b) \leq k_1 \rho_X(a, b) + k_2$.
- (A4): There is some $t \geq 0$ such that if $X, Y, Z \in \mathcal{X}$ with $Z \prec Y \prec X$ and $a \in \mathcal{R}(X)$, then $\rho_Z(a, \circ\psi_{YX}a) \leq t$.
- (A5): $(\exists t \geq 0)(\forall X \in \mathcal{X})$ if $Y, Z \in \mathcal{X}$ with $Y \wedge Z$ or $Y \prec Z$, then $\sigma_X(Y, Z) \leq t$, whenever this is defined.
- (A6): $(\exists r \geq 0)(\forall X \in \mathcal{X})(\forall a, b \in \mathcal{R}(X))$ the set of $Y \in \mathcal{X}(X)$ with $\rho_Y(a, b) \geq r$ is finite.
- (A7): $(\forall r \geq 0)(\exists r' \geq 0)(\forall X \in \mathcal{X})(\forall a, b \in \mathcal{R}(X))$ if $\sigma_Y(a, b) \leq r$ for all $Y \in \mathcal{X}(X)$, then $\rho_X(a, b) \leq r'$.
- (A8): $(\exists r \geq 0)(\forall X \in \mathcal{X})(\forall Y \in \mathcal{X}(X))(\forall a, b \in \mathcal{R}(X))$ if $\langle a, b: Y \rangle_X \geq r$ then $\sigma_Y(a, b) \leq r$.
- (A9): $(\exists r \geq 0)(\forall X \in \mathcal{X})(\forall Y, Z \in \mathcal{X}(X))$ if $Y \pitchfork Z$ and $a \in \mathcal{R}(X)$ then $\min\{\sigma_Y(a, Z), \sigma_Z(a, Y)\} \leq r$, and if $W \in \mathcal{X}(X)$ with $W \pitchfork Y$ and $W \pitchfork Z$, then $\min\{\sigma_Y(W, Z), \sigma_Z(W, Y)\} \leq r$.
- (A10): $(\exists r \geq 0)(\forall X \in \mathcal{X})$ if $\mathcal{Y} \subseteq \mathcal{X}(X)$ with $Y \wedge Z$ for all distinct $Y, Z \in \mathcal{Y}$, and if to each $Y \in \mathcal{Y}$ we have associated some $a_Y \in \mathcal{R}(X)$, then there is some $a \in \mathcal{R}(Y)$ with $\rho_Y(a_Y, \psi_{YX}a) \leq r$ and $\sigma_W(a, Y) \leq r$ for all $Y \in \mathcal{Y}$ and all $W \in \mathcal{X}(X)$ with $W \pitchfork Y$.

In [Bo4] is shown that the properties (A1)–(A10) allow us to define a ternary operation, μ_X , on each space $\mathcal{R}(X)$ to give it the structure of a coarse median space, and such that the maps χ_X and ψ_{YX} are all coarse median homomorphisms. A more precise statement of the conclusion is given as Theorem 4.7 below.

First we note that all these properties hold here. Properties (A1)–(A4) we have already observed. Property (A5) is an elementary general property of subsurface projection (given that in this case, $X \in \mathcal{X}_N$). Properties (A6), (A7) and (A8) are respectively Lemmas 4.2, 4.3 and 4.5 here. The first assertion of (A9) is Lemma 4.4 here, and the second is just the standard form of Behrstock’s lemma. It remains to prove (A10). We may as well assume that $X = \Sigma$ (since it can be interpreted as a statement intrinsic to X).

Lemma 4.6. *There is some $r_0 \geq 0$, depending only on $\xi(\Sigma)$ with the following property. Suppose that $\mathcal{Y} \subseteq \mathcal{X}$ is a collection of pairwise disjoint subsurfaces, and to each $Y \in \mathcal{Y}$, we have associated some element, $a_Y \in \mathcal{R}(Y)$. Then there is some $a \in \mathcal{R}(\Sigma)$ with $\rho(a_Y, \psi_Y a) \leq r_0$ and with $\sigma_W(a, Y) \leq r_0$ for all $W \in \mathcal{X}$ satisfying $W \pitchfork X$.*

Proof. Despite the somewhat technical statement, this is really an elementary observation about combining decorated markings on disjoint subsurfaces. It was verified in [Bo4] that the corresponding statement holds in $\mathcal{M}(\Sigma)$. In particular, we can find a marking $\bar{a} \in \text{Map}(\Sigma)$ with $\iota(\bar{a}, \tau)$ bounded, and with $\psi_Y^\diamond \bar{a} \sim \bar{a}_Y$ for all $Y \in \mathcal{Y}$. Here, τ is the union of all relative boundary components of elements of \mathcal{Y} . In fact, we can assume that \bar{a} contains each \bar{a}_Y for $Y \in \mathcal{Y} \cap \mathcal{X}_N$ as well as the core curve of each element of $\mathcal{Y} \cap \mathcal{X}_A$. Note that $\bigcup_Y \hat{a}_Y$ together with all these core curves form a multicurve in Σ , and we assign the prescribed decorations to its components. We set all other decorations equal to 0. It is now easily verified that the resulting decorated marking, a , has the required properties. \square

We have now verified the hypotheses of Theorem 1.4 of [Bo4]. We deduce:

Theorem 4.7. *There is a ternary operation, μ_X , defined on each space $\mathcal{R}(X)$ such that $(\mathcal{R}(X), \rho_X, \mu_X)$ is a coarse median space, and such that the maps $\theta_{YX} : \mathcal{R}(X) \rightarrow \mathcal{H}(Y)$ for $Y \preceq X$ are all median quasimorphisms. The median μ_X is unique with this property, up to bounded distance. The maps $\psi_{YX} : \mathcal{R}(X) \rightarrow \mathcal{R}(Y)$ for $Y \preceq X$ are also median quasimorphisms. The coarse median space $(\mathcal{R}(X), \rho_X, \mu_X)$ is finitely colourable, and has rank at most $\xi(X)$. Moreover, all bound depend only on $\xi(\Sigma)$.*

In particular, this implies Theorem 1.1.

For future reference (see Section 7), we note the following variation, due to Rafi, the distance formula of Masur and Minsky, mentioned in Section 2.

Given $a, b \in \mathcal{R}(\Sigma)$ and $r \geq 0$, let $\mathcal{A}(a, b; r) = \{X \in \mathcal{X} \mid \sigma_X(a, b) > r\}$. We have noted that this is finite. We have:

Proposition 4.8. *There is some $r_0 \geq 0$ depending only on $\xi(\Sigma)$ such that for all $r \geq r_0$, then for all $a, b \in \mathcal{R}(\Sigma)$, $\rho(a, b) \asymp \sum_{X \in \mathcal{A}(a, b; r)} \sigma_X(a, b)$.*

The corresponding statement for Teichmüller space is proven in [Ra]. A direct proof for the augmented marking complex is given in [D1]. Given that these spaces are quasi-isometric [Ra, D1] these statements are equivalent, and both are equivalent to the corresponding statement for the decorated marking complex as we have described it. We note that another proof of this distance estimate can be found in [BeHS2].

We can apply these results to the extended asymptotic cone.

We write $\mathcal{R}^\infty(\Sigma) \subseteq \mathcal{R}^*(\Sigma)$ for the asymptotic cone and extended asymptotic cone of $\mathcal{R}(\Sigma)$. We similarly have spaces $\mathcal{G}^\infty(\Sigma) \subseteq \mathcal{G}^*(\Sigma)$ etc.

For the remainder of this section, we will refer to an element of the ultraproduct, $\mathcal{U}\mathcal{G}^0(\Sigma)$, as a *curve* and to an element of $\mathcal{G}(\Sigma)$ as a *standard curve*. We will similarly refer to elements of $\mathcal{U}\mathcal{X}$ and \mathcal{X} as *subsurfaces* and *standard subsurfaces* respectively. We apply similar terminology to multicurves etc.

Note that $\mathcal{U}\text{Map}(\Sigma)$ acts by isometry of $\mathcal{R}^*(\Sigma)$, and $\mathcal{U}^0\text{Map}(\Sigma)$ acts on $\mathcal{R}^\infty(\Sigma)$. Unlike the case of the marking graph, however, these spaces are not homogeneous (as we will see in Section 6).

In fact, we have a $\mathcal{U}\text{Map}(\Sigma)$ -invariant *thick part*, \mathcal{R}_T^* , of \mathcal{R}^* (that is, the ultra-limit of the set $\mathcal{R}_T \subseteq \mathcal{R}$). We say that a component, $\mathcal{R}^\infty(\Sigma)$, of $\mathcal{R}^*(\Sigma)$ is *thick* if $\mathcal{R}^\infty(\Sigma) \cap \mathcal{R}_T^*(\Sigma) \neq \emptyset$, in which case, we denote $\mathcal{R}^\infty(\Sigma) \cap \mathcal{R}_T^*(\Sigma)$ by $\mathcal{R}_T^\infty(\Sigma)$. In fact, $\mathcal{U}\text{Map}(\Sigma)$ acts transitively on $\mathcal{R}_T^*(\Sigma)$. Note that $\mathcal{R}_T^*(\Sigma)$ contains the standard basepoint of $\mathcal{R}^*(\Sigma)$. It follows that every thick component of $\mathcal{R}^*(\Sigma)$ is a $\mathcal{U}\text{Map}(\Sigma)$ -image of the standard component.

The various coarsely lipschitz quasimorphisms we described in Section 4 now give rise to lipschitz median homomorphisms. Specifically, we have maps: $\theta_X^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{H}^*(X)$, and $\psi_X^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}^*(X)$, for all $X \in \mathcal{UX}$.

If $\gamma \in \mathcal{G}^0$ is a standard curve in Σ , we write $\mathcal{H} = \mathcal{H}(\gamma) = \mathcal{H}(X(\gamma))$. Recall that $\mathcal{H}_T^0 = \{a \in \mathcal{H}^0 \mid \eta_a(\gamma) = 0\}$ (where $\eta_a(\gamma)$ is the second coordinate of a) and that $\mathcal{H}_T \subseteq \mathcal{H}$ is the complete subgraph on \mathcal{H}_T . We define $h_\gamma : \mathcal{H} \rightarrow [0, \infty)$ by $h_\gamma(a) = \rho(a, \mathcal{H}_T)$. Thus, if $a \in \mathcal{H}^0$, then $h_\gamma(a) = \eta_a(\gamma)$.

We have noted that $\mathcal{H}(\gamma)$ is quasi-isometric to a horodisc in the hyperbolic plane, \mathbb{H}^2 ; via a quasi-isometry which sends \mathcal{H}_T to the boundary horocycle. The map h_γ then corresponds to a horofunction.

Now the extended asymptotic cone, \mathcal{H}^* , is an \mathbb{R}^* -tree. The map h_γ gives rise to a 1-lipschitz map $h_\gamma^* : \mathcal{H}^* \rightarrow \mathbb{R}^*$, which $h_\gamma^*(x) \geq 0$ for all $x \in \mathcal{H}^*$.

Let \mathcal{H}^∞ be a component of \mathcal{H}^* . We write ρ^∞ for the metric on \mathcal{H}^∞ . Thus, $(\mathcal{H}^\infty, \rho^\infty)$ is a complete \mathbb{R} -tree. The map $h_\gamma^*|_{\mathcal{H}^\infty}$ is a Busemann function on \mathcal{H}^∞ . That is, it is 1-lipschitz, and for all $x \in \mathcal{H}^\infty$ and all $t \in [0, \infty)$ there is a unique $y \in \mathcal{H}^\infty$ satisfying $h_\gamma^*(y) - h_\gamma^*(x) = \rho^\infty(x, y) = t$. (Note that $h_\gamma^*(y) - h_\gamma^*(x) \in \mathbb{R}$ for all $x, y \in \mathcal{H}^\infty$.) Writing $y = x_t$, the map $[t \mapsto x_t]$ gives a flow on \mathcal{H}^∞ for $t \in [0, \infty) \subseteq \mathbb{R}$. (Such a flow will converge on an ideal point of \mathcal{H}^∞ .)

Up to isomorphism, there are two possibilities for \mathcal{H}^∞ . The first is the “thick” case, where $h_\gamma^*(\mathcal{H}^\infty) \subseteq [0, \infty)$. Write $\mathcal{H}_T^\infty = (h_\gamma^*)^{-1}(0)$. In this case, $\mathcal{H}_T^\infty \neq \emptyset$. Every point of \mathcal{H}_T^∞ is an extreme point (has valence 1), and each point of $\mathcal{H}^\infty \setminus \mathcal{H}_T^\infty$ has valence 2^{\aleph_0} .

The second is the “thin” case. Here, $h_\gamma^*(\mathcal{H}^\infty) \cap \mathbb{R} = \emptyset$, and \mathcal{H}^∞ is the complete 2^{\aleph_0} -regular tree. (Note the Busemann cocycle $[(x, y) \mapsto h_\gamma^*(x) - h_\gamma^*(y)]$ still takes real values.)

Note that, in both cases, \mathcal{H}^∞ is almost furry (i.e. there are no points of valence 2). This will allow us to apply Proposition 4.6 (see the proof of Lemma 8.2).

By a *branch* of an \mathbb{R} -tree we mean a closed subset with exactly one point in its topological boundary. A branch is necessarily a subtree. In the above, if \mathcal{H}^∞ is thick, then every branch of \mathcal{H}^∞ intersects \mathcal{H}_T^∞ . If \mathcal{H}^∞ is thin, then every branch of \mathcal{H}^∞ contains points, y , with $h_\gamma^*(x) - h_\gamma^*(y)$ an arbitrarily large real number, where $x \in \mathcal{H}^\infty$ is an arbitrary basepoint.

We will later refer again to the special case where $\Sigma = S_{1,1}$ or $\Sigma = S_{0,4}$. In this case, we saw that $\mathcal{R}(\Sigma)$ is quasi-isometric to the hyperbolic plane. Thus, $\mathcal{R}^*(\Sigma)$ is a complete \mathbb{R}^* -tree. Any component $\mathcal{R}^\infty(\Sigma)$ is the complete 2^{\aleph_0} -regular tree.

We write $\mathcal{R}_T^\infty = \mathcal{R}_T^* \cap \mathcal{R}^\infty(\Sigma)$. Here, we just note that if $\mathcal{R}_T^\infty \neq \emptyset$, then every branch of \mathcal{R}^∞ meets \mathcal{R}_T^∞ .

5. SOME APPLICATIONS OF THE COARSE MEDIAN PROPERTY

In this section, we describe a few immediate consequences of what we have shown. In particular, we will give proofs of Theorems 1.2, 1.3, and 1.8, and the “if” part of Theorem 1.4. Theorems 1.2, 1.8 and the “only if” part of 1.3 are simple consequences of the preceding discussion, so we discuss these first.

The first result is a coarse isoperimetric inequality. This is a quasi-isometrically invariant property which one could equivalently formulate in a number of different ways. For definiteness we will say that a geodesic space, Λ , satisfies a *coarse quadratic isoperimetric inequality* if the following holds. There is some $r \geq 0$ and some $k \geq 0$ such that if γ is any closed path in Λ of length at most nr , where $n \geq 1$ is some natural number, then we can find a triangulation of the disc with at most kn^2 2-simplices and a map of its 1-skeleton into Λ such that the image of every 1-cell has length at most r and such that the map restricted to the boundary agrees with γ (thought of as a map of the boundary of the disc into Λ). One can easily check that this property is quasi-isometry invariant.

It is shown in [Bo1] (Proposition 8.2) that any coarse median space has this property. We immediately deduce from Theorem 1.1 that:

Proposition 5.1. *$\mathcal{R}(\Sigma)$ has a coarse quadratic isoperimetric inequality.*

This is, of course, equivalent to Theorem 1.2.

In fact, one can give versions which fit more naturally with the usual riemannian notion. For example:

Proposition 5.2. *There is some $\lambda > 0$, depending only on the topological type of Σ , with the following property. Suppose that γ is a riemannian circle of length l , and that $f : \gamma \rightarrow \mathbb{T}(\Sigma)$ is a 1-lipschitz map. Then, we can identify γ with the boundary of a riemannian disc, D , inducing the same riemannian metric on γ , and of area at most λl^2 , so that f extends to a 1-lipschitz map, $f : D \rightarrow \mathbb{T}(\Sigma)$.*

In fact, one can also arrange that D has bounded curvature and injectivity radius bounded below (in terms of Σ).

We suspect that λ can be made independent of Σ , but will not address that issue here.

Proof. Given the coarse quadratic inequality (Theorem 1.2) it is enough to show that the statement holds for some area bound, that is, where the bound λl^2 is replaced by some function of l which assumed to be $O(l^2)$ only on a small scale.

To see this holds, note first that $\mathbb{T}(\Sigma)$ satisfies a (linear) isodiametric inequality. In fact, we can cone any curve over any point using Teichmüller geodesics. Moreover, it is known that $\mathbb{T}(\Sigma)$ has locally bounded geometry. Specifically, Theorem 8.2 of [Mc] shows that the Teichmüller metric is bilipschitz equivalent to a

riemannian (Kähler) metric, with curvatures bounded above and below, and with injectivity radius bounded below. The set of such metrics possible on a set of bounded diameter is precompact. Now standard precompactness arguments in riemannian geometry tell us that our curve must bound a disc whose area is bounded above by some function of its length which is quadratic on a small scale. \square

We next observe that Theorem 1.8 also follows. In fact, we have already made the relevant observations in Sections 1 and 2.4. Given that $\mathcal{R}(\Sigma)$ is a finitely colourable coarse median space, it follows that $\mathcal{R}^*(\Sigma)$ with the limiting metric and ternary operation is a median algebra satisfying the hypotheses of Lemma 2.5. It follows that it is bilipschitz equivalent to a median metric space. (For more details, see [Bo2].) It in turn follows that it is bilipschitz equivalent to a CAT(0) metric [Bo3]. Note that, in view of Theorem 4.1, this applies in particular to any asymptotic cone of $\mathcal{R}(\Sigma)$, hence also any asymptotic cone of $\mathbb{T}(\Sigma)$. This proves Theorem 1.8.

Some other consequences follow on from the fact that $\mathcal{R}_T^*(\Sigma)$ is a locally convex topological median algebra of finite rank. For example, the topological dimension of any locally compact subset of any asymptotic cone of $\mathcal{R}_T^\infty(\Sigma)$ is at most $\xi(\Sigma)$ (see Theorem 2.2 and Lemma 7.6 of [Bo1]). In particular, it does not admit any continuous injective map of $\mathbb{R}^{\xi+1}$. From this we get the following. (A similar statement can be found in [EMR1].) Write B_R^n for the ball of radius r in the euclidean space \mathbb{R}^n .

Proposition 5.3. *Given parameters of quasi-isometry, there is some constant $r \geq 0$, such that there is no quasi-isometric embedding of $B_r^{\xi+1}$ into $\mathcal{R}(\Sigma)$ with these parameters.*

Proof. This is a standard argument involving asymptotic cones. Suppose that, for each $i \in \mathbb{N}$, the ball, B_i , of radius i admits a uniformly quasi-isometric embedding, $\phi_i : B_i \rightarrow \mathcal{R}(\Sigma)$. Now pass to the asymptotic cone with scaling factors, i . We end up with a bilipschitz map, $\phi^\infty : B_1 \rightarrow \mathcal{R}^\infty(\Sigma)$, contradicting the dimension bound. \square

(Indeed, the above holds in any coarse median space of rank at most ξ .)

An immediate consequence is that $\mathcal{R}(\Sigma)$ does not admit any quasi-isometric embedding of a euclidean $(\xi(\Sigma) + 1)$ -dimensional half-space. This proves the “only if” part of Theorem 1.3.

For the “if” part, we need to construct such an embedding in dimension $\xi(\Sigma)$. We use the same construction as in [EMR1], though base the proof on the arguments here. This will show, in addition, that the image can be assumed quasi-convex in the coarse median structure.

Given any $a \in \mathcal{M}^0(\Sigma)$, and any $t \subseteq a$, let $O_a(t) = \{b \in \mathcal{R}^0(\Sigma) \mid \bar{b} = a, \hat{b} \subseteq t\} \subseteq \mathcal{R}(\Sigma)$. In other words, we take all possible decorations on a subject to the constraint that all the decorated curves must lie in t .

Lemma 5.4. $O_a(t)$ is quasiconvex in $\mathcal{R}(\Sigma)$.

Proof. We define a map $\omega : \mathcal{R}(\Sigma) \rightarrow O_a(t)$ as follows. Given $x \in \mathcal{R}^0(\Sigma)$, let $\tau = \hat{x} \cap t \subseteq a$. Let $b \in \mathcal{R}^0$ be such that $\bar{b} = a$, $\eta_b|_\tau = \eta_x|_\tau$ and $\eta_b|(\bar{b} \setminus \tau) \equiv 0$, and set $\omega x = b$. In other words, we take the base marking a , and decorate curves in a if and only if they also happen to be decorated curves of x . We claim that ω is a coarse gate map.

To this end, let $c \in O_a(t)$, so $\hat{c} \subseteq t$. If $\gamma \in \tau$, then $\theta_\gamma \omega x \sim \theta_\gamma x$, so $\theta_\gamma \mu(x, \omega x, c) \sim \theta_\gamma \omega x$. If $\gamma \in \hat{c} \setminus \tau$, then $\gamma \notin \hat{x}$ (otherwise $\gamma \in \hat{c} \cap \hat{x} \subseteq t \cap \hat{x} = \tau$) and so $\theta_\gamma \omega x \sim \theta_\gamma a \sim \theta_\gamma c$ (since $\bar{c} = a$), and $\theta_\gamma \mu(x, \omega x, c) \sim \theta_\gamma \omega x$. Finally, if $X \in \mathcal{X}$ does not have the form $X(\gamma)$ for such γ , then $\theta_X c \sim \theta_X a \sim \theta_X \omega x$, so $\theta_X \mu(x, \omega x, c) \sim \theta_X \omega x$. By Lemma 4.3, $\mu(x, \omega x, c) \sim \omega x$ as claimed. \square

Given a multicurve $\tau \in \mathcal{S}$, write $\mathcal{O}(\tau) = [0, \infty)^\tau$ and $\mathcal{O}^0(\tau) = \mathbb{N}^\tau \subseteq \mathcal{O}(\tau)$. Note that $\mathcal{O}(\tau)$ is a median algebra with the product structure, and that $\mathcal{O}^0(\tau)$ is a subalgebra.

Suppose $a \in \mathcal{M}^0 \equiv \mathcal{R}_T^0$, with $\tau \subseteq a$. Define a map $\lambda = \lambda_a : \mathcal{O}^0(\tau) \rightarrow \mathcal{R}(\Sigma)$ by setting $\lambda_a(v) = b$ where $\bar{b} = \bar{a}$, $\hat{b} = \tau$, $\eta_a|_\tau = v$, and $\eta_a|(a \setminus \tau) \equiv 0$. (In other words, we take the base marking a , with decorations determined by v .) The map λ is easily seen to be a quasimorphism. Note that $O_a(\tau) = \lambda(\mathcal{O}^0(\tau))$. Also, if $x \in \mathcal{R}(\Sigma)$, then $\omega(x) = \lambda(u)$, where $u \in \mathcal{O}(\tau)$ is defined by $u|(\tau \cap \hat{x}) = \eta_x|(\tau \cap \hat{x})$ and $u|(\tau \setminus \hat{x}) \equiv 0$.

By Lemma 5.4, $O_a(\tau)$ is quasiconvex in $\mathcal{R}(\Sigma)$. Moreover, it follows that λ extends to a quasi-isometric embedding of $\mathcal{O}(\tau)$ into $\mathcal{R}(\Sigma)$.

We have shown that:

Lemma 5.5. *The map $\lambda_a : \mathcal{O}(\tau) \rightarrow \mathcal{R}(\Sigma)$ is a quasi-isometric embedding with quasiconvex image.*

If we take τ to be any complete multicurve and a to be any marking containing it, then we get a quasi-isometric embedding of a ξ -orthant (or ξ -dimensional half-space).

This proves the “if” part of Theorem 1.3.

Note that this shows that $\mathcal{R}(\Sigma)$ has coarse median rank exactly $\xi(\Sigma)$ (given the observation after Proposition 5.3).

For the “if” part of Theorem 1.4, we want to construct a quasi-isometric embedding of \mathbb{R}^ξ , in the cases described. (We will deal with the “only if” part in Section 9.)

Suppose that Υ is a finite simplicial complex. We can construct a singular euclidean space, $\mathcal{O}(\Upsilon)$, by taking an orthant for every simplex of Υ and gluing them together in the pattern determined by Υ . This has vertex o and we can identify Υ as the spherical link of o in $\mathcal{O}(\Upsilon)$. (For example, the cross polytope gives a copy of euclidean space.) Note that if Υ is bilipschitz equivalent to the standard $(n-1)$ -sphere, then $\mathcal{O}(\Upsilon)$ is bilipschitz equivalent to \mathbb{R}^n . We write $\mathcal{O}^0(\Upsilon)$ for the set of integer points in $\mathcal{O}(\Upsilon)$.

Now let $\mathcal{C}(\Sigma)$ be the curve complex associated to Σ . This is the flag complex with 1-skeleton $\mathcal{G}(\Sigma)$. In particular, $\mathcal{C}^0(\Sigma) = \mathcal{G}^0(\Sigma)$. We can identify the set of simplices of $\mathcal{C}(\Sigma)$ with the set, $\mathcal{S} \setminus \{\emptyset\}$, of non-empty multicurves in Σ .

Given a finite subcomplex, Υ , of $\mathcal{C}(\Sigma)$, write $\mathcal{S}(\Upsilon) \subseteq \mathcal{S}$ for the set of multicurves corresponding to the simplices of Υ . Write $\Upsilon^0 = \bigcup \mathcal{S}(\Upsilon)$ for the set of vertices. If $\tau \in \mathcal{S}$, we can identify $\mathcal{O}(\tau)$ as a subset of $\mathcal{O}(\Upsilon)$, and $\mathcal{O}^0(\tau)$ as a subset of $\mathcal{O}^0(\Upsilon)$.

Suppose $a \in \mathcal{M}^0$ is marking of Σ with $\Upsilon^0 \subseteq a$. We can define a map $\lambda = \lambda_a : \mathcal{O}^0(\Upsilon) \rightarrow \mathcal{R}(\Sigma)$ by combining the maps $\lambda_a : \mathcal{O}^0(\tau) \rightarrow \mathcal{R}(\Sigma)$ for $\tau \in \mathcal{S}$. Write $O(\Upsilon) = \lambda(\mathcal{O}(\Upsilon)) = \bigcup_{\tau \in \mathcal{S}(\Upsilon)} O_a(\tau)$.

Suppose now that Υ is a full subcomplex of $\mathcal{C}(\Sigma)$ (in other words, if the vertices of a simplex in $\mathcal{C}(\Sigma)$ are contained in Υ^0 , then the whole simplex is contained in Υ). In this case, we have $O(\Upsilon) = O(\Upsilon^0)$ as previously defined. In particular, Lemma 5.4 tells us that $O(\Upsilon)$ is coarsely convex.

As in the case of a single orthant, we now see:

Lemma 5.6. *If Υ is a full subcomplex of $\mathcal{C}(\Sigma)$, with $\Upsilon^0 \subseteq a \subseteq \mathcal{M}^0$, then $O(\Upsilon)$ is quasiconvex in $\mathcal{R}(\Sigma)$.*

We see that λ_a extends to a quasi-isometric embedding of $\mathcal{O}(\Upsilon)$ into $\mathcal{R}(\Sigma)$. (In fact, one can show that λ_a is a quasi-isometric embedding even if we do not assume that Υ is full, though of course, its image need not be quasiconvex in this case.)

Note that if Υ is PL homeomorphic to the standard $(n - 1)$ -sphere, we get a bilipschitz embedding of \mathbb{R}^n into $\mathcal{R}(\Sigma)$.

We will show:

Proposition 5.7. *$\mathcal{C}(\Sigma)$ contains a full subcomplex PL homeomorphic to the standard $(\xi(\Sigma) - 1)$ -sphere if and only if Σ has genus at most 1, or is the closed surface of genus 2.*

This is based on results and constructions in [Har]. (For further related discussion, see [Br].) If $\Sigma \cong S_{g,p}$, write $\xi'(\Sigma) = 2g + p - 2$ if $g, p > 0$, $\xi'(\Sigma) = 2g - 1$ if $p = 0$ and $\xi'(\Sigma) = p - 3$ if $g = 0$. (Here, we are assuming that $\xi(\Sigma) \geq 2$.)

Note that $\xi'(\Sigma) \leq \xi(\Sigma)$ with equality precisely in the cases described by Proposition 5.7.

It is shown in [Har] (Theorem 3.5) that $\mathcal{C}(\Sigma)$ is homotopy equivalent to a wedge of spheres of dimension $\xi'(\Sigma) - 1$. In particular, the homology is trivial in dimension $\xi'(\Sigma)$. (It does not matter which homology theory we use here.) Now $\mathcal{C}(\Sigma)$ has dimension $\xi(\Sigma)$, so if $\xi'(\Sigma) < \xi(\Sigma)$, it follows that $\mathcal{C}(\Sigma)$ cannot contain any $(\xi(\Sigma) - 1)$ -dimensional homology cycle, and so in particular, no topologically embedded closed $(\xi(\Sigma) - 1)$ -manifold. This proves the “only if” part of Proposition 5.7 (indeed without the “full” requirement).

For the “if” part, we need to construct such a sphere. In the planar (genus-0) case there is a simple explicit construction described in [AL], which involves doubling the arc complex of a disc. The latter is known to be homeomorphic to a

sphere, see [S]. (Another proof of this is given in [HuM]. Although not explicitly stated, it is easily checked that this gives a PL sphere.) However, it is unclear how to adapt this to the genus-1 case. Below, we give an argument which deals with all cases. It is based on an idea in Harer's proof of the result mentioned above. We first show:

Lemma 5.8. *Suppose $g, p, n \in \mathbb{N}$ and $p \geq 2$. Suppose that $\mathcal{C}(S_{g,p})$ contains a full subcomplex PL homeomorphic to the standard n -sphere. Then $\mathcal{C}(S_{g,p+1})$ contains a full subcomplex PL homeomorphic to the standard $(n+1)$ -sphere.*

Proof. For this discussion, it will be convenient to view $S_{g,p}$ as a closed surface, S , of genus g , with a set $A \subseteq S$ of p preferred points. Let $\Upsilon \subseteq \mathcal{C}(S_{g,p})$ be a full subcomplex PL homeomorphic to the standard n -sphere. Now realise the elements of Υ^0 as closed curves in $S \setminus A$, so that no three curves intersect in a point, and such that the total number of intersections is minimal, subject to this constraint. (It is well known that this necessarily minimises pairwise intersection numbers.)

Now let $I \subseteq S$ be an embedded arc meeting A precisely at its endpoints, a, b , say. We may assume that no point of I lies in two curves of Υ^0 and that (subject to this constraint) the total number, m , of intersections, $I \cap \bigcup \Upsilon^0 \subseteq \Sigma$, is minimal, in the homotopy class of I in $S \setminus A$ relative to its endpoints. Let I_0, I_1, \dots, I_m , be the components of $I \setminus (\{a, b\} \cup \bigcup \Upsilon^0)$, consecutively ordered, so that a and b lie respectively in the closures of I_0 and I_m .

Choose any point $c_i \in I_i$ and an arbitrary point $d \in I \setminus \{a, b\}$. Let $B = A \cup \{d\}$, and think of $S_{g,p+1}$ as S with the points of B removed.

The following can be thought of intuitively as sliding the point d from a to b along I . However, formally it is better expressed as keeping d fixed and applying an isotopy to the curves, as we now describe.

Given any $i \in \{0, \dots, m\}$, we obtain a map $f_i : \Upsilon \rightarrow \mathcal{C}(S_{g,p+1})$ as follows. Take an isotopy of S supported on a small neighbourhood of the interval $[d, c_i]$, fixing I setwise, and carrying c_i to d . At the end of the isotopy we get a map sending each curve in Υ^0 to a curve in $S \setminus B$, and so gives rise to a map $f_i : \Upsilon^0 \rightarrow \mathcal{C}(S_{g,p+1})$. Note that postcomposing with the map which forgets the point d , we get the inclusion of Υ^0 into $\mathcal{C}(S_{g,p})$. Now it is easily seen that f_i preserves disjointness of curves, and so extends to a map $f_i : \Upsilon \rightarrow \mathcal{C}(S_{g,p+1})$, which maps Υ isomorphically to a full subcomplex $\Upsilon_i = f_i(\Upsilon) \subseteq \mathcal{C}(S_{g,p+1})$. Now let α and β be, respectively, the boundary curves of small regular neighbourhoods of $[a, d]$ and $[d, b]$ in I . Let $\Omega^0 = \{\alpha, \beta\} \cup \bigcup_{i=0}^m \Upsilon_i^0$, and let Ω be the full subcomplex of $\mathcal{C}(\Sigma)$ with vertex set Ω_0 . We claim that Ω is PL homeomorphic to the standard $(n+1)$ -sphere.

Note first that Υ_0^0 and Υ_m^0 are respectively the sets of points adjacent to α and β .

Now, given $i \in \{0, \dots, m\}$, let $\Omega_i^0 = \{\alpha\} \cup \bigcup_{j=0}^i \Upsilon_j^0$, and let Ω_i be the full subcomplex with vertex set Ω_i^0 . Now Ω_0 is the cone on Υ_0 with vertex α , and so PL homeomorphic to a ball with boundary Υ_0 . We claim that, for all i ,

there is a PL homeomorphism of Ω_0 to Ω_i whose restriction to Υ_0 is the map $f_i \circ f_0^{-1} : \Upsilon_0 \longrightarrow \Upsilon_m$.

Suppose, inductively, that this holds for i . Moving from Υ_i to Υ_{i+1} corresponds to pushing one of the curves of Υ_i^0 across the hole, d , of $S_{g,p+1}$. In other words, there is some $\gamma \in \Upsilon^0$ such that $f_i|(\Upsilon^0 \setminus \{\gamma\}) = f_{i+1}|(\Upsilon^0 \setminus \{\gamma\})$, but with $f_i(\gamma) \neq f_{i+1}(\gamma)$. Let $\delta = f_i(\gamma)$ and $\epsilon = f_{i+1}(\gamma)$. Thus, $\Omega_{i+1}^0 = \Omega_i^0 \cup \{\epsilon\}$. Now δ and ϵ are clearly adjacent. In fact, it is easily checked that the set of curves in Ω_i adjacent to ϵ are precisely those of the form $f_i(\zeta)$, where $\zeta \in \Upsilon^0$ is equal to or adjacent to γ . Thus, Ω_{i+1} is obtained from Ω_i by attaching a cone with vertex ϵ to the star of δ in Υ_i . Given that Ω_i is a PL $(n+1)$ -ball with boundary Υ_i , we can find a PL homeomorphism of Ω_i to Ω_{i+1} which restricts to $f_{i+1} \circ f_i^{-1}$ on Υ_i . Precomposing this with the homeomorphism from Ω_0 to Ω_i proves the inductive step of the statement.

In particular, we have a PL homeomorphism from Ω_0 to Ω_m whose restriction to Υ_0 is $f_m \circ f_0^{-1}$. Now Ω is obtained from Ω_m by coning the boundary, Υ_m , with vertex β . Thus, Ω is PL homeomorphic to a suspension of Υ , hence a PL $(n+1)$ -sphere. \square

Now it is easy to find a pentagon which is a full subcomplex of $\mathcal{C}(S_{0,5})$ or equivalently in $\mathcal{C}(S_{1,2})$. Note that these cases correspond to $\xi = \xi' = 2$. Therefore, applying Lemma 5.8 inductively, we find a $(\xi - 1)$ -sphere for all surfaces of genus at most 1, where $\xi \geq 2$. Note that $\mathcal{C}(S_{2,0}) \cong \mathcal{C}(S_{0,6})$ so this deals with that case also. (An explicit description of a sphere in the case of $S_{2,0}$ can be found in [Br].)

This completes the proof Proposition 5.7.

To construct our quasi-isometric embedding of \mathbb{R}^ξ into $\mathcal{R}(\Sigma)$ in these cases, we could assume that, in defining the marking complex, we have taken the intersection bounds large enough so that some marking a contains all the vertices of a $(\xi - 1)$ -sphere, Υ , in $\mathcal{C}(\Sigma)$. Now construct $\lambda_a : \mathcal{O}(\Upsilon) \longrightarrow \mathcal{R}(\Sigma)$ as above and apply Lemma 5.6. (Alternatively, we could take different marking containing each multicurve of $\mathcal{S}(\Upsilon)$. Note that these markings are all a bounded distance apart. Thus, if $\tau' \subseteq \tau \in \mathcal{S}(\Upsilon)$, then the map of $\mathcal{O}(\tau')$ into $\mathcal{R}(\Sigma)$ agrees up to bounded distance with the restriction of the map of $\mathcal{O}(\tau)$ into $\mathcal{R}(\Sigma)$. Therefore, these maps again combine to give a quasi-isometric embedding of $\mathcal{O}(\Upsilon)$ into $\mathcal{R}(\Sigma)$.)

This proves the “if” part of Theorem 1.4.

6. PRODUCT STRUCTURE AND STRATIFICATION

In this section, we show how the extended asymptotic cone, $\mathcal{R}^*(\Sigma)$ can be partitioned into “strata” indexed by (ultralimits of) multicurves. The strata have a local product structure, where the factors correspond to extended asymptotic cones of subsurfaces. This structure arises from a kind of coarse stratification of $\mathcal{R}(\Sigma)$. The basic idea is that when there is a set of curves in a decorated marking which have large decorations then $\mathcal{R}(\Sigma)$ locally splits a direct product. (The analogous statement for Teichmüller space is that when we have a set of very short

curves on a hyperbolic surface, then the Teichmüller space in a neighbourhood of that surface splits coarsely as a product of the Teichmüller spaces of the complementary pieces, together with hyperbolic discs corresponding to twisting and shrinking these curves.)

Properties of this stratification will be used in Sections 7 and 8. We begin by describing a coarse product structure based on multicurves.

Recall that \mathcal{S} is the set of all multicurves in Σ . We allow $\emptyset \in \mathcal{S}$. If $\tau \in \mathcal{S}$, we will use the following notation (as in [Bo4]). We write $\mathcal{X}_A(\tau) = \{X(\gamma) \mid \gamma \in \tau\}$. We write $\mathcal{X}_N(\tau)$ for the set of components of $\Sigma \setminus \tau$ which are not $S_{0,3}$'s, and set $\mathcal{X}(\tau) = \mathcal{X}_N(\tau) \cup \mathcal{X}_A(\tau)$. We write $\mathcal{X}_T(\tau) = \{Y \in \mathcal{X} \mid Y \pitchfork \tau\}$; that is, there is some $\gamma \in \tau$ with $\gamma \pitchfork Y$ or $\gamma \prec Y$.

Let τ be a (possibly empty) multicurve. Let $L(\tau)$ be the set of $a \in \mathcal{R}^0(\Sigma)$ such that $\tau \subseteq \bar{a}$ and τ does not cross \hat{a} (so that $\tau \cup \hat{a}$ is a multicurve). Thus $a \in L(\hat{a})$ for all $a \in \mathcal{R}^0(\Sigma)$. Note that if $\tau \subseteq \tau'$, then $L(\tau') \subseteq L(\tau)$.

Given $r \geq 0$, let $L(\tau; r) = \{a \in \mathcal{R}(\Sigma) \mid \iota(\bar{a}, \tau) \leq r\}$. One can check that for all sufficiently large r (in relation to $\xi(\Sigma)$), the Hausdorff distance between $L(\tau)$ and $L(\tau; r)$ is bounded. As in Lemma 9.1 of [Bo4], we see that $a \in \mathcal{R}(\Sigma)$ is a bounded distance from $L(\tau)$ if and only if $\sigma_Y(\theta_Y a, \theta_Y \tau)$ is bounded for all $Y \in \mathcal{X}_T(\tau)$.

Let $\mathcal{L}(\tau) = \prod_{X \in \mathcal{X}(\tau)} \mathcal{R}(X)$. We give this the l^1 metric (though any quasi-isometrically equivalent geodesic metric would serve for our purposes). Note that this has a product coarse median structure. Combining the maps $\psi_X : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(X)$, we get a coarsely lipschitz quasimorphism $\psi_\tau : \mathcal{R}(\Sigma) \rightarrow \mathcal{L}(\tau)$.

In the other direction, Lemma 4.6 gives us a map, $v_\tau : \mathcal{L}(\tau) \rightarrow \mathcal{R}(\Sigma)$, so that $\psi_Y \circ v_\tau$ gives us a prescribed element of $\mathcal{R}(Y)$ for each coordinate $Y \in \mathcal{X}(\tau)$. In other words, $\psi_\tau \circ v_\tau$ is the identity up to bounded distance. Note that v_τ is necessarily a quasimorphism. Note that, by construction, $v_\tau(\mathcal{L}(\tau)) \subseteq L(\tau)$.

We now set $\omega_\tau = v_\tau \circ \psi_\tau : \mathcal{R}(\Sigma) \rightarrow L(\tau)$. It is now an immediate consequence of Lemma 9.3 of [Bo4] that this is a coarse gate map to the set $L(\tau)$; that is, $\rho(\omega_\tau x, \mu(x, \omega_\tau x, c))$ is bounded for all $x \in \mathcal{R}(\Sigma)$ and all $c \in L(\tau)$.

Note that $\rho(a, \mathcal{R}_T(\Sigma))$ is the sum of the decorations in a . In other words, in the notation of Section 3, we have $\rho(a, \mathcal{R}_T^*(\Sigma)) = \sum_{\gamma \in \hat{a}} \eta_a(\gamma) = \sum_{\gamma \in \hat{a}} h_\gamma(a)$.

We note that if $a \in \mathcal{L}(\tau)$, then we can write $h(v_\tau a) = \sum_{X \in \mathcal{X}(\tau)} h_X(a)$, where $h(a) = \rho(a, \mathcal{R}_T(\Sigma))$ (as defined above), and where h_X is the corresponding function defined intrinsically on each of the factors, $\mathcal{R}(X)$, of $\mathcal{L}(X)$. This comes directly from the construction of v_τ . Recall that this combines the decorations on each of the factors: every decorated curve lies in exactly one such X . Moreover, for any decorated marking, a , $h(a)$ is exactly the sum of the decorations on a . This applies in $\mathcal{R}(\Sigma)$, and each of the factors, $\mathcal{R}(X)$.

We now move on to the extended asymptotic cone, $\mathcal{R}^*(\Sigma)$. Recall that we have maps: $\theta_X^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{G}^*(X)$ and $\psi_X^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}^*(X)$ for all $X \in \mathcal{UX}$.

If τ is a multicurve, we have $\psi_\tau^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{L}^*(\tau)$ and $v_\tau^* : \mathcal{L}^*(\tau) \rightarrow \mathcal{L}^*(\Sigma)$ with $\psi_\tau^* \circ v_\tau^*$ the identity. We write $\omega_\tau^* = v_\tau^* \circ \psi_\tau^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}^*(\Sigma)$. Note that $\mathcal{L}^*(\tau)$ is a direct product of the spaces $\mathcal{R}^*(X)$ as X varies in $\mathcal{UX}(\tau)$.

Write $L^*(\tau) = v_\tau^*(\mathcal{L}^*(\tau)) = \omega_\tau(\mathcal{R}^*(\Sigma))$. This is also the limit of the sets $(L(\tau_\zeta))_\zeta$. Note that $L^*(\tau)$ is median convex in $\mathcal{R}^*(\Sigma)$, and that $\omega_\tau : \mathcal{R}^*(\Sigma) \rightarrow L^*(\tau)$ is the gate map (that is $\omega_\tau^*(x) \in [x, c]$ for all $x \in \mathcal{R}^*(\Sigma)$ and all $c \in L^*(\tau)$). In particular, $\omega_\tau^*|_{L^*(\tau)}$ is the identity. Note that if $\tau \subseteq \tau'$, then $L^*(\tau') \subseteq L^*(\tau)$.

Note that if τ is big (that is, all complementary components have complexity at most 1), then $L^\infty(\tau)$ is the direct product of ξ almost furry \mathbb{R} -trees. (Since in this case each factor of $L(\tau_\zeta)$ is quasi-isometric to either a hyperbolic plane or a horodisc.)

We now move on to describe the stratification.

Let \mathcal{S} be the set of standard multicurves in Σ . We allow $\emptyset \in \mathcal{S}$. Given $\tau \in \mathcal{S}$, let $\Theta(\tau) = \{a \in \mathcal{R}(\Sigma) \mid \hat{a} = \tau\}$. Thus $\Theta(\tau) \subseteq L(\tau)$, and $\Theta(\emptyset) = \mathcal{R}_T(\Sigma)$.

Lemma 6.1. *Given any $\tau, \tau' \in \mathcal{S}$ and any $a \in \Theta(\tau)$ and $b \in \Theta(\tau')$, there is some $c \in \Theta(\tau \cap \tau')$ with $\mu(a, b, c) \sim c$.*

Proof. Let $d \in \mathcal{R}^0(\Sigma)$ be obtained from a by setting $\bar{d} = \bar{a}$, resetting the decorations on $\tau \setminus \tau' \subseteq \hat{a}$ equal to 0, and leaving all other decorations on a unchanged. Thus $d \in \Theta(\tau \cap \tau')$. Now apply Dehn twists to d about the curves of $\tau \setminus \tau'$ to give $c \in \mathcal{R}^0(\Sigma)$, so that $\theta_\gamma c \sim \theta_\gamma b$ for all $\gamma \in \tau \setminus \tau'$. We also have $c \in \Theta(\tau \cap \tau')$. Suppose $X \in \mathcal{X}$. If $X = X(\gamma)$ for some $\gamma \in \tau \setminus \tau'$, then $\theta_X c \sim \theta_X b$. If X is not of this form, we get $\theta_X c \sim \theta_X a$. In all cases, we have $\theta_X \mu(a, b, c) \sim \mu(\theta_X a, \theta_X b, \theta_X c) \sim \theta_X c$. It follows by Lemma 4.3 that $\mu(a, b, c) \sim c$. \square

We now pass to the asymptotic cone. Write \mathcal{US} for the ultraproduct of \mathcal{S} . Recall that we have an intersection operation defined on \mathcal{US} .

Given any $\tau \in \mathcal{US}$, let $\Theta^*(\tau)$ be the limit of $(\Theta(\tau_\zeta))_\zeta$. Note that $\Theta^*(\tau)$ is closed, and $\Theta^*(\tau) \subseteq L^*(\tau)$. Also, clearly $\mathcal{R}^*(\Sigma) = \bigcup_{\tau \in \mathcal{US}} \Theta^*(\tau)$.

Lemma 6.2. *Given any $\tau, \tau' \in \mathcal{US}$, and any $a \in \Theta^*(\tau)$, $b \in \Theta^*(\tau')$, then $[a, b] \cap \Theta^*(\tau \cap \tau') \neq \emptyset$.*

Proof. Choose $a_\zeta \in \Theta(\tau)$ and $b_\zeta \in \Theta(\tau')$ with $a_\zeta \rightarrow a$ and $b_\zeta \rightarrow b$. Let $c_\zeta \in \Theta(\tau \cap \tau')$ be as given by Lemma 6.1. Then $c_\zeta \rightarrow c \in \Theta^*(\tau \cap \tau')$ and $\mu(a, b, c) = c$; that is, $c \in [a, b]$. \square

In particular, it follows that $\Theta^*(\tau) \cap \Theta^*(\tau') \subseteq \Theta^*(\tau \cap \tau')$. Therefore, given any $a \in \mathcal{R}^*$, there is a unique minimal $\tau \in \mathcal{US}$ with $a \in \Theta^*(\tau)$. We write $\tau(a) = \tau$. Since the sets $\Theta^*(\tau)$ are all closed, we see that the map $\tau : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{US}$ is lower semicontinuous. Moreover, given any $a, b \in \mathcal{R}^*(\Sigma)$, there is some $c \in [a, b]$ with $\tau(c) \subseteq \tau(a) \cap \tau(b)$.

More generally, suppose that $C \subseteq \mathcal{R}^*$ is convex. Choose $a \in C$ with $\tau(a)$ minimal. If $b \in C$, then there is some $c \in [a, b] \subseteq C$ with $\tau(c) \subseteq \tau(a) \cap \tau(b)$, so $\tau(a) = \tau(c) \subseteq \tau(b)$. We write $\tau(C) = \tau(a)$. Thus, $\tau(C)$ is uniquely determined by

the property that $\tau(C) \subseteq \tau(b)$ for all $b \in C$ and $\tau(C) = \tau(a)$ for some $a \in C$. Note that this applies in particular, if C is a component of \mathcal{R}^* . Note that a component, C , of \mathcal{R}^* is thick (i.e. $C \cap \mathcal{R}_T^*(\Sigma) \neq \emptyset$) if and only if $\tau(C) = \emptyset$.

Now, given $\tau \in \mathcal{US}$, let $\Xi(\tau) = \{a \in \mathcal{R}^*(\Sigma) \mid \tau(a) = \tau\}$. Clearly, $\Xi(\tau) \subseteq \Theta^*(\tau)$ and $\Xi(\emptyset) = \Theta^*(\emptyset) = \mathcal{R}_T^*$. Since $\tau : \mathcal{R}^* \rightarrow \mathcal{US}$ is lower semicontinuous, we have that $\Xi(\tau)$ is open in $\Theta^*(\tau)$. Also:

Lemma 6.3. *For all $\tau \in \mathcal{US}$, $\Xi(\tau)$ is dense in $\Theta^*(\tau)$.*

Proof. Let $a \in \Theta^*(\tau)$, and choose $a_\zeta \in \Theta(\tau_\zeta)$, with $a_\zeta \rightarrow a$. Thus $\tau(a_\zeta) \subseteq \tau_\zeta$. Given any $i \in \mathbb{N}$, let $a_{i,\zeta} \in \mathcal{R}$ be the decorated multicurve with $\bar{a}_{i,\zeta} = \bar{a}_\zeta$, and resetting the decoration, $\eta_{\alpha_{i,\zeta}}(\gamma)$, on each $\gamma \in \tau_\zeta$ equal to $\eta_{a_\zeta}(\gamma) + i$. Now $(a_{i,\zeta})_i$ is a (quasi)geodesic sequence in \mathcal{R} with $a_{0,\zeta} = a_\zeta$ and with $N(a_{i,\zeta}; i) \subseteq \Theta(\tau_\zeta)$ for all i . In fact, $\rho(a_i, \Theta(\tau'_i)) \geq i$ for all $\tau' \neq \tau$.

Passing to the asymptotic the cone, we see that from any $a \in \Theta^*(\tau)$, there is a bilipschitz embedded ray, λ , emanating for which $\lambda(t) \in \Xi(\tau)$ for all $t > 0$. \square

In fact, the argument shows that $\Xi(\tau) \cap C$ is dense in $\Theta^*(\tau) \cap C$ for any component, C of \mathcal{R}^* .

Write $\mathcal{S}_C \subseteq \mathcal{S}$ for the set of standard complete multicurves, and write $\mathcal{US}_C \subseteq \mathcal{US}$ for the set of complete multicurves.

Note that, if $\tau \in \mathcal{S}_C$, then any multicurve which does not cross τ must be contained in τ . Thus, $L(\tau) = \{a \in \mathcal{R}(\Sigma) \mid \hat{a} \subseteq \tau \subseteq \bar{a}\}$. In particular, we see that $L(\tau)$ is a bounded Hausdorff distance from $\Theta(\tau)$. It follows that if $\tau \in \mathcal{US}_C$, then $L^*(\tau) = \Theta^*(\tau)$.

Note that, if $a \in \mathcal{R}(\Sigma)$, then (by the assumption on our marking graph) a is a bounded distance from a marking, b , which contains a complete multicurve, $\tau \supseteq \hat{a}$. Thus $\bigcup_{\tau \in \mathcal{S}_C} \Theta(\tau)$ is cobounded in $\mathcal{R}(\Sigma)$. We deduce:

Lemma 6.4. $\mathcal{R}^*(\Sigma) = \bigcup_{\tau \in \mathcal{US}_C} \Theta^*(\tau)$.

From this, we immediately get:

Lemma 6.5. *If $\tau \in \mathcal{US}_C$, then $\Xi(\tau)$ is open in \mathcal{R}^* . Also $\bigcup_{\tau \in \mathcal{US}_C} \Xi(\tau)$ is dense in any component of \mathcal{R}^* .*

In this way, we see that $(\Xi(\tau))_{\tau \in \mathcal{US}}$ defines a stratification of \mathcal{R}^* .

We also note:

Lemma 6.6. *For all $\tau \in \mathcal{US}$, $\Xi(\tau)$ lies in the interior of $L^*(\tau)$ in $\mathcal{R}^*(\Sigma)$.*

Proof. For all $a, b \in \Xi(\tau)$, then by lower semicontinuity, there is some open $U \subseteq \mathcal{R}^*$ with $\tau(b) \supseteq \tau$ for all $b \in U$. So $b \in L^*(\tau(b)) \subseteq L^*(\tau)$. This shows that $U \subseteq L^*(\tau)$. \square

Another way to say this is that for all $a \in \mathcal{R}^*(\Sigma)$, a lies in the interior of $L^*(\tau(a))$ in $\mathcal{R}^*(\Sigma)$.

In what follows, let $\mathcal{R}^\infty(\Sigma)$ be any component of any extended asymptotic cone, $\mathcal{R}^*(\Sigma)$ of $\mathcal{R}(\Sigma)$.

Recall that \mathcal{S}_B is the set of big multicurves in Σ . If $\tau \in \mathcal{US}_B$, then $L^\infty(\tau)$ is a closed convex subset which is a direct product of ξ almost furry \mathbb{R} -trees, as defined in Section 2.4. Directly from Proposition 2.6, we get:

Lemma 6.7. *Suppose that $f : \mathcal{R}^\infty(\Sigma) \rightarrow \mathcal{R}^\infty(\Sigma)$ is a homeomorphism and that $\tau \in \mathcal{US}_B$. Then $f|L^\infty(\tau)$ is a median isomorphism onto its range, $L^\infty(\tau)$, which is convex.*

For reference in Section 8, we also note the following. Recall that we have a 1-lipschitz map, $h^* : \mathcal{R}^*(\Sigma) \rightarrow \mathbb{R}^*$, taking non-negative values. Suppose we identify $L^*(\tau)$ with $\prod_{X \in \mathcal{UX}(\tau)} \mathcal{R}^*(X)$, via the map v_τ^* , as described above. Then if $a \in L^*(\tau)$, we have $h^*(a) = \sum_{X \in \mathcal{UX}(\tau)} h_X^*(a)$, where h_X^* is the corresponding map in the factor $\mathcal{R}^*(X)$. Here the sum is taken in the ordered abelian group \mathbb{R}^* . (In practice, we are only really interested in the cocycles $[(x, y) \mapsto h^*(x) - h^*(y)]$, which take real values on any component of $\mathcal{R}^*(\Sigma)$.) The statement follows immediately from the fact that the same formula holds for the maps h and h_X defined on $\mathcal{R}(\Sigma)$ and on $\mathcal{R}(X)$. If τ is a complete multicurve, then we get $h^*(a) = \sum_{\gamma \in \tau} h_\gamma^*(a)$, where $h_\gamma(a) = \eta_a(\gamma)$, as defined in Section 3.

7. QUASI-ISOMETRIC MAPS ON THE THICK PART

In this section, we apply the results of Section 6 to quasi-isometries between Teichmüller spaces. In particular, we give proofs of Theorems 1.5 and 1.6.

Let M be a topological median algebra. We say that a subset $P \subseteq M$ is *square-free* if there is no square in M with a side contained in P .

Recall, from Section 2.4, that a *gate map* to P is a (necessarily continuous) map $\omega : M \rightarrow P$ such that for all $x \in M$ and $y \in P$ we have $\omega(x) \in [x, y]$. In particular, it follows that P is convex and $\omega|P$ is the identity.

Note that if P is square-free and $x, y \in M$ with $[x, y] \cap P = \emptyset$, then $\omega x = \omega y$ (otherwise, we would have a square $\omega x, \omega y, \mu(x, y, \omega y), \mu(x, y, \omega x)$ with side $\{\omega x, \omega y\}$ in P). If M is weakly locally convex, it then follows that $\omega : M \rightarrow P$ is locally constant on $M \setminus P$.

Lemma 7.1. *Suppose that M is a weakly locally convex topological median algebra, and $P \subseteq M$ is closed convex and square-free and admits a gate map. If $p, x, y \in P$, and p separates x from y in P , then it also separates x from y in M .*

Proof. Let $\omega : M \rightarrow P$ be the gate map, which we have seen is a locally constant retraction. Now, $\omega^{-1}(p) \setminus \{p\}$ is open in M . Thus, if $P \setminus \{p\} = U \sqcup V$ is an open partition of $P \setminus \{p\}$, then $M \setminus \{p\} = (\omega^{-1}(U \cup \{p\}) \setminus \{p\}) \sqcup \omega^{-1}V$ is an open partition of $M \setminus \{p\}$. Taking U, V so that $x \in U$ and $y \in V$, the claim follows. \square

In particular, any cut point of P will also be a cut point of M .

In our situation, \mathcal{R}^∞ is certainly (weakly) locally convex. (This holds since it is bilipschitz equivalent to a median metric, or more directly the property given as the hypotheses to Lemma 2.5 here. See [Bo1] for more explanation.)

We now consider some constructions of such P .

Definition. We say that a set $P \subseteq \mathcal{R}(\Sigma)$ is *k-square-free* if given any square, Q , and any k -quasimorphism, $\phi : Q \rightarrow \mathcal{R}(\Sigma)$, which maps some side of Q into P , there is a (possibly different) side, c, d , of P with $\rho(\phi c, \phi d) \leq k$. We say that P is *coarsely square-free* if it is k -square-free for some $k \geq 0$.

Definition. Given $k \geq 0$ and $P \subseteq \mathcal{R}(\Sigma)$, we say that P is *k-straight* if $\text{diam } \theta_X(P) \leq k$ for $X \in \mathcal{X} \setminus \{\Sigma\}$. We say P is *coarsely straight* if it is k -straight for some $k \geq 0$.

Note that by the distance formula of Rafi (Proposition 4.8 here), we see that if $a, b \in P$, then $\rho(a, b) \asymp \sigma_\Sigma(a, b)$. (Here \asymp denotes agreement to within linear bounds depending on k .) In other words, the map $\theta_\Sigma : P \rightarrow \mathcal{H}(\Sigma) = \mathcal{G}(\Sigma)$ is a quasi-isometric embedding.

Lemma 7.2. *A coarsely straight set is coarsely square-free.*

Proof. Let $Q = \{a, b, d, c\}$ be a square, with sides, $\{a, b\}$ and $\{a, c\}$, and let $\phi : Q \rightarrow \mathcal{R}$ be a quasimorphism, with $\phi a, \phi b \in P$. Now $\theta_\Sigma \circ \phi : Q \rightarrow \mathcal{H}(\Sigma)$ is also a quasimorphism. If $\rho(\phi a, \phi b) \asymp \sigma_\Sigma(\phi a, \phi b)$ is sufficiently large, then an elementary property of hyperbolic spaces (essentially the fact the median is rank-1) tells us that $\sigma_\Sigma(\phi a, \phi c)$ and $\sigma_\Sigma(\phi b, \phi d)$ are both bounded.

Suppose $X \in \mathcal{X}$. Again, if $\sigma_\Sigma(\phi a, \phi b)$ is large, then at least one of $\sigma_\Sigma(\phi a, X)$ or $\sigma_\Sigma(\phi b, X)$ is large. In the former case, the Gromov product $\langle \sigma_\Sigma \phi a, \sigma_\Sigma \phi b : \sigma_\Sigma X \rangle_\Sigma$ is large. By Lemma 4.5, $\sigma_X(\phi a, \phi b)$ is bounded. Since $\theta_X \circ \phi : Q \rightarrow \mathcal{H}(X)$ is a quasimorphism, and $\mathcal{H}(X)$ is hyperbolic, it follows (as above) that $\sigma_X(\phi b, \phi d)$ is also bounded. In the latter case (that is, $\sigma_\Sigma(\phi b, X)$ is large), it similarly follows that $\sigma_X(\phi a, \phi c)$ and $\sigma_X(\phi b, \phi d)$ are both bounded. Since this holds for all $X \in \mathcal{X}$, it follows by Lemma 4.3 that $\rho(\phi a, \phi c)$ and $\rho(\phi b, \phi d)$ are bounded.

In summary, we have shown that either $\rho(\phi a, \phi b)$ or $\rho(\phi a, \phi c)$ is bounded as required. \square

Lemma 7.3. *Suppose that P is coarsely straight. Then P is quasiconvex in $\mathcal{R}(\Sigma)$ if and only if $\theta_\Sigma P$ is coarsely median convex in $\mathcal{G}(\Sigma)$.*

Proof. The “only if” part is an immediate consequence of the fact that $\theta_\Sigma : \mathcal{R}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$ is a quasimorphism. For the converse, suppose $a, b \in P$ and $c \in \mathcal{R}$. Since P is quasiconvex, there is some $d \in P$ with $\theta_\Sigma d \sim \mu(\theta_\Sigma a, \theta_\Sigma b, \theta_\Sigma c)$, and so $\theta_\Sigma d \sim \theta_\Sigma \mu(a, b, c)$. If $X \in \mathcal{X} \setminus \{\Sigma\}$, then $\theta_X a \sim \theta_X b \sim \theta_X d \sim \theta_X \mu(a, b, c)$, and so $\theta_X d \sim \mu(\theta_X a, \theta_X b, \theta_X c) \sim \theta_X \mu(a, b, c)$. It follows by Lemma 4.3, that $d \sim \mu(a, b, c)$. In other words, this shows that the coarse interval $[a, b]$ lies in a bounded neighbourhood of P as required. \square

Clearly, if P is quasiconvex and coarsely square free, then P^∞ is closed convex and square-free in \mathcal{R} .

We can construct examples of such P from coarsely straight sequences.

Definition. We say that a bi-infinite sequence, $(a_i)_{i \in \mathbb{Z}}$, in $\mathcal{R}(\Sigma)$ is *coarsely straight* if $\rho(a_i, a_{i+1})$ is bounded above for all i , and if $\sigma_\Sigma(a_i, a_j)$ is bounded below by an increasing linear function of $|i - j|$.

Note that, since the first condition implies also that $\sigma_X(a_i, a_{i+1})$ is bounded above, the second condition is equivalent to saying that the sequence $(\theta_\Sigma a_i)_i$ is quasigeodesic in $\mathcal{G}(\Sigma)$.

Lemma 7.4. *If $(a_i)_i$ is a coarsely straight sequence, then the set $\{a_i \mid i \in \mathbb{Z}\}$ is coarsely straight in $\mathcal{R}(\Sigma)$.*

Proof. Let $X \in \mathcal{X} \setminus \{\Sigma\}$. Since $(\theta_\Sigma a_i)_i$ is quasigeodesic in the hyperbolic space $\mathcal{G}(\Sigma)$, we can find $m < n \in \mathbb{Z}$, with $n - m$ bounded, and with $\langle a_i, a_m : X \rangle_\Sigma \geq l_1$ and $\langle a_j, a_n : X \rangle_\Sigma \geq l_1$ for all $i \leq m$ and all $j \geq n$, where l_1 is the constant from Lemma 4.5. (Of course, this might hold for all $i, j \in \mathbb{Z}$.) By Lemma 4.5, it follows that $\sigma_X(a_i, a_m)$ and $\sigma_X(a_j, a_n)$ are bounded for all $i \leq m$ and all $j \geq n$. Also, if $m \leq i \leq j \leq n$, then $\rho(a_i, a_j)$ is bounded, so $\sigma_X(a_i, a_j)$ is bounded. It follows that $\sigma_X(a_i, a_j)$ is bounded for all $i, j \in \mathbb{Z}$ as required. \square

It now follows that if $(a_i)_i$ is coarsely straight, then it is also quasiconvex, and it is also quasigeodesic in $\mathcal{R}(\Sigma)$.

We can also define a coarse gate map to $P = \{a_i \mid i \in \mathbb{Z}\}$. Given $c \in \mathcal{R}$, we can find $n \in \mathbb{Z}$ such that $\mu(\theta_\Sigma a_n, \theta_\Sigma c, \theta_\Sigma a_i) \sim \theta_\Sigma a_n$ for all $i \in \mathbb{Z}$. Then $\theta_\Sigma \mu(a_n, c, a_i) \sim \theta_\Sigma a_n$. Also, for all $X \in \mathcal{X} \setminus \{\Sigma\}$, we have $\theta_X a_i \sim \theta_X a_n$, so $\theta_X \mu(a_n, c, a_i) \sim \theta_X a_n$. It follows that by setting $\omega(c) = a_n$ we obtain a coarse gate map $\omega : \mathcal{R}(\Sigma) \rightarrow P$.

Now P^* is a convex subset of $\mathcal{R}^*(\Sigma)$ median isomorphic to \mathbb{R}^* . Restricting to the standard component, $\mathcal{R}^\infty(\Sigma)$, of $\mathcal{R}^*(\Sigma)$, we see that $P^\infty = P^* \cap \mathcal{R}^\infty(\Sigma)$ is closed and convex in $\mathcal{R}^\infty(\Sigma)$ and median isomorphic to \mathbb{R} . In particular, it is a bi-lipschitz embedding of \mathbb{R} . Moreover, P^∞ is square-free and admits a gate map, $\omega^\infty : \mathcal{R}^\infty(\Sigma) \rightarrow P^\infty$. It follows, by Lemma 7.1, that any point of P^∞ is a cut point of $\mathcal{R}^\infty(\Sigma)$.

Finally, note that if $g \in \text{Map}(\Sigma)$ is pseudoanosov, then the $\langle g \rangle$ -orbit of any point of $\mathcal{G}(X)$ is quasigeodesic (see [MaM1]). We see that if $a \in \mathcal{R}(\Sigma)$, then $(g^i a)_i$ is a coarsely straight sequence, hence quasiconvex by the above. This gives rise to a line in $\mathcal{R}^\infty(\Sigma)$ all of whose points are cut points of \mathcal{R}_T^∞ .

Since $\mathcal{U} \text{Map}(\Sigma)$ acts transitively on \mathcal{R}_T^* , we deduce:

Lemma 7.5. *Each point of \mathcal{R}_T^* is a cut point of the component of \mathcal{R}^* in which it lies.*

In fact, we have a converse:

Lemma 7.6. *If $x \in \mathcal{R}^* \setminus \mathcal{R}_T^*$, then x is not a cut point of the component in which it lies.*

Proof. Let $\tau = \tau(x) \in \mathcal{S}$. By assumption, $\tau \neq \emptyset$. By Lemma 6.6, x lies in the interior of $L^\infty(\tau)$ in \mathcal{R}^∞ . Therefore, if x were a cut point of \mathcal{R}^* , it would also be an intrinsic cut point of $L^\infty(\tau)$. But $L^\infty(\tau)$ is a median, hence topological, direct product of at least two non-trivial path-connected spaces, and so any two points of $L^\infty(\tau)$ lie in an embedded disc in $L^\infty(\tau)$. \square

We see that \mathcal{R}_T^* is determined by the topology of \mathcal{R}^* as the set of cut points. We deduce:

Lemma 7.7. *Suppose that Σ, Σ' are compact surfaces, and $f : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}^*(\Sigma')$ is a homeomorphism, then $f(\mathcal{R}_T^*(\Sigma)) = \mathcal{R}_T^*(\Sigma')$.*

(Given that ξ is the locally compact dimension of $\mathcal{R}^\infty(\Sigma)$, if such a homeomorphism exists, then $\xi(\Sigma) = \xi(\Sigma')$.)

Note that the above holds for any extended asymptotic cones, for any choice of scaling factors.

Now suppose that Σ and Σ' are compact surfaces of complexity at least 2 and that $\phi : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(\Sigma')$ is a quasi-isometry. This induces a (bilipschitz) homeomorphism, $f = \phi^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}^*(\Sigma')$.

Lemma 7.8. *There is some $k \geq 0$ such that the Hausdorff distance between $\phi(\mathcal{R}_T(\Sigma))$ and $\mathcal{R}_T(\Sigma')$ is at most k .*

Proof. By symmetry (swapping the roles of Σ and Σ'), it's enough to show that $\phi(\mathcal{R}_T(\Sigma))$ lies in a bounded neighbourhood of $\mathcal{R}_T(\Sigma')$. Suppose to the contrary that we can find an \mathbb{N} -sequence, $(x_i)_{i \in \mathbb{N}}$, in $\mathcal{R}_T(\Sigma)$ with $r_i = \rho(\phi(x_i), \mathcal{R}_T(\Sigma')) \rightarrow \infty$. Let $\mathcal{R}^\infty(\Sigma)$ and $\mathcal{R}^\infty(\Sigma')$ be asymptotic cones with $\mathcal{Z} = \mathbb{N}$, scaling factors $(1/r_i)_i$ and basepoints $(x_i)_i$ and $(\phi x_i)_i$. We get a homeomorphism, $f = \phi^\infty : \mathcal{R}^\infty(\Sigma) \rightarrow \mathcal{R}^\infty(\Sigma')$ with $\rho(f(x), \mathcal{R}_T^\infty(\Sigma')) = 1$. But $f(x) \in f(\mathcal{R}_T^\infty(\Sigma)) = \mathcal{R}_T^\infty(\Sigma')$, giving a contradiction. \square

To see that k depends only on $\xi(\Sigma)$ and the constants of quasi-isometry of ϕ , we apply the usual argument — allowing the maps to vary. In other words, we get a sequence, $\phi_i : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(\Sigma')$, of uniform quasi-isometries, and $x_i \in \mathcal{R}_T(\Sigma)$, with $r_i = \rho(\phi_i(x_i), \mathcal{R}_T(\Sigma')) \rightarrow \infty$. We get a limiting homeomorphism, $f : \mathcal{R}^\infty(\Sigma) \rightarrow \mathcal{R}^\infty(\Sigma')$, and derive a contradiction as before.

This proves Theorem 1.6.

It now follows that ϕ is a bounded distance from a map from $\mathcal{R}_T(\Sigma)$ to $\mathcal{R}_T(\Sigma')$. Recall (Lemma 3.4) that $\mathcal{R}_T(\Sigma)$ is a uniformly embedded copy of the marking complex, $\mathcal{M}(\Sigma)$, of Σ . Thus ϕ gives rise to a quasi-isometry from $\mathcal{M}(\Sigma)$ to $\mathcal{M}(\Sigma')$.

Now, if $\mathcal{M}(\Sigma)$ and $\mathcal{M}(\Sigma')$ are quasi-isometric, then Σ and Σ' are homeomorphic, under the conditions described by Theorem 1.5. This follows using the result of [BeKMM, Ham], and is shown directly in [Bo4].

This proves Theorem 1.5.

8. QUASI-ISOMETRIC RIGIDITY

In this section, we complete the proof of quasi-isometric rigidity of the Teichmüller metric.

In view of Theorem 1.5, we can assume that $\Sigma = \Sigma'$. Let $\phi : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(\Sigma)$ be a quasi-isometry. By Theorem 1.6, $\phi(\mathcal{R}_T(\Sigma))$ is a bounded Hausdorff distance from $\mathcal{R}_T(\Sigma)$. As noted at the end of the previous section, this gives rise to a quasi-isometry from $\mathcal{M}(\Sigma)$ to itself. By quasi-isometric rigidity of the marking graph, [BeKMM, Ham], there is some $g \in \text{Map}(\Sigma)$ such that $\rho(gx, \phi x)$ is bounded for all $x \in \mathcal{R}_T(\Sigma)$. Postcomposing with g^{-1} , we may as well assume that g is the identity, so $\rho(x, \phi x)$ is bounded. Thus, up to bounded distance, we can assume that $\phi|_{\mathcal{R}_T(\Sigma)}$ is the identity.

Henceforth in this section we will assume that ϕ is the identity on $\mathcal{R}_T(\Sigma)$. We want to show that it is a bounded distance from the identity everywhere.

The basic idea is as follows. Points at any bounded distance from the thick part, $\mathcal{R}_T(\Sigma)$, get displaced a bounded distance by ϕ . On the other hand, “most” of $\mathcal{R}(\Sigma)$ has locally the structure of a direct product of hyperbolic spaces (namely the decorated complexes associated to curves or to complexity-1 subsurfaces). It is known that a quasi-isometry of a product of hyperbolic spaces of this type preserves the product structure up to bounded distance (and permutation of the factors). This means that if a point is moved a large distance in one of the factors, then a nearby point must get moved an even larger distance. (This arises from a general observation about quasi-isometries of hyperbolic spaces.) We can take care to arrange that this new point is no further from the thick part than the original. In this way, we can derive a contradiction to the statement about bounded displacement. Of course, to make sense of this, one would need a very careful elaboration of the results about quasi-isometries of products. Rather than attempt to formulate this, we will reinterpret these ideas in the asymptotic cone, where most of the argument will be carried out.

To begin, recall that we have a 1-lipschitz map, $h : \mathcal{R}(\Sigma) \rightarrow [0, \infty)$ defined by $h(a) = \rho(a, \mathcal{R}_T(\Sigma))$. This gives rise to a 1-lipschitz map $h^* : \mathcal{R}^*(\Sigma) \rightarrow \mathbb{R}^*$, with $h^* \geq 0$ and with $(h^*)^{-1}(0) = \mathcal{R}_T^*$. If $\mathcal{R}^\infty(\Sigma)$ is any component of $\mathcal{R}^*(\Sigma)$, then $h^*(x) - h^*(y) \in \mathbb{R}$, for all $x, y \in \mathcal{R}^\infty(\Sigma)$.

The map ϕ gives rise to a bilipschitz map, $\phi^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}^*(\Sigma)$, fixing $\mathcal{R}_T^*(\Sigma)$. Suppose $a \in \mathcal{R}^*(\Sigma)$ is moved a limited (i.e. real) distance by ϕ^∞ . Let $\mathcal{R}^*(\Sigma)$ be the component of $\mathcal{R}^*(\Sigma)$ containing a , and let $f = \phi^*|_{\mathcal{R}^\infty(\Sigma)}$. Then f is a bilipschitz self-homeomorphism of $\mathcal{R}^\infty(\Sigma)$. As usual, we write ρ^∞ for the metric on $\mathcal{R}^\infty(\Sigma)$. (At this stage, $\mathcal{R}^\infty(\Sigma)$ may, or may not, be the component of $\mathcal{R}^*(\Sigma)$ containing $\mathcal{R}_T^*(\Sigma)$.)

The following technical lemma will constitute the bulk of what remains of the proof.

Lemma 8.1. *Suppose that $f : \mathcal{R}^\infty(\Sigma) \rightarrow \mathcal{R}^\infty(\Sigma)$ is a bilipschitz homeomorphism, and that there is some $a \in \mathcal{R}^\infty(\Sigma)$ with $fa \neq a$. Then, given any $k \in [0, \infty) \subseteq \mathbb{R}$, there is some $b \in \mathcal{R}^\infty(\Sigma)$ with $h^*(b) \leq h^*(a)$ and with $\rho^\infty(b, fb) \geq k$.*

We will eventually apply this when $\rho^\infty(a, fa) = 1$ and $k = 2$, though this does not matter to the discussion at present.

Before starting on the proof of the lemma, we recall some more general facts about $\mathcal{R}^\infty = \mathcal{R}^\infty(\Sigma)$, as we have outlined in Section 2. We have noted that \mathcal{R}^∞ is a topological median algebra of rank $\xi = \xi(\Sigma)$. From [Bo2], we know that all intervals in \mathcal{R}^∞ are compact, and so any closed convex subset of \mathcal{R}^∞ admits a gate map.

It will also be convenient to note, again from [Bo2] (see Lemma 2.5 here) that $\mathcal{R}^\infty(\Sigma)$ admits a bilipschitz equivalent median metric, ρ^M , which induces the same median structure. Note that Lemma 8.1 is equivalent to the same statement with ρ^∞ replaced by ρ^M , which is what we will actually prove. (Though we retain then original definition of h^* .)

By Lemma 6.4, we have $a \in L^*(\tau)$ for some complete multicurve, $\tau \in \mathcal{US}_C$. Writing $\tau = \{\gamma_1, \dots, \gamma_\xi\}$, we see that $L^*(\tau)$ can be identified, via a bilipschitz median isomorphism with $\prod_{i=1}^\xi \mathcal{H}^*(\gamma_i)$. As described at the end of Section 6, if $x \in L^*(\tau)$, then $h^*(x) = \sum_{i=1}^\xi h_i^*(x)$, where $h_i^* = h_{\gamma_i}^* : \mathcal{H}^*(\gamma_i) \rightarrow \mathbb{R}^*$ is the map described towards the end of Section 4.

Now let $D = L^\infty(\tau) = L^*(\tau) \cap \mathcal{R}^\infty(\Sigma)$. For each $i \in \{1, \dots, \xi\}$, let Δ_i be the factor of D parallel to $\mathcal{H}^\infty(\gamma_i)$ which contains a . In this way, we can identify $D \cong \prod_{i=1}^\xi \Delta_i$, via a bilipschitz median isomorphism.

(We remark that since D is convex, and ρ^∞ is already a median metric on D (as the l^1 product of \mathbb{R} -trees), we can assume that ρ^M is equal to ρ^∞ on D . In this case, the map $[(x, y) \mapsto h^*(x) - h^*(y)]$ is a Busemann cocycle with respect to the metric ρ^M . However, this is not formally needed the proof we give below.)

Now D is a direct product of ξ almost furry trees, so by Proposition 2.6, we see that $f|D$ is a median isomorphism onto its range, $D' = f(D)$. Moreover, D' is closed and convex in \mathcal{R}^∞ . In particular, it follows that for each i , $f|\Delta_i$ is a median isomorphism to $\Delta'_i = f(\Delta_i)$, which is also closed and convex in \mathcal{R}^∞ . As observed above, there are gate maps, $\omega_i : \mathcal{R}^\infty \rightarrow \Delta_i$ and $\omega'_i : \mathcal{R}^\infty \rightarrow \Delta'_i$.

We first reduce to the case where Δ_i and Δ'_i are parallel. We state this formally as follows:

Lemma 8.2. *Suppose that the conclusion of Lemma 8.1 fails. Given any i , then Δ_i and Δ'_i are parallel, and $f|\Delta_i = \omega'_i|\Delta_i$.*

Proof. Following the discussion of Section 2.4, we let $C_i = \omega_i \Delta'_i \subseteq \Delta_i$ and $C'_i = \omega'_i \Delta_i \subseteq \Delta'_i$. These are closed convex subsets. Let $\lambda_i = \omega_i \omega'_i : \mathcal{R}^\infty \rightarrow C_i$ and $\lambda'_i = \omega'_i \omega_i : \mathcal{R}^\infty \rightarrow C'_i$. These are also gate maps: in this case, just the nearest point projections to the subtrees C_i and C'_i . Now C_i, C'_i are parallel, with inverse parallel isomorphisms, $\omega'_i|C_i$ and $\omega_i|C'_i$. (Of course, possibly $C_i = C'_i$.) Also, as observed

in Section 2.4, if $b \in \Delta_i$ and $b' \in \Delta'_i$, then $\rho^M(b, b') \geq \rho^M(b, C_i) + \rho^M(b', C'_i)$.

Case(1): Δ_i is thin.

We first consider the case where Δ_i is thin (in the sense defined at the end of Section 4 — recall that Δ_i is an isometric copy of $\mathcal{H}^\infty(\gamma_i)$). In this case, Δ_i is a complete 2^{\aleph_0} -regular \mathbb{R} -tree. Moreover, the map $[x \mapsto h^*(x) - h^*(y)]$ has no lower bound on any branch of Δ_i .

If $C_i \neq \Delta_i$, then there is a branch, T , of Δ_i , with $T \cap C_i = \emptyset$. This must contain points b , with $h_i^*(a) - h_i^*(b)$ arbitrarily large (real) and also $\rho^M(b, C_i)$ arbitrarily large. In particular, we can find $b \in T$ with $h_i^*(b) \leq h_i^*(a)$ and with $\rho^M(b, C_i) \geq k$. Now, $b \in \Delta_i$, $fb \in \Delta'_i$, and so, as observed in Section 2.4, we have $\rho^M(b, fb) \geq \rho^M(b, C_i) + \rho^M(fb, C'_i)$. In particular, $\rho^M(b, fb) \geq \rho^M(b, C_i) \geq k$. Moreover, we have $h_j^*(b) = h_j^*(a)$ for all $j \neq i$. Since $h^* = \sum_{i=1}^{\xi} h_i^*$, we have $h^*(b) \leq h^*(a)$. We have arrived at the conclusion of Lemma 8.1 in this case. We therefore have that $C_i = \Delta_i$.

Suppose now that $C'_i \neq \Delta'_i$. In this case, we have a branch T' of Δ'_i , with $T' \cap C'_i = \emptyset$. Now $f^{-1}T'$ is a branch of Δ_i . We can find points $b \in f^{-1}T'$ with $h_i^*(a) - h_i^*(b)$ arbitrarily large, and with $\rho^M(fb, C'_i)$ arbitrarily large. In particular, we can suppose that $h_i^*(b) \leq h_i^*(a)$ so $h^*(b) \leq h^*(a)$, and that $\rho^M(fb, C'_i) \geq k$. So as in the previous case, we have $\rho^M(b, fb) \geq \rho^M(fb, C_i) \geq k$, again proving the the conclusion of Lemma 8.1 in this case (which is formally a contradiction to the hypotheses of Lemma 8.2).

We can therefore assume that $C_i = \Delta_i$ and $C'_i = \Delta'_i$. In other words, Δ_i, Δ'_i are parallel (hence either equal or disjoint). The parallel map $\omega_i : \Delta'_i \rightarrow \Delta_i$ is an isometry in the metric ρ^M . Note that if $x \in \Delta_i$ and $y \in \Delta'_i$, then $\rho^M(x, y) = \rho^M(x, \omega_i y) + \rho^M(\Delta_i, \Delta'_i)$. Consider the map $g = \omega_i f : \Delta_i \rightarrow \Delta_i$. This is a bilipschitz self-homeomorphism. If g is not the identity, then we can easily find a branch, T , of Δ_i with $T \cap gT = \emptyset$. We can now find points $b \in T$, with $h_i^*(a) - h_i^*(b)$ and $\rho^M(b, gb)$ both arbitrarily large (real). In particular, we can suppose $h^*(b) \leq h^*(a)$ and that $\rho^M(b, fb) \geq \rho^M(b, gb) \geq k$, as required.

We can therefore assume that Δ_i, Δ'_i are parallel, and that g is the identity map. In other words, $f|_{\Delta_i} = \omega'_i|_{\Delta_i}$, proving Lemma 8.2 in this case.

Case(2): Δ_i is thick.

This means that $h_i^*(\Delta_i) = [0, \infty) \subseteq \mathbb{R}$. Now, $(h_i^*)^{-1}(0)$ is the set of extreme points of Δ_i . All other points have valence 2^{\aleph_0} .

In this case, let $\tau' = \tau \setminus \{\gamma_i\}$. This is a big multicurve. Let $X \in \mathcal{UX}$ be the complementary component of τ' which contains γ_i . This is a (non-standard) $S_{0,4}$ or $S_{1,1}$. This case was discussed at the end of Section 4. In particular, $\mathcal{R}^\infty(X)$ is a complete regular 2^{\aleph_0} -tree, and every branch of $\mathcal{R}^\infty(X)$ meets $\mathcal{R}_T^\infty(X)$. Let $\hat{D} = L^\infty(\tau') = L^*(\tau') \cap \mathcal{R}^\infty(\Sigma)$. We can identify \hat{D} as the direct product, $\mathcal{R}^\infty(X) \times \prod_{j \neq i} \mathcal{H}^*(\gamma_j)$ of almost furry \mathbb{R} -trees, via a bilipschitz median isomorphism. Let

\hat{D} be the factor of \hat{D} parallel to $\mathcal{R}^\infty(X)$ and containing a . Thus, we can identify $\hat{D} \equiv \hat{\Delta} \times \prod_{j \neq i} \Delta_j$. Now, $\hat{\Delta}$ is a closed convex subset, isometric to an \mathbb{R} -tree, and containing Δ_i . Note that the map $h_i^* : \Delta_i \rightarrow [0, \infty)$ extends to a map $h_i^* : \hat{\Delta} \rightarrow [0, \infty)$ (defined as for h^* on $\mathcal{H}^*(\Sigma)$, intrinsically to $\hat{\Delta}$, as discussed at the end of Section 4). Note that $\mathcal{R}_T^*(X)$ gets identified with $(h_i^*)^{-1}(0)$, and so every branch of $\hat{\Delta}$ meets this set. Also, from the discussion at the end of Section 6, we have $h^*(x) = \sum_{i=1}^{\xi} h_i^*(x)$ for all $x \in \hat{D}$.

By Proposition 2.6 (similarly as with D above) we see that $f|_{\hat{D}}$ is a median isomorphism onto its range, $\hat{D}' = f(\hat{D})$. Moreover, \hat{D}' is closed and convex in \mathcal{R}^∞ . Let $\hat{\Delta}' = f(\hat{\Delta})$. This is also closed and convex, and $f|_{\hat{\Delta}} : \hat{\Delta} \rightarrow \hat{\Delta}'$ is a median isomorphism.

We now proceed to argue as before, with $\hat{\Delta}$ playing the role of Δ_i . Instead of saying that we can find b with $h_i^*(a) - h_i^*(b)$ arbitrarily large, we now claim that our b will satisfy $h_i^*(b) = 0$. Since $h_i^*(a) \geq 0$ and $h_j^*(b) = h_j^*(a)$ for all $j \neq i$, we again get that $h^*(b) \leq h^*(a)$. We finally conclude, as before, that (we can assume that) $\hat{\Delta}$ and $\hat{\Delta}'$ are parallel, and that $f|_{\hat{\Delta}} = \omega'_i|_{\hat{\Delta}}$. Restricting to Δ_i , we get also that Δ_i, Δ'_i are parallel, and $f|_{\Delta_i} = \omega'_i|_{\Delta_i}$. \square

Proof of Lemma 8.1. We again assume that the conclusion of Lemma 8.1 fails. Since Lemma 8.2 holds for all $i \in \{1, \dots, \xi\}$, it follows that D and D' are parallel, and that $f|_D$ is the parallel (gate) map from D to D' .

Now either $D = D'$ or $D \cap D' = \emptyset$. But in the latter case, $\mathcal{R}^\infty(\Sigma)$ would contain a $(\xi + 1)$ -cube. (Take any ξ -cube, Q in D , then $Q \cup fQ$ would be a $(\xi + 1)$ -cube.)

This contradicts the fact that the rank of \mathcal{R}^∞ is equal to ξ .

Thus, $D = D'$, and $f|_D$ is the identity map on D . Since, by construction, $a \in D$, this contradicts the hypothesis that $fa \neq a$. \square

Proof of Theorem 1.7. Let $\mathcal{R}(\Sigma) \rightarrow \mathcal{R}(\Sigma)$ be a quasi-isometry, which we can assume fixes $\mathcal{R}_T(\Sigma)$. Note that, for $a \in \mathcal{R}(\Sigma)$, $\rho(a, \phi a)$ is (linearly) bounded above in terms of $h(a) = \rho(a, \mathcal{R}_T(\Sigma))$. Given $n \in \mathbb{N}$, let $m(n) = \max\{\rho(a, \phi a) \mid a \in \mathcal{R}^0(\Sigma), h(a) \leq n\}$. For simplicity, we assume that this is attained, though it makes no essential difference to the argument. We want to show that $m(n)$ is bounded, independently of n .

Suppose, for contradiction, that $m(n) \rightarrow \infty$ as $n \rightarrow \infty$. Choose $a_n \in \mathcal{R}^0(\Sigma)$, with $h(a_n) \leq n$ and with $\rho(a_n, \phi a_n) = m(n)$. Let $\mathcal{Z} = \mathbb{N}$, with any non-principal ultrafilter, and let $\mathcal{R}^*(\Sigma)$ be the extended asymptotic cone with scaling factors $1/m(n)$. Let $a \in \mathcal{R}^*(\Sigma)$ be the limit of $(a_n)_{n \in \mathcal{Z}}$. By construction, we have $\rho^*(a, \phi^* a) = 1$. Let $\mathcal{R}^\infty(\Sigma)$ be the component containing a (hence also $\phi^* a$). Writing $f = \phi^*|_{\mathcal{R}^\infty(\Sigma)}$, we have that $f : \mathcal{R}^\infty(\Sigma) \rightarrow \mathcal{R}^\infty(\Sigma)$ is a bilipschitz homeomorphism with respect to the metric \mathcal{R}^∞ .

We are now in the situation described by Lemma 8.1, with $\rho^\infty(a, fa) = 1$. Set $k = 2$, and let b be the point obtained with $h^*(b) \leq h^*(a)$ and with $\rho^\infty(b, fb) = 2$. (Any real number bigger than 1 would do.)

We claim that we can find a sequence, $(b_n)_n$ in $\mathcal{R}(\Sigma)$, with $b_n \rightarrow b$ (in the sense of ultralimits) and with $h(b_n) \leq h(a_n)$ for almost all n . To see this, start with any sequence, $(c_n)_n$, converging on b . Note that $(h(a_n) - h(c_n))/m(n)$ tends to $h(a) - h(b) \in [0, \infty) \subseteq \mathbb{R}$. We take b_n to be the point a distance $\max\{0, h(c_n) - h(a_n)\}$ along a shortest geodesic in $\mathcal{R}(\Sigma)$ from c_n to $\mathcal{R}_T(\Sigma)$. In this way, $h(b_n) \leq h(a_n)$, and $\rho(c_n, b_n)/m(n) = \max\{0, (h(c_n) - h(a_n))/m(n)\} \rightarrow \max\{0, h^*(b) - h^*(a)\} = 0$. Since $c_n \rightarrow b$, we also have $b_n \rightarrow b$ as required.

In particular, $h(b_n) \leq n$ and so $\rho(b_n, \phi b_n) \leq m(n)$. Passing to the ultralimit, we get $\rho^\infty(b, fb) \leq 1$, contradicting $\rho^\infty(b, fb) = 2$.

We conclude that $m(n)$ is bounded, by some constant, m , say. It now follows that $\rho(a, \phi a) \leq m$ for all $a \in \mathcal{R}^0(\Sigma)$.

It remains to note that the bound, m , above depends only on $\xi(\Sigma)$ and the quasi-isometric constants of ϕ . This follows by the usual modification of the argument. If it were not bounded, we could find a sequence of uniform quasi-isometries, $\phi_n : \mathcal{R}(\Sigma) \rightarrow \mathcal{R}(\Sigma)$, all fixing $\mathcal{R}_T(\Sigma)$ but with arbitrarily large displacement somewhere. The argument now proceeds as before, except that the map ϕ^* arises as a limit of the maps $(\phi_n)_n$, rather than from a single map. We again, derive a contradiction in the same way. \square

9. CONCLUSION OF THE PROOF OF THEOREM 1.4

In this section, we prove the ‘‘only if’’ part of Theorem 1.4. Recall that this states that if $\mathbb{R}^{\xi(\Sigma)}$ quasi-isometrically embeds in $\mathbb{T}(\Sigma)$, then Σ has genus at most 1, or is a closed surface of genus 2.

We begin with some observations, which are valid in a quite general context (cf. [Bo4]). Recall that we have a collection of spaces, $\mathcal{R}(X)$, $\mathcal{H}(X)$ and maps, θ_X , χ_X and ψ_{YX} indexed by the set, \mathcal{X} , of subsurfaces of Σ , and satisfying properties (A1)–(A10) as discussed in Section 4. In fact, we have in addition, the distance formula of [Ra, D1], given as Proposition 4.8 here (which we have noted implies (A6) and (A7)). This allows us to bring various results of [Bo4] into play, some of which we have already noted. In particular, $\mathcal{R}^\infty(\Sigma)$ is a rank- ξ topological median algebra, and ρ^∞ is bilipschitz equivalent to a median metric ρ^M .

If $Q \subseteq \mathcal{R}^*(\Sigma)$ is an n -cube, we can partition its sides into n parallel classes. In particular, if c, d and c', d' are i th faces, then the intervals, $[c, d]$ and $[c', d']$ are parallel: that is, the maps $[x \mapsto \mu(c', d', x)]$ and $[x \mapsto \mu(c, d, x)]$ are inverse median isomorphisms between $[c, d]$ and $[c', d']$.

Now, if $n = \xi$, and c, d is an i th face of Q , the $[c, d]$ is a rank-1 median algebra (that is, a totally ordered set), in this case, isomorphic to an interval in \mathbb{R}^* . Moreover, the convex hull, $\text{hull}(Q)$, of Q is a median direct product of such intervals.

In particular, if $Q \subseteq \mathcal{R}^\infty(\Sigma)$, then $\text{hull}(Q)$ is median-isomorphic to the real cube $[-1, 1]^\xi$.

In general, by a *real n -cube* in a topological median algebra, M , we mean a closed convex subset median-isomorphic to $[-1, 1]^n$. If M is median metric space (for example, $(\mathcal{R}^\infty(\Sigma), \rho^M)$), this is isometric to an l^1 product of compact real intervals. We say that a subset of M is *culubated* if it is a locally finite union of cubes. After subdivision, we can assume that, in the neighbourhood of any given point, these cubes form the cells of a cube complex embedded in M . The following was proven in [Bo4] (see [BeKMM] for a related statement in the case of the mapping class group).

Proposition 9.1. *Let M be a complete median metric space of rank n , and let $\Phi \subseteq M$ be closed subset homeomorphic to \mathbb{R}^n . Then Φ is culubated.*

In particular, this applies to $\mathcal{R}^\infty(\Sigma)$ with $n = \xi$.

Returning to $\mathcal{R}^*(\Sigma)$, recall that we have maps $\psi_X^* = \psi_{X\Sigma}^* : \mathcal{R}^*(\Sigma) \rightarrow \mathcal{R}(X)$ for all $X \in \mathcal{X}$. Given any subset, $P \subseteq \mathcal{R}^*(\Sigma)$, write $D(P) = \{X \in \mathcal{UX} \mid \psi_X^*|_P \text{ is injective}\}$. We write $D^0(P) = D(P) \cap \mathcal{UG}^0(X)$, which we can identify with the set of curves $\gamma \in \mathcal{UG}^0$ such that $\theta_\gamma^*|_P$ is injective. (Recall that $\theta_\gamma^* = \psi_\gamma^*$.)

Now, if $Q \subseteq \mathcal{R}^*(\Sigma)$ is a cube, and c, d and c', d' are i th faces, then since $[c, d]$ and $[c', d']$ are parallel, we have $D([c, d]) = D([c', d'])$. We denote this set by $D_i(Q)$. We similarly define $D_i^0(Q) = D^0([c, d]) = D_i(Q) \cap \mathcal{UG}^0$.

The following was shown in [Bo4].

Proposition 9.2. *Let $Q \subseteq \mathcal{R}^*(\Sigma)$ be a $\xi(\Sigma)$ -cube. For any i , the set $D_i^0(Q)$ is either empty or consists of a single curve $\gamma_i \in \mathcal{UG}^0$. If it is empty, then there is a unique complexity-1 subsurface $Y_i \in D_i(Q)$. If the γ_i are all disjoint, and they form a big multicurve $\tau(Q)$. The Y_i are also disjoint, and are precisely the complexity-1 components of $\Sigma \setminus \tau(Q)$.*

We note, in particular, that γ_i or Y_i is completely determined by any face of Q , without reference to Q itself.

For the proof of Theorem 1.4, we also note the following:

Lemma 9.3. *Suppose that $\gamma \in \mathcal{UG}^0$, and that $a, b, c \in \mathcal{R}^*(\Sigma)$, with $c \in [a, b] \setminus \{a, b\}$, and with θ_γ^*a , θ_γ^*b and θ_γ^*c all distinct. Then $c \notin \mathcal{R}_T^*(\Sigma)$.*

Proof. Take $\gamma_\zeta \in \mathcal{G}^0(\Sigma)$ with $\gamma_\zeta \rightarrow \gamma$ and $a_\zeta, b_\zeta, c_\zeta \in \mathcal{R}(\Sigma)$ with $a_\zeta \rightarrow a$, $b_\zeta \rightarrow b$ and $c_\zeta \rightarrow c$. If $c \in \mathcal{R}_T^\infty$, then we could also take $c_\zeta \in \mathcal{R}_T$. Let $d_\zeta = \mu(a_\zeta, b_\zeta, c_\zeta)$. Thus $d_\zeta \rightarrow c$. Since θ_{γ_ζ} is a median quasimorphism, $\theta_{\gamma_\zeta}d_\zeta$ is a median of $\theta_{\gamma_\zeta}a_\zeta$, $\theta_{\gamma_\zeta}b_\zeta$, $\theta_{\gamma_\zeta}c_\zeta$ in $\mathcal{H}(\gamma_\zeta)$. Since $\mathcal{H}(\gamma_\zeta)$ is quasi-isometric to a horodisc, it is easily seen that $\min\{\sigma_{\gamma_\zeta}(a_\zeta, c_\zeta), \sigma_{\gamma_\zeta}(b_\zeta, c_\zeta)\}$, is bounded above by a linear function of $\sigma_{\gamma_\zeta}(c_\zeta, d_\zeta)$. Passing to the limit, we see that $\min\{\sigma_\gamma^\infty(a, c), \sigma_\gamma^\infty(b, c)\} = 0$, so $\theta_\gamma^*c \in \{\theta_\gamma^*a, \theta_\gamma^*b\}$, giving a contradiction. \square

We now proceed to the proof of the ‘‘only if’’ part of Theorem 1.4.

Suppose that $\phi : \mathbb{R}^\xi \rightarrow \mathcal{R}(\Sigma)$ is a quasi-isometric embedding. Passing to an asymptotic cone with fixed basepoint, we get a map $f = \phi^\infty : \mathbb{R}^\xi \rightarrow \mathcal{R}^\infty(\Sigma)$, which is bilipschitz onto its image, $\Phi = f(\mathbb{R}^\xi)$. Note that the basepoint, o , of $\mathcal{R}^\infty(\Sigma)$ lies in $\Phi \cap \mathcal{R}_T^\infty(\Sigma)$.

By Proposition 9.1, Φ is cubulated. In fact, we can find a neighbourhood of the basepoint $o \in \Phi$, which has the structure of a finite cube complex, where each ξ -dimensional cell is the convex hull of a ξ -cube.

Now consider the link, Δ , of o in Φ . This is a simplicial complex which is a homology $(\xi - 1)$ -sphere. In particular, the $(\xi - 1)$ th dimensional homology of Δ is non-trivial.

Let Δ^0 be its vertex set. Each $x \in \Delta^0$ corresponds to a 1-cell of Φ , with one vertex o and the other denoted $a(x) \in \Phi \subseteq \mathcal{R}^\infty(\Sigma)$. Note that this 1-cell is precisely the median interval, $[o, a(x)]$. Now, as noted after Proposition 9.2, we can canonically associate to x , either a curve $\gamma(x) \in \mathcal{UG}^0$, or a complexity-1 subsurface, $Y(x) \in \mathcal{UX}$. (To be specific, either $D^0([o, a(x)]) = \{\gamma(x)\}$, or else $D^0([o, a(x)]) = \emptyset$ and $Y(x)$ is the unique complexity-1 surface in $D([o, a(x)])$.) Any $(\xi - 1)$ -simplex in Δ corresponds to a ξ -cube, and so, by Proposition 9.2, we see that the curves $\gamma(x)$ or subsurfaces $Y(x)$ are all disjoint, as x ranges over the vertices of the simplex.

Lemma 9.4. *Suppose that $x, y \in \Delta^0$ are distinct, and that both correspond to curves, $\gamma(x)$ and $\gamma(y)$. Then $\gamma(x) \neq \gamma(y)$.*

Proof. Suppose, for contradiction, that $\gamma(x) = \gamma(y) = \gamma$, say. Since the intervals $[o, a(x)]$ and $[o, a(y)]$ meet precisely in o , it follows that $o = \mu(a(x), a(y), o)$; in other words, $o \in [a(x), a(y)]$. Moreover, $\theta_\gamma^\infty|_{[o, a(x)]}$ and $\theta_\gamma^\infty|_{[o, a(y)]}$ are both injective (by construction of $\gamma(x)$ and $\gamma(y)$), and so, in particular, it follows that $\theta_\gamma^\infty o$, $\theta_\gamma^\infty a(x)$ and $\theta_\gamma^\infty a(y)$ are all distinct. We now apply Lemma 9.3 to give the contradiction that $o \notin \mathcal{R}_T^\infty$. \square

Lemma 9.5. *Each $x \in \Delta^0$ corresponds to some $\gamma(x) \in \mathcal{UG}^0(\Sigma)$.*

Proof. Suppose, to the contrary, that $x \in \Delta^0$ corresponds to a complexity-1 subsurface, $Y(x)$. Let Q be any ξ -cube of the cubulation containing $a(x)$; so that $\{o, a(x)\}$ is a face. (This corresponds to a $(\xi - 1)$ -simplex in Δ .) Let $\tau = \tau(Q)$ be the big multicurve described by Proposition 9.2. Let $\gamma \in \tau$ be a boundary curve of $Y(x)$ in Σ . Thus, $\gamma = \gamma(y)$ for some $y \in Q$ (adjacent to x in Δ^0). Let $Q_0 \subseteq Q$ be the $(\xi - 1)$ -face containing o but not containing $a(x)$. (This corresponds to a $(\xi - 2)$ -simplex of Δ .) Now, given that Φ is homeomorphic to \mathbb{R}^ξ , it is easily seen that there must be a (unique) ξ -cube, Q' , of the cubulation with $Q \cap Q' = Q_0$. Let $z \in \Delta^0$ be the unique vertex with $a(z) \in Q' \setminus Q_0$. Let $\tau' = \tau(Q')$. Now τ' is obtained from τ by replacing $\gamma(y)$ by $\gamma(z)$ and leaving all other curves alone. Note also that all complexity-1 components of the complement also remain unchanged (since these are the subsurfaces $Y(w)$ for those vertices w which do not correspond

to curves). It therefore follows that, in fact, we must have $\gamma(z) = \gamma(y) = \gamma$, contradicting Lemma 9.4. In other words, this situation can never arise, and so each $x \in \Delta^0$ must correspond to a curve $\gamma(x)$. \square

By Lemmas 9.5 and 9.4, we therefore have an injective map $[x \mapsto \gamma(x)] : \Delta^0 \rightarrow \mathcal{UG}^0(\Sigma) = \mathcal{UC}^0(\Sigma)$. Now (again since Φ is homeomorphic to \mathbb{R}^ξ), every edge of Δ lies inside some $(\xi - 1)$ -simplex. So if $x, y \in \Delta^0$ are adjacent, $a(x), a(y)$ lie in some cube of the cubulation, and so $\gamma(x), \gamma(y)$ are distinct and disjoint. Now the ultraproduct, $\mathcal{UC}(\Sigma)$, of the curve complex $\mathcal{C}(\Sigma)$ is a flag complex (since $\mathcal{C}(\Sigma)$ is) so we get an injective simplicial map of Δ into $\mathcal{UC}(\Sigma)$. This gives us an injective map of Δ into $\mathcal{C}(\Sigma)$. Since $\mathcal{C}(\Sigma)$ has dimension $\xi - 1$, it follows that $\mathcal{C}(\Sigma)$ has non-trivial homology in dimension $\xi - 1$. But by the result of [Har] referred to in Section 5, the homology is trivial in all dimensions at least ξ' (as discussed after Proposition 5.7). It now follows that $\xi = \xi'$, and so Σ has genus at most 1, or is a closed surface of genus 2.

This proves the “only if” part of Theorem 1.4.

REFERENCES

- [AL] J.Aramayona, C.J.Leininger, *Finite rigid sets in curve complexes* : J. Topol. Anal. **5** (2013) 183–203.
- [BaH] H.-J.Bandelt, J.Hedlikova, *Median algebras* : Discrete Math. **45** (1983) 1–30.
- [Be] J.A.Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space* : Geom. Topol. **10** (2006) 1523–1578.
- [BeHS1] J.A.Behrstock, M.F.Hagen, A.Sisto, *Hierarchically hyperbolic spaces I: curve complexes for cubical groups* : preprint, 2014, posted at arXiv:1412.2171.
- [BeHS2] J.A.Behrstock, M.F.Hagen, A.Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula* : preprint, 2015, posted at arXiv:1509.00632.
- [BeKMM] J.A.Behrstock, B.Kleiner, Y.Minsky, L.Mosher, *Geometry and rigidity of mapping class groups* : Geom. Topol. **16** (2012) 781–888.
- [BeM] J.A.Behrstock, Y.N.Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.
- [Bo1] B.H.Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.
- [Bo2] B.H.Bowditch, *Embedding median algebras in products of trees* : Geom. Dedicata **170** (2014) 157–176.
- [Bo3] B.H.Bowditch, *Some properties of median metric spaces* : Groups, Geom. Dyn. **10** (2016) 279–317.
- [Bo4] B.H.Bowditch, *Large-scale rigidity properties of the mapping class groups* : preprint, Warwick, 2015.
- [Bo5] B.H.Bowditch, *Large-scale rank and rigidity of the Weil-Petersson metric* : preprint, Warwick, 2015.
- [Br] N.Broadbent, *Homology of the curve complex and the Steinberg module of the mapping class group* : Duke Math. J. **161** (2012) 1943–1969.
- [D1] M.G.Durham, *The augmented marking complex of a surface* : preprint, 2013, posted at arXiv:1309.4065.
- [D2] M.G.Durham, *The asymptotic geometry of the Teichmüller metric: Dimension and rank* : preprint, 2014, posted at arXiv:1501.00200.

- [EMR1] A.Eskin, H.Masur, K.Rafi, *Large scale rank of Teichmüller space* : preprint, 2013, posted at arXiv:1307.3733.
- [EMR2] A.Eskin, H.Masur, K.Rafi, *Rigidity of Teichmüller space* : preprint, 2015, posted at arXiv:1506.04774.
- [G] M.Gromov, *Asymptotic invariants of infinite groups* : “Geometric group theory, Vol. 2” London Math. Soc. Lecture Note Ser. No. 182, Cambridge Univ. Press (1993).
- [Ham] U.Hamenstädt, *Geometry of the mapping class groups III: Quasi-isometric rigidity* : preprint, 2007, posted at arXiv:0512429.
- [Har] J.L.Harer, *The virtual cohomological dimension of the mapping class group of an orientable surface* : Invent. Math. **84** (1986) 157–176.
- [HuM] J.Hubbard, H.Masur, *Quadratic differentials and foliations* : Acta. Math. **142** (1979) 221–274.
- [KR] M.Kapovich, K.Rafi, *Teichmüller space is semi-hyperbolic* : in preparation.
- [MaM1] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103–149.
- [MaM2] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.
- [Mc] C.T.McMullen, *The moduli space of Riemann surfaces is Kähler hyperbolic* : Ann. of Math. **151** (2000) 327–357.
- [Mi1] Y.N.Minsky, *Extremal length estimates and product regions in Teichmüller space* : Duke Math. J. **83** (1996) 249–286.
- [Mi2] Y.N.Minsky, *Quasiprojections in Teichmüller space* : J. Reine Angew. Math. **473** (1996) 121–136.
- [Mo] L.Mosher, *Stable Teichmüller quasigeodesics and ending laminations* : Geom. Topol. **7** (2003) 33–90.
- [Ra] K.Rafi, *A combinatorial model for the Teichmüller metric* : Geom. Funct. Anal. **17** (2007) 936–959.
- [Ro] H.L.Royden, *Automorphisms and isometries of Teichmüller space* : in “Advances in the Theory of Riemann Surfaces”, Ann. of Math. Studies No. 66, Princeton Univ. Press (1971) 369–383.
- [S] J.D.Stasheff, *Homotopy associativity of H-spaces. I* : Trans. Amer. Math. Soc. **108** (1963) 275–292.
- [VaW] L.Van den Dries, A.J.Wilkie, *On Gromov’s theorem concerning groups of polynomial growth and elementary logic* : J. Algebra **89** (1984) 349–374.
- [Ve] E.R.Verheul, *Multimedians in metric and normed spaces* : CWI Tract, 91, Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam (1993).
- [Y] S.-K.Yeung, *Quasi-isometry of metrics on Teichmüller spaces* : Int. Math. Res. Not. **4** (2005) 239–255.

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