

Original citation:

Bowditch, B. H. (2018) Large-scale rigidity properties of the mapping class groups. *Pacific Journal of Mathematics*, 293 (1). pp. 1-73.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/91088>

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:

<http://doi.org/10.2140/pjm.2018.293.1>

A note on versions:

The version presented in WRAP is the published version or, version of record, and may be cited as it appears here.

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

*Pacific
Journal of
Mathematics*

**LARGE-SCALE RIGIDITY PROPERTIES
OF THE MAPPING CLASS GROUPS**

BRIAN H. BOWDITCH

Volume 293 No. 1

March 2018

LARGE-SCALE RIGIDITY PROPERTIES OF THE MAPPING CLASS GROUPS

BRIAN H. BOWDITCH

We study the coarse geometry of the mapping class group of a compact orientable surface. We show that, apart from a few low-complexity cases, any quasi-isometric embedding of a mapping class group into itself agrees up to bounded distance with a left multiplication. In particular, such a map is a quasi-isometry. This is a strengthening of the result of Hamenstädt and of Behrstock, Kleiner, Minsky and Mosher that the mapping class groups are quasi-isometrically rigid. In the course of proving this, we also develop the general theory of coarse median spaces and median metric spaces with a view to applications to Teichmüller space, and related spaces.

1. Introduction

One of the main aims of this paper is to give an account of the quasi-isometric rigidity of the mapping class group of a closed orientable surface. Quasi-isometric rigidity was established in [Hamenstädt 2005; Behrstock, Kleiner, Minsky and Mosher 2012]. Here, we give a strengthening of this result which applies to quasi-isometric embeddings (see Theorem 1.1 below).

Many of our arguments have parallels with those of [Behrstock, Kleiner, Minsky and Mosher 2012], though the details are different. Another aim of this paper is to set these arguments in a broader context. The key observation, made in [Bowditch 2013], is that the mapping class group admits a “coarse median” structure. The median in this case is the *centroid* constructed in [Behrstock and Minsky 2011]. Here, in Section 7, we list a set of axioms related to subsurface projection (cf., [Masur and Minsky 2000]) which imply the existence of medians (see Theorem 1.4 below). The point is that the same axioms apply in other situations, notably to Teichmüller space in either the Teichmüller metric or the Weil–Peterson metric. It then follows that these also admit a coarse median structure. This is explained, respectively, in [Bowditch 2016a] and [Bowditch 2015], where various consequences of this observation for the large-scale geometry of these spaces are explored. Again, many of the arguments follow along similar lines, and several general results of this

MSC2010: 20F65.

Keywords: mapping class group, quasi-isometry, rigidity, median.

paper are used in those papers (see, for example, Propositions 1.2 and 1.3 below, as well as the structure of cubes discussed in Sections 10–12).

We begin by outlining the main results of this paper.

Let Σ be a compact orientable surface of genus g with p boundary components. Let $\xi(\Sigma) = 3g + p - 3$ be the *complexity* of Σ . Write $\text{Map}(\Sigma)$ for the mapping class group. When this is viewed as a geometric object, we will use different notation. In particular, we will write $\mathbb{M}(\Sigma)$ for the “marking graph” of Σ as discussed in Section 8. In fact, any proper geodesic space on which $\text{Map}(\Sigma)$ acts isometrically, properly discontinuously and compactly (such the Cayley graph with respect to any finite generating set) would serve for the present discussion. Any two such spaces will be $\text{Map}(\Sigma)$ -equivariantly quasi-isometric, by the Schwarz–Milnor Lemma.

It is shown in [Hamenstädt 2005; Behrstock and Minsky 2008] that $\mathbb{M}(\Sigma)$ has coarse rank equal to $\xi(\Sigma)$; that is the maximal dimension ν such that $\mathbb{M}(\Sigma)$ admits a quasi-isometric embedding of \mathbb{R}^ν (see also [Eskin, Masur and Rafi 2017, Corollary C] and [Bowditch 2013, Theorem 2.6]). Note that it follows that if Σ and Σ' are compact orientable surfaces with $\mathbb{M}(\Sigma)$ quasi-isometric to $\mathbb{M}(\Sigma')$, then $\xi(\Sigma) = \xi(\Sigma')$.

We will show:

Theorem 1.1. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 4$, and that $\phi : \mathbb{M}(\Sigma) \rightarrow \mathbb{M}(\Sigma')$ is a quasi-isometric embedding. Then $\Sigma = \Sigma'$, and ϕ is a bounded distance from the isometry of $\mathbb{M}(\Sigma)$ induced by some element of $\text{Map}(\Sigma)$.*

It immediately follows that ϕ is, in fact, a quasi-isometry. One can also deal, modulo some qualifications, with lower-complexity cases (see the discussion after Theorem 15.2 here). As observed above, if one assumes that ϕ is a quasi-isometry (and that $\Sigma = \Sigma'$), then this statement is given in [Hamenstädt 2005] and [Behrstock, Kleiner, Minsky and Mosher 2012].

We remark that if one assumes quasi-isometric rigidity as given in those papers, then one recovers (indirectly) that the quasi-isometry type of $\mathbb{M}(\Sigma)$ determines the topological type of Σ (modulo a few low-dimensional exceptional cases) since it determines $\text{Map}(\Sigma)$ up to isomorphism (see, for example, [Rafi and Schleimer 2011] for a proof that $\text{Map}(\Sigma)$ determines Σ). Given this, Theorem 1.1 would be equivalent to asserting that any quasi-isometric embedding of $\mathbb{M}(\Sigma)$ into itself is necessarily a quasi-isometry (at least when $\xi(\Sigma) \geq 4$). However, we will give another proof of the rigidity statement in this paper.

As noted above, we base our account around the notion of a coarse median space, as defined in [Bowditch 2013]. This is a geodesic metric space equipped with a ternary operation satisfying certain conditions. Roughly speaking, these say that when dealing with a finite number of points in the space, the ternary operation behaves, up to bounded distance, like the standard median operation on the vertex

set of a (finite) CAT(0) cube complex. Such a space comes with a notion of “rank” which is the maximal dimension of such a cube complex needed for the hypothesis.

A related, but different, notion is that of a median metric space, which is also central to our discussion. The definition of a median metric space is quite simple, and is given in Section 2. For further discussion, see [Verheul 1993; Chatterji, Druţu and Haglund 2010; Bowditch 2016c]. This has also been studied from a combinatorial viewpoint; see for example, [Chepoi 2000]. In a median metric space, any triple of points has a unique “median”, that is, a point lying between any pair in the triple. This defines a continuous ternary operation, and gives the space the structure of a topological median algebra. (For expositions of the theory of median algebras, see [Isbell 1980; Bandelt and Hedlíková 1983; Roller 1998].) Again, one can associate a “rank” to such a space as the maximal dimension of an embedded cube. (Any CAT(0) cube gives rise example of such a space, after one replaces the euclidean metric on each cube with the l^1 metric. The vertex set is then also such a space.) One can show that a complete connected median metric space of finite rank is canonically bi-Lipschitz equivalent to a CAT(0) metric, [Bowditch 2016c].

The asymptotic cone (see [van den Dries and Wilkie 1984; Gromov 1993]) of a coarse median space is a topological median algebra. If the space has finite rank ν , then the asymptotic cone is bi-Lipschitz equivalent to a median metric space of rank at most ν (see [Behrstock, Druţu and Sapir 2011; Bowditch 2014a] and Theorem 6.9 here). Also, the dimension of any compact subset thereof has dimension at most ν . (This follows from [Bowditch 2013] as we discuss in Section 2.) From the fact that $\text{Map}(\Sigma)$ is a coarse median space one gets a median on its asymptotic cone. This was previously obtained by other means in [Behrstock, Druţu and Sapir 2011]. Much of this is elaborated upon in [Bowditch 2013; 2014a]. Here we obtain more information about the flats in such spaces, which we use for the rigidity result of Theorem 1.1. Similar statements can be found in [Behrstock, Kleiner, Minsky and Mosher 2012], though more specifically for the mapping class group.

We remark that, in [Rafi and Schleimer 2011], the rigidity of the mapping class group is used to deduce the rigidity of the curve graph. Again, it would be interesting to generalise this to quasi-isometric embeddings. As the authors observe, much of their paper works for such embeddings. However there is a key point (aside from their references to [Hamenstädt 2005; Behrstock, Kleiner, Minsky and Mosher 2012]) where an inverse quasi-isometry is needed.

We briefly state a few of the key results proven in this paper, which are used in proving Theorem 1.1, and/or have applications elsewhere.

The first two relate to a median metric space. By a *real cube* in such a space, we will mean a median-convex subset isometric to a finite l^1 -product of compact real intervals. (See Section 3 for more precise definitions.) A (closed) subset is *cubulated* if it is a locally finite union of real cubes. We show:

Proposition 1.2. *Suppose that M is a complete median metric space of rank $v < \infty$, and that $\Phi \subseteq M$ is a closed subset homeomorphic to \mathbb{R}^v . Then Φ is cubulated.*

This is proven in Section 4 (see Proposition 4.3). Under additional topological assumptions one can show that Φ is median-convex and isometric to \mathbb{R}^v with the l^1 metric (see Proposition 4.6). Using this, one gets a result about products of \mathbb{R} -trees:

Proposition 1.3. *Suppose that M is a complete median metric space of rank $v < \infty$, with $v \geq 2$. Suppose that D is a direct product of v \mathbb{R} -trees, and that none of the factors has a point of valence 2 (i.e., a point which separates the \mathbb{R} -tree into exactly 2 components). Suppose that $f : D \rightarrow M$ is a continuous injective map, with closed image, $f(D) \subseteq M$. Then f is a median homomorphism, and $f(D)$ is median-convex in M .*

This is shown at the end of Section 4 (Proposition 4.8). This result is used in [Bowditch 2015; 2016a]. A more direct proof in a specific case is given in [Bowditch 2016b] (see Proposition 2.1 thereof). Analogous, but different, statements can be found in [Kleiner and Leeb 1997] and [Kapovich, Kleiner and Leeb 1998].

We make much use of subsurface projections from the marking graph, $\mathbb{M}(\Sigma)$, to curve graphs associated to subsurfaces of Σ . In Section 7, we condense the essential information we need into a set of axioms, (A1)–(A10). This means that much of the argument can be put in a more general setting. In particular, we have the following paraphrasing of a result which will be stated more formally in Section 7; see Theorems 7.1 and 7.2.

Theorem 1.4. *Suppose that to each subsurface, X , of Σ , we have associated geodesic metric spaces, $\mathcal{M}(X)$ and $\mathcal{G}(X)$, together with a collection of projection maps between them satisfying axioms (A1)–(A10). Then each $\mathcal{M}(X)$ has the natural structure of a coarse median space in such a way that each projection map is a quasimorphism (i.e., a median homomorphism up to bounded distance).*

Note that this includes the case where $X = \Sigma$. Here, the spaces $\mathcal{G}(X)$ are (assumed to be) uniformly hyperbolic and the median is the usual centroid in such a space. The various constants involved in the conclusion depend only on those of the hypotheses (A1)–(A10).

In this paper, we are interested mainly in the case where $\mathcal{M}(X)$ and $\mathcal{G}(X)$ are respectively the marking graph, $\mathbb{M}(X)$ and $\mathbb{G}(X)$ and the curve graph of the subsurface X . The same axioms can also be applied to Teichmüller space in either the Teichmüller metric [Bowditch 2016a] or the Weil–Petersson metric [Bowditch 2015].

A simple consequence of Theorem 1.4 is that the asymptotic cone of $\mathcal{M}(\Sigma)$ is a topological median algebra. In fact, it is bi-Lipschitz equivalent to a median metric space, which then allows us to bring the results mentioned above into play.

In outline this paper is structured as follows. Sections 2 to 4 are devoted to a general discussion of median metric spaces. In Section 5 we review properties of asymptotic cones. In Section 6 we discuss general coarse median spaces. In Section 7 we give a set of hypotheses relating to subsurface projection which imply that a geodesic metric space admits a coarse median, and give a precise formulation and proof of Theorem 1.4. This is then applied to the marking graph in Section 8. In Sections 9 to 13, we explore further properties of the marking graph and its asymptotic cone, setting as much as possible in a general context (so that it can be applied elsewhere to Teichmüller space). In Section 14, we explain, in general terms, how the asymptotic cone can be used to control Hausdorff distance. Finally, in Section 15, this is applied to the marking complex, to give a proof of Theorem 1.1, together with some discussion of the lower complexity cases.

Notation. Throughout this paper, we will use $\mathbb{G}(X)$ and $\mathbb{M}(X)$ respectively to denote the curve graph and marking graph of a subsurface X of Σ . (We allow $X = \Sigma$, and we need to modify the definitions in the case where X is an annulus, as discussed in Section 8.) We will use the notation $\mathcal{G}(X)$ and $\mathcal{M}(X)$ when the statements apply to the more general spaces satisfying the axioms laid out in Section 7. (This symbol \mathcal{M} will generally denote a coarse median space, as in Section 6.) Note that the curve graph $\mathbb{G}(X)$ plays two slightly different roles: it is one of the family of spaces satisfying these axioms; also its vertex set can be identified with the set of annular subsurfaces of X , and which in this capacity can be viewed as an indexing set. (We remark that in [Bowditch 2013], and some other references, the notation $\Theta(X)$ and $\Lambda(X)$ was used respectively for $\mathcal{G}(X)$ and $\mathcal{M}(X)$.) For the main applications in the present paper, there would be no loss in interpreting $\mathcal{G}(X)$, $\mathcal{M}(X)$ as $\mathbb{G}(X)$, $\mathbb{M}(X)$, respectively.

2. Median metric spaces

We begin with some general discussion of median metric spaces. For elaboration relevant to this paper, see for example, [Verheul 1993; Chatterji, Druţu and Haglund 2010; Bowditch 2016c].

Let (M, ρ) be a metric space. Given $a, b \in M$, let

$$[a, b] = [a, b]_\rho = \{x \in M \mid \rho(a, b) = \rho(a, x) + \rho(x, b)\}.$$

Thus, $[a, b] = [b, a]$ and $[a, a] = \{a\}$.

Definition. We say ρ is a *median metric* if, for all $a, b, c \in M$, $[a, b] \cap [b, c] \cap [c, a]$ consists of exactly one element of M .

We denote this element by $\mu(a, b, c)$ — the *median* of a, b, c . It follows using [Sholander 1954] that (M, μ) is a median algebra (see [Verheul 1993; Chatterji,

[Druţu and Haglund 2010] and Section 2 of [Bowditch 2016c]). Moreover, $[a, b]$ is exactly the median interval between a and b , i.e., $[a, b] = [a, b]_\mu = \{x \in M \mid \mu(a, b, x) = x\}$. Conversely, note that if (M, μ) is a median algebra, and ρ is a metric satisfying $[a, b]_\mu = [a, b]_\rho$ for all $a, b \in M$, then ρ is a median metric inducing μ . Also, the map $\mu : M^3 \rightarrow M$ is continuous (that is, M is a “topological median algebra”).

The following definitions only require the median structure on M .

Definition. A subset $B \subseteq M$ is a *subalgebra* if it is closed under μ . It is *convex* if $[a, b] \subseteq B$ for all $a, b \in B$. An *n-cube* is a subset of M median-isomorphic to the direct product of n two-point median algebras: $\{-1, 1\}^n$. (Note that any two-point set admits a unique median structure.) We refer to a 2-cube as a *square*. The *rank* of M is the maximal n such that M contains an n -cube. The rank is deemed to be infinite if there are cubes of all dimensions.

Given $A \subseteq M$ write $\langle A \rangle$ and $\text{hull}(A)$ respectively for the subalgebra generated by A and the convex hull of A , that is, respectively, the smallest subalgebra and smallest convex set in M containing A . Clearly $\langle A \rangle \subseteq \text{hull}(A)$. If A is finite, then so is $\langle A \rangle$. In fact, $|\langle A \rangle| \leq 2^{2^{|A|}}$. Any finite median algebra can be canonically identified as the vertex set of a finite CAT(0) complex; see [Chepoi 2000].

Note that any interval in a connected median metric space is connected (since the map $x \mapsto \mu(a, b, x)$ is a continuous retraction to $[a, b]$). It follows that any connected component of a median metric space is convex.

Definition. We say that a median metric space is *proper* if it is connected, complete and has finite rank.

Henceforth we will assume that M is a proper median metric space, though as we will comment, many of the constructions only require it to be a median metric space, or indeed just a median algebra.

It was shown in [Bowditch 2014a, Corollary 1.3] that if M is proper, then every interval $[a, b]$ in M is compact. (One can go on to deduce that the convex hull of any compact set is compact.)

We say that a topological median algebra is *locally convex* if every point has a base of convex neighbourhoods.

Lemma 2.1. *Any median metric space M of finite rank is locally convex.*

Proof. This follows since M is “weakly locally convex” in the sense of [Bowditch 2013, Section 7]. (Note that if $a, b \in M$, then the diameter of $[a, b]$ is equal to $\rho(a, b)$.) Since it has finite rank, Lemma 7.1 of [loc. cit.], tells us that it is locally convex. \square

(In the case of interest here, namely the asymptotic cone of a finite rank coarse median space, the conclusion also follows from Lemma 9.2 of [loc. cit.])

It was also shown there (Theorem 2.2) that any locally compact subset of M has topological dimension at most $\text{rank}(M)$. (For more discussion of dimension, see Section 4 of the present paper.)

Theorem 2.2 [Bowditch 2016c, Theorem 1.1]. *If (M, ρ) is a proper median metric space, then there is a canonically associated bi-Lipschitz equivalent metric, σ_ρ , on M for which (M, σ_ρ) is CAT(0).*

In fact, we can arrange that $\rho/\sqrt{\text{rank}(M)} \leq \sigma_\rho \leq \rho$.

Note that it immediately follows that M is contractible.

A simple example is \mathbb{R}^n with the l^1 metric. In this case, σ_ρ recovers the euclidean metric on \mathbb{R}^n . Any convex subset of \mathbb{R}^n has the form $P = \prod_{i=1}^n I_i$ where $I \subseteq \mathbb{R}$ is a real interval (possibly unbounded). If each I_i is either a singleton or all of \mathbb{R} , we refer to P as a *coordinate plane*. If each $I_i = [a_i, b_i]$ with $a_i < b_i$, we refer to P as an l^1 *cube*. We refer to $P = \prod_{i=1}^n (a_i, b_i)$ as the *relative interior* of P , and we refer to the elements of $Q = \prod_{i=1}^n \{a_i, b_i\}$ as the *corners* of P . Note that these are determined by the intrinsic geometry of P . Also $P = \text{hull}(Q)$. In fact, $P = [a, b]$ where a, b are any pair of opposite corners of P .

Another class of examples arise from CAT(0) complexes. Suppose that Υ is (the topological realisation of) a finite CAT(0) complex. Suppose that each cell is given the structure of an l^1 -cube. This induces a path metric, ρ , on Υ , so that (Υ, ρ) is a median metric space. In this case, (Υ, σ_ρ) is a euclidean CAT(0) cube complex, where we can allow the cells to be rectilinear parallelepipeds.

Definition. We refer to a space of the form (Υ, ρ) as an l^1 -*cube complex*.

There is a sense in which any proper metric median space can be approximated by subspaces of this form. The following was shown in [Bowditch 2016c, Lemmas 7.5 and 7.6].

Lemma 2.3. *Let (M, ρ) be a complete connected median metric space. Suppose that $\Pi \subseteq M$ is a finite subalgebra. Then there is a closed subset $\Upsilon \subseteq M$ which has the structure of a finite l^1 -cube complex in the induced metric ρ , and such that $\Pi \subseteq \Upsilon$ is exactly the set of vertices of this complex.*

The statement is taken to imply that the metric ρ restricted to Υ is already a path metric on Υ . In general, Υ will not be unique. (One can make a canonical choice by taking cells to be totally geodesic in the metric σ_ρ on M , but we will not need this here.) Note that we do not assume here that the cells of Υ are convex in M . (If that were the case, we refer to Υ as a “straight” cube complex, as we will define more formally in Section 3.)

We continue with some more general observations. For the moment, M can be any median metric space.

Given $a, b \in M$, we define $\phi = \phi_{a,b} : M \rightarrow [a, b]$ by $\phi(x) = \mu(a, b, x)$. This is a 1-Lipschitz median epimorphism.

Definition. We say that two pairs $(a, b), (c, d)$ in M^2 are *parallel* if $[b, c] = [a, d]$.

It is equivalent to saying that both $b, c \in [a, d]$ and $a, d \in [b, c]$. When a, b, c, d are all distinct, it is also equivalent to saying that a, b, d, c is a square. Note that parallelism is an equivalence relation on M^2 . If a, b and c, d are parallel, then $\phi_{a,b}|_{[c, d]}$ is an isometry (hence a median isomorphism) from $[c, d]$ to $[a, b]$. Its inverse is $\phi_{c,d}|_{[a, b]}$.

The following is a standard notion for median algebras.

Definition. If $C \subseteq M$ is closed and convex, we say that $\phi : M \rightarrow C$ is a *gate map* of M to C if $\phi(x) \in [x, c]$ for all $x \in M$ and $c \in C$.

A more detailed discussion of gate maps can be found in [Bowditch 2016a, Section 2.4].

One verifies that ϕ is a 1-Lipschitz retraction of M to C , and a median homomorphism. If ϕ exists then it is unique. Note that the map $\phi_{a,b}$ of the previous paragraph is a gate map to $[a, b]$. In fact, if M is proper, then gate maps to closed convex sets always exist. This can be seen using the fact that intervals are compact, though we will not need this here.

Definition. A *wall* in M is a partition of M into two nonempty convex subsets.

This is equivalent to a median epimorphism $\phi : M \rightarrow \{-1, 1\}$, where the partition is given by $\{\phi^{-1}(-1), \phi^{-1}(1)\}$. We can speak about an *oriented* or *unoriented* wall according to whether we consider the partition as an ordered or an unordered pair. Any two disjoint convex subsets, C, D , of M are separated by some wall, that is, $C \subseteq \phi^{-1}(-1)$ and $D \subseteq \phi^{-1}(1)$. We say two walls, ϕ, ψ , *cross* if the map $(\phi, \psi) : M \rightarrow \{-1, 1\}^2$ is surjective. The rank of M can be equivalently defined as the maximal cardinality of a set of pairwise crossing walls. We say that M is *n-colourable* if we can colour the walls of M with n colours such that no two walls of the same colour cross. This implies that $\text{rank}(M) \leq n$. (See [Bowditch 2013, Section 12] for more discussion of colourability.)

These notions only require the median structure on M . If Π is a finite median algebra, then we can identify the set of (unoriented) walls with the set of hyperplanes in the associated finite CAT(0) complex. In this case, two walls cross if and only if the corresponding hyperplanes intersect.

If $a, b \in M$, then $[a, b]$ admits a partial order defined by $x \leq y$ if $x \in [a, y]$ (or equivalently $y \in [x, b]$). If $[a, b]$ has rank 1, this is a total order. If M is connected and metrisable, then $[a, b]$ is isometric to a compact real interval. In particular, any connected median metric space of rank 1 is an \mathbb{R} -tree. (In this case, the metric σ_ρ , described above, agrees with ρ .)

We also note the following construction of quotient median algebras. Suppose that M is a median algebra, and that \sim is an equivalence relation on M such that whenever $a, b, c, d \in M$ with $c \sim d$, then $\mu(a, b, c) \sim \mu(a, b, d)$. Let $P = M/\sim$. Given $x, y, z \in P$, set $\mu_P(x, y, z)$ to be the equivalence class of $\mu(a, b, c)$, where a, b, c are representatives of x, y, z respectively. This is well defined, the quotient (P, μ_P) is a median algebra, and the quotient map is an epimorphism. Indeed any epimorphism of median algebras arises in this way.

We finish this section with the following proposition which will be applied to asymptotic cones of finite-rank coarse median spaces (see Lemma 6.6).

Proposition 2.4. *Let (M, μ) be a median algebra with $\text{rank}(M) \leq v$, and let ρ be a geodesic metric on M . Suppose that there is some $k \geq 1$ such that for all $a, b, c, d \in M$, we have $\rho(\mu(a, b, c), \mu(a, b, d)) \leq k\rho(c, d)$. Then there is a median metric λ on M , bi-Lipschitz equivalent to ρ , and which induces the median μ . Moreover, the bi-Lipschitz constants depend only on v and k .*

The second hypothesis asserts that the projection to intervals is uniformly Lipschitz. (It is precisely axiom (L2) of [Bowditch 2014a, Section 1].) It implies that the median operation is Lipschitz hence continuous (so M is a topological median algebra). In fact, we can weaken the geodesic condition to assert that M is Lipschitz path-connected, in the sense of axiom (L3) there. The proof is the same, but we won't need the more general statement here.

In the same work, it was shown that if (M, μ) is also finitely colourable, then it embeds in a finite product of trees, so the induced metric is median. The proof below amounts to observing that, under the weaker hypothesis of finite rank, the same construction gives a median metric directly.

Proof of Proposition 2.4. We write $\langle a, b:c \rangle_\lambda = \frac{1}{2}(\lambda(a, c) + \lambda(b, c) - \lambda(a, b))$ (i.e., the ‘‘Gromov product’’). Then $[a, b]_\lambda = \{x \in M \mid \langle a, b:x \rangle_\lambda = 0\}$. Thus, λ is a median metric inducing μ if and only if $[a, b]_\lambda = [a, b]$ for all $a, b \in M$. In this case, given any $a, b, c \in M$ we have $\rho(c, d) = \langle a, b:d \rangle_\lambda$, where $d = \mu(a, b, c)$ (see, for example, [Bowditch 2014a, Section 2]).

Now, if $\Pi \subseteq M$ is a finite subalgebra, we put a metric, λ_Π , on Π as in [Bowditch 2014a, Section 5]. (To each wall, W , of Π we associate a ‘‘width’’, $\lambda(W)$, and set $\lambda_\Pi(a, b) = \sum_W \lambda(W)$, as W ranges over the set, $\mathcal{W}(a, b)$, of walls separating a from b .) It is easily seen that $[a, b]_{\lambda_\Pi} = [a, b] \cap \Pi$, the latter being the intrinsic median interval in Π . (This holds since $\mathcal{W}(a, b) \subseteq \mathcal{W}(a, c) \cup \mathcal{W}(b, c)$, with equality if and only if $c \in [a, b] \cap \Pi$.) Therefore λ_Π is a median metric on Π .

Moreover, λ_Π is uniformly bi-Lipschitz equivalent to ρ restricted to Π . This follows as in [Bowditch 2014a, Section 5]. Note that if $x, y \in M$, then $[x, y] \cap \Pi$ has rank at most v . It follows by Dilworth's lemma that $[x, y] \cap \Pi$ is v -colourable (Lemma 2.3 of the same work) and so embeddable as a subalgebra of the cube

$[0, 1]^v$ (Proposition 1.4 there). Lemmas 5.2 and 5.4 there then respectively give us lower and upper bounds on $\lambda(x, y)$ in terms of $\rho(x, y)$. (Note that, in the notation of that work, $T \leq \rho(x, y)$, if we assume that M is a geodesic space.)

As in Section 6 of [Bowditch 2014a], we note that the set of finite subalgebras of M , ordered by inclusion, is cofinal in the set of all finite subsets of M . Therefore, by Tychonoff's theorem, we find a cofinal set of finite subalgebras, Π , so that $\lambda_\Pi(a, b) \rightarrow \lambda(a, b)$ for all $a, b \in M$, where λ is a metric on M , bi-Lipschitz equivalent to ρ .

To see that λ is a median metric inducing ρ , we need to check that $[a, b]_\lambda = [a, b]$. To this end, suppose $a, b, c \in M$ and let $d = \mu(a, b, c)$. Note that $a, b, c, d \in \Pi$ for a cofinal subset of those Π in our cofinal set of subalgebras. If $c \in [a, b]$, then $\langle a, b:c \rangle_{\lambda_\Pi} = 0$ for all such Π , so $\langle a, b:c \rangle_\lambda = 0$, so $c \in [a, b]_\lambda$. Conversely, if $c \in [a, b]_\lambda$, then $\lambda_\Pi(c, d) = \langle a, b:c \rangle_{\lambda_\Pi} \rightarrow 0$, so $\lambda(c, d) = 0$, so $c = d$, so $c \in [a, b]$. \square

3. Blocks

In this section, we describe top-dimensional cubes in median metric spaces.

Let M be a proper median metric space. Throughout this section, we will use v to denote $\text{rank}(M)$.

Definition. An n -block in M is a convex subset isometric to an l^1 -product of n nontrivial compact real intervals.

This is equivalent to saying that it is convex and median-isomorphic to $[-1, 1]^n$. Clearly, $n \leq v$.

We write $P \equiv \prod_{i=1}^n I_i$, where each I_i is a compact real interval, and can be identified with a 1-face of P .

Let $Q(P)$ be the set of corners of P , that is, $Q(P) = \prod_i \{a_i, b_i\}$ where $I_i = [a_i, b_i]$. It is clear that $Q(P)$ is intrinsically an n -cube in P , hence an n -cube in M . We see $P = \text{hull}(Q(P))$. In fact, $P = [a, b]$, where a, b are any pair of opposite corners of Q .

Lemma 3.1. *Let M be a proper median metric space of rank v . The following are equivalent for a subset $P \subseteq M$:*

- (1) P is v -block.
- (2) P is the convex hull of a v -cube in M .
- (3) P is isometric to a v -dimensional l^1 -cube.

Proof. The fact that (2) implies (1) was proven in [Bowditch 2016c, Proposition 5.6]. Suppose (3) holds. Let a, b be opposite corners of P (defined intrinsically). Directly from the definition of intervals in M , we can see that $P \subseteq [a, b]$, and so $P \subseteq \text{hull}(Q)$, where Q is the set of corners of P . By the observation preceding the lemma, we

know that $\text{hull}(Q)$ is a ν -block, and it now follows easily that we must have $P = \text{hull}(Q)$. \square

In (3) here, we are assuming that P is isometric to an l^1 -cube in the induced metric. We suspect that it would be sufficient to assume that this were the case for the induced path-metric. We will show this to be the case under some regularity assumptions (see Lemma 3.4 below).

Lemma 3.2. *Let M be a proper median metric space of rank ν . Suppose that $P, P' \subseteq M$ are ν -blocks, and that $P \cap P'$ is a common codimension-1 face. Then $P \cup P'$ is also a ν -block.*

Proof. Let $R_0 = Q(P \cap P') = Q(P) \cap Q(P')$. Let $R = Q(P) \setminus R_0$ and $R' = Q(P') \setminus R_0$. Thus R_0, R, R' are parallel $(\nu - 1)$ -cubes. In particular, $R \cup R'$ is a ν -cube. Let $P'' = \text{hull}(R \cup R')$. By Lemma 3.1, this is a ν -block. We claim that $R_0 \subseteq P''$. For if $r_0 \in R_0$, let $r \in R$ and $r' \in R'$ be adjacent vertices of $Q(P)$ and $Q(P')$ respectively. Thus, $[r_0, r]$ and $[r_0, r']$ are 1-faces of $Q(P)$ and $Q(P')$. In particular, $[r_0, r] \cap [r_0, r'] = \{r_0\}$ and so $r_0 \in [r, r'] \subseteq P''$ as claimed. It now follows that $P \cup P' = P''$. \square

More generally, if P, P' are any two blocks, then so is $P \cap P'$ provided it is nonempty. In fact, $P \cap P' = \text{hull}(Q)$, where Q is the projection (image of the gate map) of $Q(P')$ to P . In particular, $Q(P \cap P') \subseteq \langle Q(P) \cup Q(P') \rangle$.

We have the following procedure for subdividing blocks. Suppose $P \equiv \prod_{i=1}^n I_i$. If $F_i \subseteq I_i$ are finite subsets containing the endpoints, then $F = \prod_{i=1}^n F_i$ is a finite subalgebra of P . In fact, any finite subalgebra of P containing Q has this form. We can represent P as an l^1 -cube complex whose vertex set is exactly F . We refer to this as a *subdivision* of P .

Lemma 3.3. *Suppose that \mathcal{P} is a finite set of blocks in M . Then we can subdivide these blocks to find another set of blocks, \mathcal{P}' , with $\bigcup \mathcal{P} = \bigcup \mathcal{P}'$ such that any two blocks of \mathcal{P}' meet, if at all, in a common face.*

Proof. Let $A = \bigcup_{P \in \mathcal{P}} Q(P)$ and let $\Pi = \langle A \rangle$. If $P \in \mathcal{P}$, then $P \cap \Pi$ is a subalgebra of P containing $Q(P)$ and so determines a subdivision of P . We subdivide each element of \mathcal{P} in this way to give us our new collection \mathcal{P}' . Now if $P, P' \in \mathcal{P}'$, then $Q(P \cap P') \subseteq \langle Q(P) \cup Q(P') \rangle \subseteq \Pi$. But by construction, $P \cap \Pi \subseteq Q(P)$ and $P' \cap \Pi \subseteq Q(P')$, so $Q(P \cap P') \subseteq P \cap P' \cap \Pi \subseteq Q(P) \cap Q(P')$. It now follows that $P \cap P'$ is a common face of P and P' as claimed. \square

In other words, we can realise $\bigcup \mathcal{P}$ as an l^1 -cube complex in M all of whose cells are blocks.

Definition. A *straight cube complex* in M is an embedding of a locally finite cube complex in M such that each cell is a block (necessarily of the corresponding dimension).

The following is equivalent to the informal definition of “cubulated” given in Section 1.

Definition. A *cubulated set* is a subset of M which is a locally finite union of blocks.

A cubulated set, Φ , is clearly closed, and by the above, we see that any point $x \in \Phi$ has a neighbourhood in Φ which is a straight cube complex contained in Φ . In fact, we can assume that x is a vertex of this cube complex. Note also that a finite union or a finite intersection of cubulated sets is also cubulated.

In fact, if Φ_1, \dots, Φ_n is a finite set of cubulated sets, with $x \in \bigcap_i \Phi_i$, then we can find a straight cube complex, $\Upsilon \subseteq \bigcup_i \Phi_i$ as above, with each $\Upsilon \cap \Phi_i$ a subcomplex of Υ . (This is a consequence of the construction of Lemma 3.3.)

Lemma 3.4. *Suppose that $\Phi \subseteq M$ is cubulated. Suppose that $P \subseteq \Phi$ is isometric to a ν -dimensional l^1 -cube in the path-metric induced from ρ . Then P is a ν -block in M .*

Proof. By Lemma 3.3 we can find a straight cube complex $\Upsilon \subseteq \Phi$, with $P \subseteq \Upsilon$. We can assume that the intrinsic corners of P are all vertices of Υ . It now follows that P is a union of ν -blocks of M , which are ν -cells of Υ . These determine a subdivision of P in the induced path metric on P . Applying Lemma 3.2 inductively, we see that P is a block in M . \square

Definition. Suppose that $\Phi \subseteq M$ is cubulated. We say that a point $x \in \Phi$ is *regular* if it has a neighbourhood in Φ which is a ν -block in M . Otherwise, we say that x is *singular*. We write Φ_S for the set of singular points of Φ .

Note that Φ_S is a cubulated set of dimension at most $\nu - 1$.

Suppose now that Φ is cubulated and homeomorphic to \mathbb{R}^ν . If $K \subseteq \Phi$ is compact, then K lies inside a straight cube complex, Υ , in Φ . Moreover, we can assume that any $(\nu - 1)$ -cell of Υ meeting K lies in exactly two ν -cells of Υ . By Lemma 3.2, the union of these two cells is also a ν -block in M . From this, we deduce:

Lemma 3.5. *Suppose that $\Phi \subseteq M$ is cubulated and homeomorphic to \mathbb{R}^ν . Then Φ_S is a cubulated set of dimension at most $\nu - 2$.*

Note that, if P is any block in Φ , then the relative interior of P in Φ is exactly the intrinsic relative interior of P , as defined earlier.

Definition. A *leaf segment* of Φ is a closed subset, L , of Φ homeomorphic to a real interval such that if $x \in L$, then there is a block $P \subseteq \Phi$ containing x in its relative interior, with $L \cap P$ lying in a coordinate line of P . If the real interval is the whole real line, we refer to L as a *leaf*.

Clearly this implies that $L \cap \Phi_S = \emptyset$. We note:

Lemma 3.6. *Every leaf segment of Φ is convex in M .*

Proof. Let $L \subseteq \Phi$ be a leaf segment, and suppose $I \subseteq L$ is a compact subinterval. Since $I \cap \Phi_S = \emptyset$, we can find a subset $P \subseteq \Phi$ which is a block in the intrinsic path metric on P , and with $I \subseteq P$ an intrinsic coordinate line with respect to that structure. But by Lemma 3.4, P is a block in M , and so I is convex. It now follows that L is convex. \square

Definition. A *flat* in M is a closed convex subset isometric to \mathbb{R}^ν with the l^1 metric.

(Note that we always take a flat to be of maximal dimension; that is, $\nu = \text{rank}(M)$.)

In fact (as with blocks), we see that any closed subset of M which is isometric to \mathbb{R}^ν in the induced metric is flat. (Indeed, we suspect this remains true if we substituted “induced path-metric” for “induced metric” in the above.) Also, any closed convex subset of M median isomorphic to \mathbb{R}^ν , with the standard product structure, is a flat. In particular, the notion depends only on the topology and median structure.

Clearly a flat is a cubulated set with empty singular set. Conversely, we have:

Lemma 3.7. *Suppose that $\Phi \subseteq M$ is a cubulated set homeomorphic to \mathbb{R}^ν , and with $\Phi_S = \emptyset$. Then Φ is a flat.*

Proof. First note that, in the intrinsic path metric, Φ is locally isometric to \mathbb{R}^ν in the l^1 metric. Since it is complete, it must be globally isometric. By Lemma 3.4 any subset of Φ that is intrinsically a block is indeed a block in M , and so, in particular, convex. Since any two points of Φ are contained in such a subset, it follows that Φ is convex. The induced path metric is therefore the same as the induced metric. \square

Here is a criterion for recognising that a cubulated set is indeed nonsingular:

Lemma 3.8. *Suppose that $\Phi \subseteq M$ is cubulated, and that there is a homeomorphism $f : \mathbb{R}^\nu \rightarrow \Phi$ such that if $H \subseteq \mathbb{R}^\nu$ is any codimension-1 coordinate plane in \mathbb{R}^ν then $f(H)$ is cubulated. Then Φ is a flat, and f is a median isomorphism.*

Proof. Suppose that $L \subseteq \mathbb{R}^\nu$ is a coordinate line, and that $x \in L$ with $f(x) \notin \Phi_S$. Let H_1, H_2, \dots, H_n be the codimension-1 coordinate planes through x , with $L = \bigcap_{i=2}^n H_i$, and with H_1 orthogonal to L . As noted after Lemma 3.3, we can find a neighbourhood, Υ , of $f(x)$ in Φ , which is a straight cube complex, with $f(x)$ a vertex, and each $f(H_i) \cap \Upsilon$ a subcomplex of Υ . In particular, $f(L) = \bigcap_{i=2}^n f(H_i)$ is a 1-dimensional subcomplex, and so meets $f(x)$ in a pair of 1-cells of Υ . Let Δ be the link of $f(x)$ in Υ . Since $f(x) \notin \Phi_S$, this is a cross polytope. Note that $f(L)$ determines two vertices, p, q , of Δ . Now $f(H_1)$ separates the two rays of $f(L)$ with basepoint $f(x)$ in Φ . It therefore determines a subcomplex of Δ separating p from q in Δ . It follows that p and q must be opposite vertices of Δ . We see that the union of the two 1-cells of $f(L)$ meeting x is convex.

In summary, we have shown that, away from Φ_S , the images of coordinate lines are locally convex, that is, leaf segments of Φ . By a simple compactness argument,

it now follows that if $I \subseteq \mathbb{R}^v$ is a compact interval lying in a coordinate line with $f(I) \cap \Phi_S = \emptyset$, then $f(I)$ is a leaf segment of Φ . We can now deduce that if $P \subseteq \Phi$ is any v -block in $\Phi \setminus \Phi_S$, then $f^{-1}|_P$ is a median isomorphism to a block $f^{-1}(P)$ in \mathbb{R}^v . In fact, it is enough that P should not meet Φ_S in its relative interior.

Suppose now that $y \in \Phi$. Let $\Upsilon \subseteq \Phi$ be a straight cube complex that is a neighbourhood of y , and with y as a vertex. To simplify notation, suppose that $f^{-1}(y)$ is the origin in \mathbb{R}^v . Let P be a cube of Υ with y a corner of P . Then $f^{-1}P$ has the form $\prod_{i=1}^v [0, \pm t_i]$ for some $t_i > 0$. Since there are only finitely many such P , after shrinking them, we can assume that the preimages all have the form $\prod_{i=1}^v [0, \pm t]$ for some $t > 0$. We see that there are exactly 2^v such cubes, which fit together into a bigger cube of the form $f([-t, t]^v)$. In particular, the link of y in Υ is a cross polytope, and y is regular.

We have shown that $\Phi_S = \emptyset$, and so by Lemma 3.7, Φ is a flat. \square

For reference elsewhere (see Proposition 4.8 below) we note that there is a variation on Lemma 3.8, where \mathbb{R}^v is replaced by a real cube, $[-1, 1]^v$, and $\Phi = f([-1, 1]^v)$. In fact, it is enough to assume that the relative interior, $f((-1, 1)^v) \subseteq \Phi$, is cubulated (in the sense that any compact subset of $f((-1, 1)^v)$ lies inside another compact subset of Φ which is cubulated). Also, we only need to consider coordinate planes restricted to the interior of Φ . The argument is essentially the same, this time applied to the relative interior of Φ and then taking the closure.

4. Cubulating planes

In this section, we discuss the regularity of “top-dimensional manifolds” in M . These play an important role in [Kleiner and Leeb 1997; Kapovich, Kleiner and Leeb 1998; Behrstock, Kleiner, Minsky and Mosher 2012]. Our argument is analogous to those to be found there, though set in a somewhat different context. Here, we interpret this in terms of cubulations. We will only use dimension of (locally) compact sets, so all the standard definitions are equivalent. For definiteness, we can interpret the dimension of a topological space to be its covering dimension. (Note that this differs from the notion of “topological rank” used in [Kleiner and Leeb 1997].)

Suppose that (M, ρ) is a complete median metric space. We first note:

Lemma 4.1. *Any locally compact subset of M has topological dimension at most $\text{rank}(M)$.*

Proof. First note that by Lemma 2.1, M is locally convex. The statement then follows by Theorem 2.2 and Lemma 7.6 of [Bowditch 2013]. \square

From this we see that if M is homeomorphic to \mathbb{R}^v , then $v = \text{rank}(M)$. (The fact that $v \geq \text{rank}(M)$ is an immediate consequence of Lemma 4.1. For the other

direction, note that by Lemma 2.3, any n -cube in M is the vertex set of an embedded l^1 -cube in M , and so $n \leq \nu$, and it follows that $\text{rank}(M) \leq \nu$.)

In fact, we can say a lot more about the regularity of such a space:

Lemma 4.2. *If M is a complete median metric space homeomorphic to \mathbb{R}^ν , then M is cubulated.*

In particular, we see that M is locally isometric to \mathbb{R}^ν with the l^1 metric away from a cubulated singular set of dimension at most $\nu - 2$. (Note that we are not claiming that the cubulation is combinatorial in the sense of PL manifolds. Certainly the link of any cell in the cubulation will be a homology sphere. It is not clear whether it need be a topological sphere in this situation.)

Proof of Lemma 4.2. Let $B_1 \subseteq B_0$ be topological ν -balls in M . We suppose that $N(B_1; 2u) \subseteq B_0$, where $N(\cdot; r)$ denotes the metric r -neighbourhood with respect to the metric ρ . Let $0 < s < t < u$ be sufficiently small depending on u , as described below. We take a topological triangulation of ∂B_0 , all of whose simplices have diameter at most s . Let $A \subseteq \partial B_0 \subseteq M$ be the set of vertices of this triangulation, and let $\Pi = \langle A \rangle \subseteq M$. By Lemma 2.3, Π is the vertex set of an l^1 -cube complex Υ embedded in M . We extend the inclusion of A into Υ to a continuous map $f : \partial B_0 \rightarrow \Upsilon$. Provided s is small enough in relation to t , we can arrange that the ρ -diameter of the image of each simplex is at most t . (For example, take the corresponding euclidean metric, σ_Υ , on Υ . Then $(\Upsilon, \sigma_\Upsilon)$ is CAT(0), and we can map in simplices, inductively on the 1-skeleta by taking geodesic rulings. In this way the σ_Υ -diameter of the image of any simplex is at most s . Now $\rho \leq \sigma_\Upsilon \sqrt{\nu}$, so this works provided $s\sqrt{\nu} \leq t$.) Now $\rho(x, f(x)) \leq s + t$ for all ∂B_0 . Again, provided t is small enough in relation to u , we can find a homotopy, $F : \partial B_0 \times [0, 1] \rightarrow M$, between f and the inclusion of B_0 into M whose trajectories all have length at most u . In particular, the image of the homotopy lies in $N(\partial B_0; u)$ and is therefore disjoint from B_1 . For this, it is convenient to take the CAT(0) metric, σ , on M , as given by Theorem 2.2. We can then use linear isotopy in this metric, that is, the trajectory from x to $f(x)$ is the σ -geodesic segment. Again we note that $\rho \leq \sigma \sqrt{\nu}$, so this works provided $(s + t)\sqrt{\nu} \leq u$.

Now Υ is a CAT(0) complex in the euclidean metric, and so in particular is contractible. We can therefore extend $f : \partial B_0 \rightarrow \Upsilon$ arbitrarily to a continuous map $f : B_0 \rightarrow \Upsilon$. We combine this with the homotopy constructed above to give a continuous map $g : B_0 \rightarrow M$ which restricts to inclusion on ∂B_0 . More formally, if $x \in B_0 \setminus \{0\}$, write $x = \lambda \hat{x}$, where $\lambda \in (0, 1]$ and $\hat{x} \in \partial B_0$ (via any homeomorphism of B_0 with the unit ball in \mathbb{R}^ν). If $\lambda \leq \frac{1}{2}$, then set $g(x) = f(2\lambda \hat{x})$. If $\lambda \geq \frac{1}{2}$, then set $g(x) = F(\hat{x}, 2\lambda - 1)$. We set $g(0) = 0$. Note that $g(B_0) = f(B_0) \cup \text{image}(F)$, and we have noted that $B_1 \cap \text{image}(F) = \emptyset$, and so $B_1 \cap g(B_0) \subseteq f(B_0)$. But now, $B_0 \subseteq g(B_0)$ (since $g|_{\partial B_0}$ is just inclusion). It therefore follows that $B_1 \subseteq f(B_0) \subseteq \Upsilon$.

We do not know a priori that Υ is a straight complex. However, every ν -cell of Υ must be a ν -block. Moreover, B_1 must lie in the union of these ν -cells. (For if $x \in B_1$, then any cell of Υ must lie in a ν -cell, otherwise some neighbourhood of x in B_1 would have dimension at most $\nu - 1$.) Since B_1 was an arbitrary ν -ball in M , we see that every compact subset of M lies in a finite union of ν -blocks of M . It follows that M is cubulated. \square

We can give a more general version of this for subsets of a proper median metric space as follows (given as Proposition 1.2 in Section 1).

Proposition 4.3. *Suppose that M is a complete median metric space of rank at most ν , and that $\Phi \subseteq M$ is a closed subset homeomorphic to \mathbb{R}^ν . Then Φ is cubulated.*

Clearly, in this case, the rank will be exactly ν . As before, we see that Φ is locally isometric to \mathbb{R}^ν in the l^1 metric away from a codimension-2 singular set (see Lemma 3.5).

Note that there is no loss in assuming that M is connected (hence “proper” in the terminology of Section 3) since we can simply restrict to the component containing Φ . We have already observed in Section 2 that this is convex, hence intrinsically a complete median metric space.

For the proof, will need the following two topological lemmas:

Lemma 4.4. *Suppose that X is a Hausdorff topological space and that $B, P \subseteq X$ are embedded topological n -balls, with intrinsic boundary spheres $S(B)$ and $S(P)$ respectively. Suppose that $P \setminus S(P)$ is open in X , that $P \cap S(B) = \emptyset$ and that $B \cap P \setminus S(P) \neq \emptyset$. Then $P \subseteq B$.*

Proof. Write $I(B) = B \setminus S(B)$ and $I(P) = P \setminus S(P)$ for the relative interiors. These are both homeomorphic to \mathbb{R}^n . Let $U = I(P) \cap B = I(P) \cap I(B)$. By assumption, $U \neq \emptyset$. Now $I(P)$ is open in X , so U is open in $I(B)$. Thus, U is homeomorphic to an open subset of \mathbb{R}^n , hence, by invariance of domain, it is also open in $I(P)$. But $U = I(P) \cap B$, so U is also closed in $I(P)$, and so, by connectedness, $U = I(P)$. In other words, $I(P) \subseteq I(B)$, and it follows that $P \subseteq B$ as claimed. \square

For the second topological lemma, we need the following definition.

Definition. The (locally) compact dimension of a Hausdorff topological space is the maximal topological dimension of any (locally) compact subset.

Clearly the compact dimension is at most the locally compact dimension, which in turn is at most the “separation dimension” as defined in [Bowditch 2013, Section 7].

Lemma 4.5. *Suppose that M is a Hausdorff topological space of compact dimension at most ν . Suppose that B is a topological ν -ball with boundary ∂B . Suppose that $f_0, f_1 : B \rightarrow M$ are continuous and homotopic relative to ∂B , and that f_0 is injective. Then $f_0(B) \subseteq f_1(B)$.*

The proof is based on an argument in Section 6.1 of [Kleiner and Leeb 1997]. A related, but slightly different statement can be found in [Behrstock, Kleiner, Minsky and Mosher 2012, Section 6]. In what follows, H_r will denote Čech homology with coefficients in a field (say \mathbb{Z}_2 to be specific). We will only deal with compact spaces, so that the usual homology axioms, in particular, homotopy, excision and exactness, hold. We need compact spaces and field coefficients for exactness; see [Eilenberg and Steenrod 1952, Chapter IX]. (Note that in [Kleiner and Leeb 1997], it is implicit from context that singular homology is being used. As a consequence they use open sets instead of compact sets.) Note that, if K is compact and of dimension at most ν , then $H_n(K, A)$ is trivial for any compact $A \subseteq K$ and any $n > \nu$.

Proof. Let $C = f_0(B)$, $D = f_1(B)$, $S = f_0(\partial B) = f_1(\partial B)$ and let $E \subseteq M$ be the image of a homotopy from f_0 to f_1 . Thus, $S \subseteq C \cap D \subseteq C \cup D \subseteq E$ are all compact. Suppose, for contradiction, that $p \in C \setminus D$. Let $N \subseteq C$ be an open neighbourhood of p in C , whose closure is homeomorphic to a closed ν -ball disjoint from D . Now $H_\nu(C, C \setminus N) \cong H_{\nu-1}(S) \cong \mathbb{Z}_2$, but the image of $H_\nu(C, C \setminus N)$ in $H_\nu(E, C \cup D \setminus N)$ is trivial. (Note that this corresponds to the image of $H_{\nu-1}(\partial B)$ under that map induced by $f_1 \simeq f_0$.) Now the natural map $H_\nu(C, C \setminus N) \rightarrow H_\nu(C \cup D, C \cup D \setminus N)$ is an isomorphism, by excision. Also, since $H_{\nu+1}(E, C \cup D)$ is trivial, the exact sequence of triples tells us that the natural map, $H_\nu(C \cup D, C \cup D \setminus N) \rightarrow H_\nu(E, C \cup D \setminus N)$ is injective. Composing, we get that the natural map $H_\nu(C, C \setminus N) \rightarrow H_\nu(E, C \cup D \setminus N)$ is injective, giving a contradiction. \square

We can now give the proof of Proposition 4.3. We have already observed that we can assume M to be connected. We recall that M is contractible (see Theorem 2.2), and has locally compact dimension at most ν (Lemma 4.1).

Proof of Proposition 4.3. This is an extension of the argument for Lemma 4.2. This time, we take three closed topological balls, $B_2 \subseteq B_1 \subseteq B_0 \subseteq \Phi \subseteq M$. We assume that B_2 is contained in the relative interior of B_1 , and that $N(B_1; 2u) \subseteq B_0$ (in the metric ρ on M). We start as before, triangulating ∂B_0 , to give us a complex $\Upsilon \subseteq M$, a continuous map $f : B_0 \rightarrow \Upsilon$, and a homotopy in M from $f|_{\partial B_0}$ to the inclusion of ∂B_0 . We can arrange that the homotopy does not meet B_1 . We combine f with this homotopy to give a continuous map, $g : B_0 \rightarrow M$, which restricts to the identity on ∂B_0 .

Since M is contractible, g is homotopic to the inclusion of B_0 in M , relative to ∂B_0 . Therefore, Lemma 4.5 tells us that $B_0 \subseteq g(B_0)$. Moreover, as observed above, the homotopy part of g does not meet B_1 and so we see that $B_1 \subseteq f(B_0) \subseteq \Upsilon$.

In summary, we have $B_2 \subseteq B_1 \subseteq \Upsilon$. After subdividing, we can suppose that any cell of Υ meeting B_2 is disjoint from the spherical boundary, $S(B_1)$, of B_1 . Let \mathcal{P} be the set of ν -cells of Υ meeting B_2 in their relative interiors. Each of these is a

ν -block, and by the same dimension argument as in the proof of Lemma 4.2, we have $B_2 \subseteq \bigcup \mathcal{P}$. We claim that $\bigcup \mathcal{P} \subseteq \Phi$.

In fact suppose that $P \in \mathcal{P}$. We apply Lemma 4.4 with $X = \Upsilon$, $B = B_1$. Since Υ is a complex of dimension ν , we have $P \setminus S(P)$ open in Υ . Also, $P \setminus S(B_1) = \emptyset$, and by assumption $B_2 \cap P \setminus S(P) \subseteq B_1 \cap P \setminus S(P)$ is nonempty. It follows that $P \subseteq B_1$, so in particular, $P \subseteq \Phi$.

Since B_2 can be chosen arbitrarily, we see that any compact subset of Φ is contained in a finite union of ν -blocks contained in Φ , and so Φ is cubulated as required. \square

Remark. In fact, the argument shows that if $B \subseteq M$ is homeomorphic to a closed ν -ball, and $K \subseteq B \setminus \partial B$ is a compact subset of the relative interior, then there is a compact cubulated set, Υ , with $K \subseteq \Upsilon \subseteq B$.

Combining Proposition 4.3 and Lemma 3.8, we get:

Proposition 4.6. *Suppose that M is a complete median metric space, and that $\Phi \subseteq M$ is a closed subset and that there is a homeomorphism $f : \mathbb{R}^\nu \rightarrow \Phi$ with the following property. For each codimension-1 coordinate plane, $H \subseteq \mathbb{R}^\nu$, there is a closed subset, $\Psi \subseteq M$, homeomorphic to \mathbb{R}^ν such that $f(H) = \Phi \cap \Psi$. Then Φ is a flat, and f is a median isomorphism.*

Note that the hypotheses on Φ only depend on the topological structure of M .

We conclude, in particular, that Φ is isometric to \mathbb{R}^ν with the l^1 metric.

This is all we will need for our discussion of the marking graph. We also include the following results which will be relevant to applications elsewhere; see [Bowditch 2015; 2016a].

Definition. We say that an \mathbb{R} -tree is *furry* if every point has valence at least 3.

Proposition 4.7. *Suppose that M is a complete median metric space of rank ν , that D is a direct product of ν furry \mathbb{R} -trees, and that $f : D \rightarrow M$ is a continuous injective map with closed image. Then f is a median homomorphism. Moreover, $f(D)$ is convex.*

Proof. By a *product flat* in D we mean a direct product of bi-infinite geodesics in each of the factors. If every point in each factor has valence at least 4 (as in the cases of genuine interest) then we see that every product flat Φ satisfies the hypotheses of Proposition 4.6, and so $f|_\Phi$ is a median homomorphism. Now any two points, a, b lie in some such product flat Φ and $[a, b] \subseteq \Phi$. Thus, if $c \in [a, b]$, then $fc \in [fa, fb]$, and it follows that f is a median homomorphism on all of D .

If we allow for vertices of valence 3, then we just note that any codimension-1 coordinate plane in Φ is the intersection of three product flats, hence cubulated. We can then apply Lemma 3.8 directly, to see that f is a median homomorphism on Φ , hence, as above, everywhere. \square

We remark that Proposition 4.7 applies in particular if M is also a product of ν \mathbb{R} -trees. It follows that f splits as a direct product of embeddings, up to permutation of the factors. Some further discussion of this, with applications, can be found in [Bowditch 2016b].

Definition. A *tree product*, T , in M is a convex subset median isomorphic to a direct product of ν nontrivial rank-1 median algebras. It is *maximal* if it is not contained in any strictly larger tree product.

Note that T is an l^1 -product of \mathbb{R} -trees. It is easily seen that the closure of a tree product is a tree product, and so any maximal tree product is closed.

Note that in the above terminology, any closed subset of M homeomorphic to a direct product of ν furry \mathbb{R} -trees for $\nu \geq 2$ is a tree product by Proposition 4.7.

For applications elsewhere, in particular in [Bowditch 2016a], we note that we can relax the “furriness” condition somewhat.

Definition. An \mathbb{R} -tree is *almost furry* if it is infinite, and no point has valence equal to 2.

In this case, by removing the extreme (valence-1) points, we obtain the maximal furry subtree. The following was given as Proposition 1.3 in the introduction.

Proposition 4.8. *Suppose that M is a complete median metric space of rank ν , that D is a direct product of ν 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. Moreover, $f(D)$ is convex.*

Proof. We can apply the arguments to the maximal subset which is a product of furry trees, and then take its closure. We have already observed that the key statements, in particular Lemma 3.8 and Proposition 4.3, have local versions which can be applied to this case. \square

5. Ultraproducts

In this section, we give some general background to the theory of ultraproducts and asymptotic cones. The notion of an asymptotic cone was introduced in [van den Dries and Wilkie 1984]; see also [Gromov 1993]. The idea behind this is to keep rescaling the metric so that points move closer and closer together, and then pass to an “ultralimit” of the resulting spaces. (Here, the term “ultralimit” is used in the sense of [Gromov 1993], rather than in the usual sense of model theory.) We then factor out “infinitesimals” to give what we call here an “extended asymptotic cone”. If we also throw away the “unlimited” parts (beyond infinity), we get the usual asymptotic cone. In principle, this may depend on the choice of rescaling factors and (if the continuum hypothesis fails) on the choice of ultrafilter, but such ambiguity will not matter to us here.

Let \mathcal{Z} be a countable set equipped with a nonprincipal ultrafilter. We can think of this as a finitely additive measure on \mathcal{Z} , taking values in $\{0, 1\}$, such that \mathcal{Z} itself has measure 1, and any finite subset of \mathcal{Z} has measure 0. If a predicate, $P(\zeta)$, depends on $\zeta \in \mathcal{Z}$, we say that P holds *almost always* if the set of ζ for which it holds has measure 1.

We refer to a sequence of objects indexed by \mathcal{Z} as a \mathcal{Z} -sequence. Typically, we will use the notation $\vec{X} = (X_\zeta)_\zeta$ for such a sequence. If these are all sets, we write $\prod \vec{X} = \prod_\zeta X_\zeta$ for their product. Given $\vec{x}, \vec{y} \in \prod \vec{X}$, we write $\vec{x} \approx \vec{y}$ to mean that $x_\zeta = y_\zeta$ almost always. Thus, \approx is an equivalence relation on $\prod \vec{X}$, and we write $\mathcal{U}\vec{X} = \prod \vec{X} / \approx$ for the quotient.

Definition. We refer to $\mathcal{U}\vec{X}$ as the *ultraproduct* of the \mathcal{Z} -sequence \vec{X} .

Note that we only need to have x_ζ defined almost always to determine an element of $\mathcal{U}\vec{X}$. We write $\mathbf{x} = [\vec{x}]$ for this element.

We write $\mathcal{P}(\vec{X})$ for the \mathcal{Z} -sequence $(\mathcal{P}(X_\zeta))_\zeta$, where \mathcal{P} denotes power set. There is a natural map $\mathcal{U}\mathcal{P}(\vec{X}) \rightarrow \mathcal{P}(\mathcal{U}\vec{X})$, defined by sending \vec{Y} to the set of $\mathbf{x} = [\vec{x}] \in \mathcal{U}\vec{X}$ such that $x_\zeta \in Y_\zeta$ almost always. We can identify the image of this map with $\mathcal{U}\vec{Y}$. Note that we can define unions and intersections in $\mathcal{P}(\vec{X})$ (by taking unions and intersections on each ζ -coordinate). These operations are respected by the above map.

Given two \mathcal{Z} -sequences of sets, \vec{X} and \vec{Y} , we can form the direct product $\vec{X} \times \vec{Y}$ as $(X_\zeta \times Y_\zeta)_\zeta$, and we see that $\mathcal{U}(\vec{X} \times \vec{Y})$ is naturally identified with $\mathcal{U}\vec{X} \times \mathcal{U}\vec{Y}$. A \mathcal{Z} -sequence of relations on $X_\zeta \times Y_\zeta$ give rise to a relation on $\mathcal{U}\vec{X} \times \mathcal{U}\vec{Y}$ via the map from $\mathcal{U}\mathcal{P}(\vec{X} \times \vec{Y})$ to $\mathcal{P}(\mathcal{U}\vec{X} \times \mathcal{U}\vec{Y})$. In other words, \mathbf{x} is related to \mathbf{y} if x_ζ is almost always related to y_ζ . A particular case is when relation on $X_\zeta \times Y_\zeta$ is almost always the graph of a function. In fact, the following is a simple exercise:

Lemma 5.1. *Given any \mathcal{Z} -sequence of functions, $f_\zeta : X_\zeta \rightarrow Y_\zeta$, there is a unique function $\mathcal{U}f : \mathcal{U}\vec{X} \rightarrow \mathcal{U}\vec{Y}$, such that $\mathbf{y} = \mathcal{U}f(\mathbf{x})$ if and only if $y_\zeta = f_\zeta(x_\zeta)$ almost always.*

We also note that the discussion of relations also applies to finite products of sets, and so to n -ary relations and n -ary operations for any finite n . For example, if $\vec{\Gamma}$ is a sequence of groups, then $\mathcal{U}\vec{\Gamma}$ has the structure of a group. If each Γ_ζ acts on a set X_ζ , then $\mathcal{U}\vec{\Gamma}$ acts on $\mathcal{U}\vec{X}$.

Suppose that $X_\zeta = X$ is constant. In this case, we write $\mathcal{U}X = \mathcal{U}\vec{X}$.

Definition. We refer to $\mathcal{U}X$ as the *ultrapower* of the set X .

There is a natural injection X into $\mathcal{U}X$ obtained by taking constant sequences. We refer to the image of this map as the *standard part* of $\mathcal{U}X$. We usually identify X with the standard part of $\mathcal{U}X$. If X is finite, then $\mathcal{U}X$ is equal to its standard part.

Note that the ultrapower, $\mathcal{U}\mathbb{R}$, of the real numbers is an ordered field. We say that $\mathbf{x} \in \mathcal{U}\mathbb{R}$ is *limited* if $|\mathbf{x}| \leq \mathbf{y}$ for some $\mathbf{y} \in \mathbb{R} \subseteq \mathcal{U}\mathbb{R}$ (where $|\mathbf{x}| = \max\{\mathbf{x}, -\mathbf{x}\}$).

Otherwise it is *unlimited*. We say that \mathbf{x} is *infinitesimal* if $|\mathbf{x}| \leq \mathbf{y}$ for all positive standard \mathbf{y} . Note that 0 is the only standard infinitesimal, and that the nonzero infinitesimals are exactly the reciprocals of unlimited numbers.

There is a well defined map $\text{st} : \mathcal{U}\mathbb{R} \rightarrow [-\infty, \infty] = \mathbb{R} \cup \{-\infty, \infty\}$ such that $\text{st}(\mathbf{x}) = \infty$ if \mathbf{x} is positive unlimited, $\text{st}(\mathbf{x}) = -\infty$ if \mathbf{x} is negative unlimited, and $\mathbf{x} - \text{st}(\mathbf{x})$ is infinitesimal if \mathbf{x} is limited. We refer to $\text{st}(\mathbf{x})$ as the *standard part* of \mathbf{x} . We will usually restrict attention to nonnegative numbers, so we get a map $\text{st} : \mathcal{U}[0, \infty) \rightarrow [0, \infty]$. If $(x_\zeta)_\zeta$ is a \mathcal{Z} -sequence of real numbers, we write $x_\zeta \rightarrow x \in \mathbb{R} \cup \{\infty\}$ to mean that $x = \text{st}(\mathbf{x})$. (This is the same as taking limits in \mathbb{R} with respect to the ultrafilter.)

In the case of the natural number, there are no infinitesimals, and \mathbb{N} is an initial segment of $\mathcal{U}\mathbb{N}$. We get a map $\text{st} : \mathcal{U}\mathbb{N} \rightarrow \mathbb{N} \cup \{\infty\}$ which is the identity on \mathbb{N} .

Given any set M we define an *nonstandard metric* on M to be a metric with values in $\mathcal{U}\mathbb{R}$. In other words, it is a map $M^2 \rightarrow \mathcal{U}[0, \infty)$ satisfying the same axioms as a metric, except with \mathbb{R} replaced by $\mathcal{U}\mathbb{R}$. Note that, if σ is a nonstandard metric, the composition $\hat{\sigma} = \text{st} \circ \sigma : M^2 \rightarrow [0, \infty]$ is an idealised pseudometric on M . Here, we use the term *idealised* to mean that we allowing points to be an infinite distance apart. As with usual pseudometric spaces, we can take the Hausdorffification, \hat{M} , of M . In other words, given $x, y \in M$, we write $x \simeq y$ to mean that $\hat{\sigma}(x, y) = 0$. Thus \simeq is an equivalence relation on M , and we set $\hat{M} = M/\simeq$. The induced map, $\hat{\sigma} : \hat{M}^2 \rightarrow [0, \infty]$ is an idealised metric on \hat{M} . Note that the relation on \hat{M} given by deeming x to be equivalent to y if $\hat{\sigma}(x, y) < \infty$ is an equivalence relation.

Definition. A *component* of an idealised metric space, \hat{M} , is an equivalence class under the above relation.

Note that components are both open and closed in the topology induced on \hat{M} . Also any component is a metric space in the usual sense. Note that we can speak about an extended metric space as being *complete*; that is, all its components are complete metric spaces. (We will see that, in the cases of interest in this paper, this notion “component” will coincide with the usual notion of a connected component).

Suppose that $((X_\zeta, \sigma_\zeta))_\zeta$ is a \mathcal{Z} -sequence of metric spaces. This gives rise to a nonstandard metric, $\mathcal{U}\sigma$, on $\mathcal{U}\vec{X}$, and hence to an idealised pseudometric, $\hat{\sigma}$, on $\mathcal{U}\vec{X}$. Let \hat{X} be the Hausdorffification, with idealised metric $\hat{\sigma} : \hat{X}^2 \rightarrow [0, \infty]$.

If $x_\zeta \in X_\zeta$, we write $x_\zeta \rightarrow x$ to mean that $x \in \hat{X}$ is the image of the sequence \mathbf{x} under the natural maps. We think of x as the limit of the x_ζ . By construction, every sequence has a unique limit.

For the following lemma, we use the fact that \mathcal{Z} is countable to find a \mathcal{Z} -sequence $(n_\zeta)_\zeta$ in \mathbb{N} with $n_\zeta \rightarrow \infty$. For example let, $n : \mathcal{Z} \rightarrow \mathbb{N}$ be any injective map. (This is true of a broader class of cardinals than \aleph_0 , though we won’t pursue that issue here — we will have no need of any uncountable indexing sets.)

Lemma 5.2. $(\hat{X}, \hat{\sigma})$ is complete.

Proof. Let $(x^i)_{i \in \mathbb{N}}$ be a Cauchy sequence in \hat{X} . It is enough to show that $(x^i)_i$ has a convergent subsequence. We can suppose that $\hat{\sigma}(x^i, x^{i+1}) \leq 1/2^{i+1}$ for all i . Given $i \in \mathbb{N}$, let $(x_\zeta^i)_\zeta$ be some representative of x^i in $\vec{X} = (X_\zeta)_\zeta$. Let $\mathcal{Z}_0(j) = \{\zeta \in \mathcal{Z} \mid \sigma_\zeta(x_\zeta^j, x_\zeta^{j+1}) \leq 1/2^j\}$, let $\mathcal{Z}(i) = \bigcap_{j \leq i} \mathcal{Z}_0(j)$ and $\mathcal{Z}(\infty) = \bigcap_{j=0}^\infty \mathcal{Z}_0(j)$. Given $\zeta \in \mathcal{Z} \setminus \mathcal{Z}(\infty)$, let $i(\zeta) = \max\{i \mid \zeta \in \mathcal{Z}(i)\}$, and let $y_\zeta = x_\zeta^{i(\zeta)}$. Note that if $\zeta \in \mathcal{Z}(i)$, then $\sigma_\zeta(y_\zeta, x_\zeta^i) \leq \sum_{j \geq i} (1/2^j) \leq 2/2^i$. We distinguish two cases.

If $\mathcal{Z}(\infty)$ has measure 0, then y_ζ is defined almost always. Let y be the image of $(y_\zeta)_\zeta$ in \hat{X} . Now $\sigma_\zeta(y_\zeta, x_\zeta^i) \leq 2/2^i$ almost always, and so $\hat{\sigma}(y, x^i) \leq 2/2^i$, showing that x^i converges to y .

If $\mathcal{Z}(\infty)$ has measure 1, we set $y_\zeta = x_\zeta^{n_\zeta}$, where $n_\zeta \rightarrow \infty$, and argue as before. \square

Suppose that $A_\zeta \subseteq X_\zeta$ (almost always). As discussed earlier, this gives rise to a subset of $\mathcal{U}\vec{X}$ which can be identified with $\mathcal{U}\vec{A}$. We denote its image in \hat{X} by \hat{A} . In fact, restricting the metrics, $(\hat{A}, \hat{\sigma})$ is the limit of the subspaces (A_ζ, σ_ζ) constructed intrinsically. Note that $x \in \hat{A}$ if and only if $\sigma_\zeta(x_\zeta, A_\zeta) \rightarrow 0$ (where we are taking limits with respect to the ultrafilter on \mathcal{Z}). We also note that \hat{A} is closed in the induced topology on \hat{X} . This can be seen by a similar argument to Lemma 5.2, or simply by noting that \hat{A} is complete in the induced metric. Note that $\hat{\mathbb{R}}$ is an ordered abelian group, which we refer to as the *extended reals*. (In this paper, this will usually be denoted instead by \mathbb{R}^* , for the reasons given below.)

Suppose that $f_\zeta : X_\zeta \rightarrow Y_\zeta$ is a \mathcal{Z} -sequence of maps between the metric spaces (X_ζ, σ_ζ) and (Y_ζ, σ'_ζ) . We have a map, $\mathcal{U}f : \mathcal{U}\vec{X} \rightarrow \mathcal{U}\vec{Y}$ given by Section 1. Suppose there is a constant, $k \in [0, \infty)$, and a \mathcal{Z} -sequence, $(h_\zeta)_\zeta$, in $[0, \infty)$ with $h_\zeta \rightarrow 0$, such that for almost all ζ and all $x, y \in X_\zeta$ we have $\sigma'_\zeta(f_\zeta(x), f_\zeta(y)) \leq k\sigma_\zeta(x, y) + h_\zeta$. Then, $\mathcal{U}f$ induces a k -Lipschitz map $\hat{f} : \hat{X} \rightarrow \hat{Y}$. (The graph of \hat{f} is the limit of the graphs of the f_ζ , taking the l^1 metrics on $X_\zeta \times Y_\zeta$.) The image $\hat{f}(\hat{Y})$ is the limit of the images, $f_\zeta(X_\zeta)$, in the sense of the previous paragraph.

Suppose that $((X_\zeta, \mathcal{Z}))_\zeta$ is a \mathcal{Z} -sequence of geodesic metric spaces. Then the components of $(\hat{X}, \hat{\sigma})$ are precisely the connected components, and each such component is a geodesic space. (This can be seen by applying the previous paragraph to geodesics, thought of as uniformly Lipschitz maps of a compact real interval into the spaces X_ζ .)

Suppose that $(X_\zeta, \rho_\zeta) = (X, \rho)$ a constant sequence. In this case, we get a natural injective map of (X, ρ) into the limit $(\hat{X}, \hat{\rho})$, which is an isometry onto its range. The closure of this range in \hat{X} is just the metric completion of X .

More interestingly, we can take a positive infinitesimal, $\mathbf{t} \in \mathcal{U}\mathbb{R}$, and set $\sigma_\zeta = \mathbf{t}_\zeta \rho$ to be the rescaled pseudometric. In this case, we write $(X^*, \rho^*) = (\hat{X}, \hat{\sigma})$ for the limiting space. Note that this is the same as taking the rescaled metric space $(\hat{X}, \mathbf{t}\hat{\rho})$ and passing to its Hausdorffification.

Definition. We refer to (X^*, ρ^*) as the *extended asymptotic cone* of X with respect to \mathbf{t} .

Note that X^* has a preferred basepoint, namely that given by any constant sequence in \vec{X} . This, in turn, determines a preferred component, X^∞ , of X^* , namely that containing this basepoint.

Definition. We refer to X^∞ as the *asymptotic cone* of X with respect to \mathbf{t} .

By Lemma 5.2, the asymptotic cone is always complete. If X is a geodesic space, so is X^∞ .

One can generalise the above to a \mathcal{Z} -sequence of metric spaces, (X_ζ, ρ_ζ) , rescaled by an infinitesimal \mathbf{t} , to give an extended asymptotic cone, (X^*, ρ^*) . In this case, one needs a sequence of basepoints, $e_\zeta \in X_\zeta$ to determine a basepoint and base component of X^* .

Definition. We say that a \mathcal{Z} -sequence of maps, $f_\zeta : X_\zeta \rightarrow Y_\zeta$ between metric spaces are *uniformly coarsely Lipschitz* if there are constants, $k, h \geq 0$, such that for almost all $\zeta \in \mathcal{Z}$ and all $x, y \in X_\zeta$, we have $\sigma'_\zeta(f_\zeta x, f_\zeta y) \leq k\sigma_\zeta(x, y) + h$. They are *uniform quasi-isometric embeddings* if also $\sigma_\zeta(x, y) \leq k\sigma'_\zeta(f_\zeta x, f_\zeta y) + h$. They are *uniform quasi-isometries* if also $Y_\zeta = N(f_\zeta(X_\zeta); h)$.

Lemma 5.3. *A \mathcal{Z} -sequence of uniformly coarsely Lipschitz maps, $f_\zeta : X_\zeta \rightarrow Y_\zeta$, induces a Lipschitz map, $f^* : X^* \rightarrow Y^*$, which restricts to a map $f^\infty : X^\infty \rightarrow Y^\infty$. If the maps f_ζ are uniform quasi-isometric embeddings, then f^* and f^∞ are bi-Lipschitz onto their range. If they are quasi-isometries, then f^* and f^∞ are bi-Lipschitz homeomorphisms.*

Proof. By Lemma 5.1, we have a map $\mathcal{U}f : \mathcal{U}X \rightarrow \mathcal{U}Y$. This descends to a map $f^* : X^* \rightarrow Y^*$ (since $t_\zeta \sigma_\zeta(x_\zeta, y_\zeta) \rightarrow 0$ implies $t_\zeta \sigma'_\zeta(f_\zeta(x_\zeta), f_\zeta(y_\zeta)) \rightarrow 0$). The fact that f^* and its restriction f^∞ are (bi-)Lipschitz follows since the $ht_\zeta \rightarrow 0$, so the additive constant disappears in the limit. \square

In particular, quasi-isometric spaces have bi-Lipschitz equivalent asymptotic cones (for the same scaling sequence).

An example of the above construction is given by a sequence, $\vec{G} = (G_\zeta)_\zeta$ of graphs. Let $V_\zeta = V(G_\zeta)$ be the vertex sets. The adjacency relations on the V_ζ determine an adjacency relation on $\mathcal{U}\vec{V}$, so as to give it the structure of the vertex set $V(\mathcal{U}\vec{G})$ of a graph $\mathcal{U}\vec{G}$. If each G_ζ is connected, the combinatorial distance functions on V_ζ give us a limiting nonstandard metric and hence an idealised metric on $\mathcal{U}\vec{V}$, with values in $\mathbb{N} \cup \{\infty\}$. This is the same as the combinatorial idealised metric given by $\mathcal{U}V = V(\mathcal{U}\vec{G})$. In particular, the components are again the connected components. (Note that we lose some information in the standardisation process, since different pairs of components might be at different unlimited distances apart.)

Suppose that $\vec{\Gamma} = (\Gamma_\zeta)_\zeta$ is a \mathcal{Z} -sequence of groups. Then $\mathcal{U}\vec{\Gamma}$ is also a group. If each Γ_ζ acts on a set X_ζ , then $\mathcal{U}\vec{\Gamma}$ acts on $\mathcal{U}\vec{X}$. If Γ_ζ acts by isometry in some metric space, then so does $\mathcal{U}\vec{\Gamma}$. If Γ and X are fixed, then any two points of $X \subseteq \mathcal{U}\vec{X}$ in the same $\mathcal{U}\Gamma$ -orbit also lie in the same Γ -orbit (since if $y = gx$ for some $g \in \mathcal{U}\Gamma$, then $y = g_\zeta x$ for almost all g_ζ , and so certainly for some g_ζ).

If Γ is a fixed group acting on a metric space, X , we get an induced action of $\mathcal{U}\Gamma$ on the extended asymptotic cone, X^* (with respect to any infinitesimal \mathbf{t}). Note that we can identify Γ as a normal subgroup of $\mathcal{U}\Gamma$. In fact, we have normal subgroups, $\Gamma \triangleleft \mathcal{U}^1\Gamma$, $\mathcal{U}^1\Gamma \triangleleft \mathcal{U}^0\Gamma$ and $\mathcal{U}^0\Gamma \triangleleft \Gamma_\infty$ of Γ_∞ , where $\mathcal{U}^1\Gamma$ is the stabiliser of the basepoint of X^* , and $\mathcal{U}^0\Gamma$ is the setwise stabiliser of the asymptotic cone, X^∞ . Note that $\mathcal{U}^1\Gamma$ and $\mathcal{U}^0\Gamma$ may depend on \mathbf{t} .

If the action of Γ on X is cobounded (i.e., X is a bounded neighbourhood of some, hence any, Γ -orbit), then the actions of $\mathcal{U}\Gamma$ on X^* and of $\mathcal{U}^0\Gamma$ on X^∞ are transitive. In particular, X^* and X^∞ are homogeneous (extended) metric spaces.

Note that a special case of this construction is \mathbb{R}^* , which is always isomorphic to the extended reals, $\hat{\mathbb{R}}$. If X is a Gromov hyperbolic space, then X^* is an \mathbb{R}^* -tree, and X^∞ is an \mathbb{R} -tree. Of course, this also applies to the asymptotic cone of a sequence of uniformly hyperbolic spaces.

Terminology. To briefly summarise our terminology, we use “nonstandard” to refer to ultralimits, “extended” to refer to the standard part of a nonstandard number (quotienting out by infinitesimals), and “idealised” to mean we are adjoining $\pm\infty$. In this way, we have the extended reals \mathbb{R}^* as a subset (or quotient) of the nonstandard reals $\mathcal{U}\mathbb{R}$. We can view the idealised reals, $[-\infty, \infty]$, as a quotient of \mathbb{R}^* .

6. Coarse median spaces

Coarse median spaces were defined in [Bowditch 2013]. The main point here is that they give a means of talking about (quasi)cubes or (quasi)flats in a geodesic space. Following the construction of [Behrstock and Minsky 2011], this is applicable to the mapping class group, as shown in [Bowditch 2013]. It also applies to Teichmüller space in either the Teichmüller metric, [Bowditch 2016a] or the Weil–Petersson metric [Bowditch 2015]. We remark that another class of space which encompasses these cases, and which implies coarse median, is described in [Behrstock, Hagen and Sisto 2015; 2017].

Before continuing, we introduce the following general conventions.

Conventions. Given two points, x, y , in a metric space, and $r \geq 0$, we will write $x \sim_r y$ to mean that the distance between them is at most r . We will often simply write $x \sim y$, and behave as though this relation were transitive. Here is understood that, at any given stage, the bound r depends only on the constants introduced at

the beginning of an argument. It can be explicitly determined by following through the steps of the argument, though we will not usually explicitly estimate it.

Similarly, given two functions f, g , we will write $f \sim g$ to mean that $f(x) \sim g(x)$ for all x in the domain.

We are often only interested in maps defined up to bounded distance. For a graph it would therefore be enough to specify a map on the set of vertices. When referring to a finite product of metric spaces, we can always take the l^1 metric. For a finite product of graphs, we can always restrict to the 1-skeleton of the product cube complex. In any case, we will only be interested in the product metric defined up to bi-Lipschitz equivalence.

We will sometimes adopt a similar convention for linear bounds. Given $\lambda \geq 1$ and $r \geq 0$, we write $x \asymp_{\lambda,r} y$ to mean that $\lambda^{-1}(x - r) \leq y \leq \lambda x + r$. Again, we usually omit λ, r from the notation, and write $x \asymp y$.

When we come to discuss marking graphs, the constants implicit in the notation \sim and \asymp will ultimately depend only on the complexity, $\xi(\Sigma)$, of our surface, Σ , as defined in Section 7. We will make this explicit at the relevant points.

Let (\mathcal{M}, ρ) be a geodesic metric space.

Definition. We say that a ternary operation, $\mu : \mathcal{M}^3 \rightarrow \mathcal{M}$, is a “coarse median” if it satisfies the following:

- (C1) There are constants, $k, h(0)$, such that for all $a, b, c, a', b', c' \in \mathcal{M}$ we have $\rho(\mu(a, b, c), \mu(a', b', c')) \leq k(\rho(a, a') + \rho(b, b') + \rho(c, c')) + h(0)$, and
- (C2) There is a function, $h : \mathbb{N} \rightarrow [0, \infty)$, with the following property. Suppose that $A \subseteq \mathcal{M}$ with $1 \leq |A| \leq p < \infty$. Then there is a median algebra, (Π, μ_Π) and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \mathcal{M}$ such that for all $x, y, z \in \Pi$ we have $\rho(\lambda\mu_\Pi(x, y, z), \mu(\lambda x, \lambda y, \lambda z)) \leq h(p)$ and for all $a \in A$, we have $\rho(a, \lambda\pi a) \leq h(p)$.

We say that \mathcal{M} has *rank at most n* if we can always take Π to have rank at most n (as a median algebra). We say that \mathcal{M} is *n -colourable* if we can always take Π to be n -colourable. We refer to (\mathcal{M}, ρ, μ) as a *coarse median space*. We refer to k, h as the *parameters* of \mathcal{M} .

From (C2) we can deduce that, if $a, b, c \in \mathcal{M}$, then $\mu(a, b, c)$, $\mu(b, a, c)$ and $\mu(b, c, a)$ are a bounded distance apart, and that $\rho(\mu(a, a, b), a)$ is bounded. Since we are only really interested in μ up to bounded distance, we can assume that μ is invariant under permutation of a, b, c and that $\mu(a, a, b) = a$.

Note that in (C2), we can always assume that $\Pi = \langle \pi A \rangle$ (in particular, that it is finite). Also, if we are not concerned about rank, we can always take Π to be the free median algebra on A , and π to be the inclusion of A in Π .

Note that a direct of product of coarse median spaces is also a coarse median space.

For future reference, we note:

Lemma 6.1. *Suppose $a, b, c \in \mathcal{M}$, 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 and k_2 depend only on the parameters of \mathcal{M} .*

Proof. Using property (C1), we see that the maps $x \mapsto \mu(a, c, x)$ and $x \mapsto \mu(b, c, x)$ are coarsely Lipschitz, and so we get linear bounds on $\rho(a, c)$ and $\rho(b, c)$ in terms of $\rho(a, b)$. \square

With the conventions introduced earlier, this shows $\rho(a, b) \asymp \rho(a, c) + \rho(c, b)$. (where the implicit constants depend only on the parameters of \mathcal{M}).

Given two spaces X, Y , equipped with ternary operations μ_X and μ_Y , together with a metric ρ on Y , we say that a map $\phi : X \rightarrow Y$ is an l -quasimorphism if $\rho(\phi\mu_X(x, y, z), \mu_Y(\phi x, \phi y, \phi z)) \leq l$ for all $x, y, z \in X$. Typically, Y will be a coarse median space, and X will be either a median algebra or a coarse median space. (Note that the map λ featuring in (C2) is an $h(p)$ -quasimorphism.)

Lemma 6.2. *Suppose that Π is a median algebra generated by a finite subset, $B \subseteq \Pi$. Suppose that $\lambda, \lambda' : \Pi \rightarrow \mathcal{M}$ are l -quasimorphisms with $\rho(\lambda b, \lambda' b) \leq l$ for all $b \in B$. Then, for all $x \in \Pi$, $\rho(\lambda x, \lambda' x)$ is bounded above by some linear function of l , depending only on the parameters of \mathcal{M} and the cardinality of B .*

Proof. Define $B_i \subseteq \Pi$ inductively by $B_0 = B$ and $B_{i+1} = \mu(B_i^3)$. We see inductively that $\lambda|_{B_i}$ and $\lambda'|_{B_i}$ are a bounded distance apart, where the bound depends on i and is linear in l . Now $|\Pi| \leq q = 2^{2^p}$ where $p = |B|$, and so certainly, $\Pi = B_q$, and the result follows. \square

In particular, in clause (C2) of the definition, if we assume that $\Pi = \langle \pi A \rangle$, then the map λ is unique up to bounded distance depending only on the parameters and p .

The following will allow us to assume that quasimorphisms of cubes are in fact uniform quasimorphisms:

Lemma 6.3. *Given $n \in \mathbb{N}$, there are constants k_0, h_0 and h_1 depending only on n and the parameters of \mathcal{M} such that the following holds. Suppose that $Q = \{-1, 1\}^n$ and that $\psi : Q \rightarrow \mathcal{M}$ is an l -quasimorphism for some $l \geq 0$. Then there is an h_0 -quasimorphism, $\phi : Q \rightarrow \mathcal{M}$, with $\rho(\phi x, \psi x) \leq k_0 l + h_1$ for all $x \in Q$.*

Proof. Let Π be the free median algebra on the set Q , and let $\theta : \Pi \rightarrow Q$ be the unique median homomorphism extending the identity on Q (thought of as a map from a set to a median algebra). Now there is a median monomorphism, $\omega : Q \rightarrow \Pi$ with $\theta \circ \omega$ the identity on Q . (To see this, we can think of Π as the vertex set of a finite CAT(0) cube complex. Every pair of intrinsic faces of $Q \subseteq \Pi$ are separated

by some hyperplane of Π , and these must all intersect in some n -cell of Π . Each element, $x \in Q$, determines a unique vertex $\omega(x)$ of this n -cell. This gives us a homomorphism $\omega : Q \rightarrow \Pi$, with $\omega(Q)$ equal to the vertex set of the n -cell. Note that ω is not canonically determined: it might depend on the choice of cell.)

Now apply (C2) to $\psi(Q) \subseteq \mathcal{M}$ to give an $h(2^n)$ -quasimorphism $\lambda : \Pi \rightarrow \mathcal{M}$ with $\lambda|_Q \sim_{h(2^n)} \psi$. Let $\phi = \lambda \circ \omega : Q \rightarrow \mathcal{M}$. This is an h_0 -quasimorphism, where $h_0 = h(2^n)$.

Let $\lambda' = \lambda \circ \theta : \Pi \rightarrow \mathcal{M}$. Thus λ' is a l -quasimorphism, and $\lambda'|_Q = \psi = \lambda|_Q$. By Lemma 6.2, we have $\rho(\lambda x, \lambda' x) \leq k_0 l + h_2$ for all $x \in \Pi$, where k_0, h_2 depend only on the parameters of \mathcal{M} . But $\lambda' \circ \omega|_Q = \lambda \circ \theta \circ \omega|_Q = \lambda|_Q \sim_{h(2^n)} \psi$, and so we see that $\rho(\phi x, \psi x) \leq k_0 l + h_1$ for all $x \in Q$, where $h_1 = h_2 + h(2^n)$. \square

The following two lemmas will be used to establish that statements that hold in a median algebra hold up to bounded distance in a coarse median space.

Lemma 6.4. *Suppose that (\mathcal{M}, ρ, μ) is a coarse median space. Suppose that Π is a finite median algebra with $|\Pi| \leq q < \infty$, and that $\lambda : \Pi \rightarrow \mathcal{M}$ is an l -quasimorphism. Given $t \geq 0$, there is a finite median algebra Π' , a map $\lambda' : \Pi' \rightarrow \mathcal{M}$ and an epimorphism $\theta : \Pi \rightarrow \Pi'$ such that for all distinct $x, y \in \Pi'$, $\rho(\lambda' x, \lambda' y) > t$, and for all $z \in \Pi$, $\rho(\lambda z, \lambda' \theta z) \leq s$, where s depends only on q, h, t and the parameters of \mathcal{M} .*

Proof. Define a relation, \approx , on Π , by setting $x \approx y$ if $\rho(\lambda x, \lambda y) \leq t$. Let \simeq be the smallest equivalence relation on Π containing \approx with the property that whenever $x, y, z, w \in \Pi$ with $z \simeq w$, we have $\mu_\Pi(x, y, z) \simeq \mu_\Pi(x, y, w)$. Let $\Pi' = \Pi / \simeq$ be the quotient median algebra as defined at the end of Section 2, and let $\theta : \Pi \rightarrow \Pi'$ be the quotient map. Define $\lambda' : \Pi' \rightarrow \mathcal{M}$ by setting $\lambda'(x)$ to be the λ -image of any representative of the \simeq -class of x in Π . Since \simeq includes \approx , we see that $\rho(\lambda' x, \lambda' y) > t$ for all distinct $x, y \in \Pi'$.

We claim that if $x, y \in \Pi$, with $x \simeq y$, then $\rho(\theta \lambda x, \theta \lambda y)$ is bounded above in terms of q, h, t and the parameters of \mathcal{M} . To see this, note that \simeq can be constructed from \approx by iterating two operations. We start with \approx . Whenever $z \approx w$, then we set $\mu_\Pi(x, y, z)$ related to $\mu_\Pi(x, y, w)$ for all $x, y \in \Pi$. Also, if $a \approx b$ and $b \approx c$, then we set a to be related to c . We continue again with the relation thus defined. After at most q steps, this process stabilises on the relation \simeq . From the fact that λ is a quasimorphism, and from property (C1) for \mathcal{M} , we see that at each stage the maximal distance between the λ -images of related elements of Π can increase by at most a linear function which depends only on l and the parameters of \mathcal{M} . This now proves the claim.

Suppose that $z \in \Pi$. By construction, $\lambda' \theta z = \lambda w$ for some $w \simeq z$. By the above, $\rho(\lambda z, \lambda' \theta z) = \rho(\lambda z, \lambda w)$ is bounded as required. \square

Note that λ' is itself an l' -quasimorphism, where l' depends only on q, h, t and the parameters of \mathcal{M} . This enables us to give a refinement of (C2) as follows:

Corollary 6.5. *Suppose that (\mathcal{M}, ρ, μ) is a coarse median space, and $t \geq 0$. Then there is a function, $h_t : \mathbb{N} \rightarrow [0, \infty)$ with the following property. Suppose that $A \subseteq \mathcal{M}$ with $1 \leq |A| \leq p < \infty$. Then there is a finite median algebra, (Π, μ_Π) , and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \mathcal{M}$ such that λ is a $h_t(p)$ -quasimorphism with $\rho(\lambda x, \lambda y) > t$ for all distinct $x, y \in \Pi$, and such that $\rho(a, \lambda \pi a) \leq h_t(p)$ for all $a \in A$.*

Proof. Start with Π, π, λ as given by (C2) for \mathcal{M} (so that $\lambda : \Pi \rightarrow \mathcal{M}$ is an $h(p)$ -quasimorphism, with $h(p)$ independent of t). We can assume that $|\Pi| \leq 2^{2^p}$. We now apply Lemma 6.4 to give Π', λ' and $\theta : \Pi \rightarrow \Pi'$. Now replace Π by Π', π by $\theta \circ \pi$, and λ by λ' . \square

By an “identity” in a median algebra, we mean an expression equating two terms featuring only the median operation. We refer to it as a “tautological identity” if it holds, whatever the arguments in any median algebra, M . (For example, we have the tautological identity: $\mu(a, b, \mu(a, b, c)) = \mu(a, b, c)$, for all $a, b, c \in M$.) We remark that an identity can easily be verified algorithmically: it is sufficient to check it for all possible assignments of the arguments in the two-point median algebra $\{-1, 1\}$. We make the following general observation.

General Principle. Any tautological median identity holds up to bounded distance in any coarse median space, \mathcal{M} .

More formally, this says that if P and Q are formulae defining $P(a_1, \dots, a_n)$ and $Q(a_1, \dots, a_n)$ in terms of μ , and the identity $P(a_1, \dots, a_n) = Q(a_1, \dots, a_n)$ holds for any $a_1, \dots, a_n \in M$ in any median algebra (M, μ) then it follows that $\rho(P(a_1, \dots, a_n), Q(a_1, \dots, a_n))$ is bounded for any $a_1, \dots, a_n \in \mathcal{M}$ in any coarse median space (\mathcal{M}, μ, ρ) . The bound only depends on (the complexity of) the formulae P, Q and the parameters of \mathcal{M} .

For example, for all $a, b, c \in \mathcal{M}$, $\rho(\mu(a, b, \mu(a, b, c)), \mu(a, b, c))$ is bounded above by a constant depending only on the parameters of \mathcal{M} for all $a, b, c \in \mathcal{M}$.

To prove this principle, let $A \subseteq \mathcal{M}$ be the set of elements occurring as arguments, and let $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \mathcal{M}$ be as given by (C2) of the hypotheses. Now apply either side of the identity to the π -images in Π to give an element $x \in \Pi$ (by assumption, this will be the same element for either side). We can also apply each side of the same identity to the elements of A , using the median structure, μ , on \mathcal{M} . In this way, we get two elements of \mathcal{M} . Using (C1) and (C2) directly, we see that these are both a bounded distance from λx , and so, a bounded distance from each other. The claim follows.

A more general statement holds for conditional identities. Suppose that some finite set of identities (the “input identities”) imply another identity (the “derived

identity”) in any median algebra. (For example, $d \in [a, b] \cap [b, c] \cap [c, a]$ implies $d = \mu(a, b, c)$.) We have the following generalisation.

General Principle. Given a finite set of input identities, and a derived identity, if we suppose that the input identities hold up to bounded distance for a particular set of elements in a coarse median space, \mathcal{M} , then the derived identity also holds up to bounded distance for this set of elements.

(So, for example, if $a, b, c, d \in \mathcal{M}$, with the three distance $\rho(\mu(a, b, d), d)$, $\rho(\mu(b, c, d), d)$ and $\rho(\mu(c, a, d), d)$ all bounded, then $\rho(\mu(a, b, c), d)$ is also bounded.)

The argument is essentially the same. This time, we apply Corollary 6.5 to the set A instead of (C2) directly. Suppose that $x, y \in \Pi$ are respectively the π -images of the left and right sides of one of the input identities. As in the previous argument, we see that λx and λy are respectively a bounded distance from the result of applying the same formulae in \mathcal{M} , which by assumption, are a bounded distance apart in \mathcal{M} . It follows that $\rho(\lambda x, \lambda y)$ is bounded. By choosing the constant t in Corollary 6.5 to be larger than this bound, we see that we must have $x = y$. In other words, this input identity holds exactly in Π , for the π -images of the elements of A . We can assume this is true of all the input identities. Therefore, the derived identity must hold too. Now, again, as in the previous argument, we see that the derived identity holds up to bounded distance in \mathcal{M} . This proves the claim.

We can apply the these principles in the following discussion.

Given $a, b \in \mathcal{M}$, we define the *coarse interval* between a and b as $[a, b] = \{\mu(a, b, x) \mid x \in \mathcal{M}\}$. By the observation above, we see that is a bounded Hausdorff distance from $\{x \in \mathcal{M} \mid \rho(\mu(a, b, x), x) \leq r\}$ for any fixed sufficiently large $r \geq 0$.

Definition. We say that a subset, $C \subseteq \mathcal{M}$, of a coarse median space, (\mathcal{M}, ρ, μ) , is *k-(median) quasiconvex* if for all $a, b \in C$ and $x \in \mathcal{M}$, $\rho(\mu(a, b, x), C) \leq k$.

From property (C1) we see that any quasiconvex set is quasi-isometrically embedded in \mathcal{M} (or more precisely, some uniform neighbourhood of C is quasi-isometrically embedded with respect to the induced path-metric). Note also quasiconvexity of C is equivalent to asserting that for all $a, b \in C$, the coarse interval $[a, b]$ lies in a uniform neighbourhood of C . Note that Lemma 6.1 implies that if $a, b, c \in \mathcal{M}$ and $c \in [a, b]$, then $\rho(a, b)$ agrees with $\rho(a, c) + \rho(c, b)$ up to linear bounds.

We next recall the following standard notion for any median algebra, M . Suppose that $C \subseteq M$ is (a priori) any subset. We say that a map $f : M \rightarrow C$ is a *gate map* if $fx \in [x, c]$ for all $x \in M$ and $c \in C$. Note that if $a, b \in M$ and $c \in [a, b]$ then $fc \in [a, c] \cap [b, c] \in \{c\}$, so $fc = c$. It follows immediately that $f|_C$ is the identity (since $c \in [c, c]$), and that C is convex (since $c = fc \in C$). We also claim that f is a homomorphism. For this, it is enough to show that if $c \in [a, b]$, then $fc \in [fa, fb]$. But now the identities $c \in [a, b]$, $fc \in [c, fb]$ and $fb \in [b, fc]$ together imply

$fc \in [a, fb]$. Thus (by the same observation, with a, b, c replaced by fb, a, fc) we get $fc = ffc \in [fa, fb]$ as required. We also note that if a gate map exists for a given C , then it is unique.

We can now define the corresponding notion in a coarse median space, \mathcal{M} .

Definition. A map $\phi : \mathcal{M} \rightarrow C$ to a subset $C \subseteq \mathcal{M}$ is an r -coarse gate map if for all $x \in \mathcal{M}$ and $c \in C$, we have $\rho(x, \mu(x, \phi x, c)) \leq r$.

Lemma 6.6. *If $\phi : \mathcal{M} \rightarrow C$ is an r -coarse gate map, then C is k -quasiconvex, ϕ is an l -quasimorphism, and $\rho(c, \phi c) \leq h$ for all $c \in C$, where k, l, h depend only on r and the parameters of \mathcal{M} .*

Proof. We follow the same argument as for a median algebra described above, except that now equalities and inclusions are assumed to hold up to bounded distance, depending only on r , and the parameters of \mathcal{M} . By the general principles described above, any deduction (based on a finite sequence of identities) in a median algebra holds also in a coarse median algebra, interpreting everything up to bounded distance. \square

Suppose now that $((\mathcal{M}_\zeta, \rho_\zeta, \mu_\zeta))_\zeta$ is a \mathcal{Z} -sequence of uniformly coarse median spaces (i.e., with parameters independent of ζ). Let $t \in \mathcal{U}\mathbb{R}$ be a positive infinitesimal. We get a limiting space, $(\mathcal{M}^*, \rho^*, \mu^*)$, where (\mathcal{M}^*, ρ^*) is the extended asymptotic cone, and where (\mathcal{M}^*, μ^*) is a topological median algebra (that is, the map $\mu^* : (\mathcal{M}^*)^3 \rightarrow \mathcal{M}^*$ is continuous). If each \mathcal{M}_ζ has rank at most n (as a coarse median space), then \mathcal{M}^* has rank at most n (as a median algebra). Note that (\mathcal{M}^*, μ^*) need not be a median metric space, though it satisfies a weaker metric condition described in [Bowditch 2013; 2014a], namely that the maps $x \mapsto \mu(a, b, x)$ are uniformly Lipschitz, for all $a, b \in \mathcal{M}^*$, (see Proposition 9.1 and Lemma 9.2 of [Bowditch 2013]). Note that this is the hypothesis of Proposition 2.4 here.

In those papers, we restricted attention to the asymptotic cone, \mathcal{M}^∞ , but that does not affect the above observations.

Lemma 6.7. *Let \mathcal{M}^* be an extended asymptotic cone of a \mathcal{Z} -sequence of uniformly coarse median spaces. Suppose that $Q \subseteq \mathcal{M}^*$ is an n -cube. Then we can find a sequence of l_0 -quasimorphisms $\phi_\zeta : Q \rightarrow \mathcal{M}_\zeta$ such that for all $x \in Q$, $\phi_\zeta x \rightarrow x$, where l_0 depends only on n and the uniform parameters of the \mathcal{M}_ζ .*

Proof. To begin, take any sequence of maps, $\psi_\zeta : Q \rightarrow \mathcal{M}_\zeta$, with $\psi_\zeta x \rightarrow x$ for all $x \in Q$. (Such maps exist directly from the definition of the asymptotic cone.) Since μ^* is, by definition, the limit of the μ_ζ , it follows that ψ_ζ is a h_ζ -quasimorphism, where $t_\zeta h_\zeta \rightarrow 0$ (since they must converge to a monomorphism in \mathcal{M}^*). Let $\phi_\zeta : Q \rightarrow \mathcal{M}_\zeta$ be the l_0 -quasimorphism given by Lemma 6.3. For all $x \in Q$, $\rho_\zeta(\phi_\zeta x, \psi_\zeta x) \leq kh_\zeta + h_1$ so $t_\zeta \rho_\zeta(\phi_\zeta x, \psi_\zeta x) \leq kt_\zeta h_\zeta + h_1 t_\zeta \rightarrow 0$. Thus $\phi_\zeta x \rightarrow x$, as required. \square

Note that, if we have a sequence of uniformly quasiconvex sets, $C_\zeta \subseteq \mathcal{M}_\zeta$, we have a limiting bi-Lipschitz embedded closed convex subset $C^* \subseteq \mathcal{M}^*$ in the extended asymptotic cone \mathcal{M}^* . If $\phi_\zeta : \mathcal{M}_\zeta \rightarrow C_\zeta$ are a sequence of uniform coarse gate maps, the limiting map $\phi^* : \mathcal{M}^* \rightarrow C^*$ is a gate map.

As in [Bowditch 2013, Section 12], we say that a median algebra, Π , is *n-colourable* if there is an *n-colouring* of the walls such that no two walls of the same colour cross. We say that a coarse median space \mathcal{M} is *n-colourable* if in (C2) we can always choose Π to be *n-colourable* as a median algebra. Clearly this implies that \mathcal{M} has rank at most *n*. The following was shown in [Bowditch 2014a, Proposition 12.5].

Theorem 6.8. *Suppose that $((\mathcal{M}_\zeta, \rho_\zeta, \mu_\zeta))_\zeta$ is a sequence of *n-colourable* uniform coarse median spaces, for some fixed *n*. Then \mathcal{M}^* admits a metric ρ' , bi-Lipschitz equivalent to ρ^* , such that (\mathcal{M}^*, ρ') is an (extended) median metric space with median μ^* . Moreover, \mathcal{M}^* is *n-colourable* as a median algebra.*

In fact, the bi-Lipschitz constant only depends on the parameters of the coarse median spaces.

The construction however is not canonical. Note that the median metric space arising is necessarily proper.

In particular, we see that the asymptotic cone of a sequence of finitely colourable coarse median spaces is bi-Lipschitz equivalent to a proper median metric space, and hence in turn to a CAT(0) space (by Theorem 2.2). In fact, the same holds for a sequence of finite-rank coarse median spaces. This relies on the following variation of Theorem 6.8:

Theorem 6.9. *Suppose that $((\mathcal{M}_\zeta, \rho_\zeta, \mu_\zeta))_\zeta$ is a sequence of coarse median spaces of rank *n*, for some fixed *n*. Then \mathcal{M}^* admits an extended metric, ρ' , bi-Lipschitz equivalent to ρ^* , such that (\mathcal{M}^*, ρ') is an (extended) median metric space of rank *n*, with median μ^* .*

Proof. As already observed, it is easily seen from axiom (C1) that ρ^* satisfies the hypotheses of Proposition 2.4, when restricted to any component of \mathcal{M}^* . We can therefore apply Proposition 2.4 to each component separately. \square

We finish this section by briefly discussing the special case of a Gromov hyperbolic space (\mathcal{M}, ρ) . See [Bowditch 2013, Section 3] for elaboration.

Given $a, b, c \in \mathcal{M}$, write

$$\langle a, b:c \rangle = \frac{1}{2}(\rho(a, c) + \rho(b, c) - \rho(a, b))$$

for the ‘‘Gromov product’’. Up to bounded distance, this is the same as the distance of c to some (or any) geodesic from a to b .

Definition. We say that $m \in \mathcal{M}$ is an *r-centroid* for $a, b, c \in \mathcal{M}$ if $\langle a, b:m \rangle \leq r$, $\langle b, c:m \rangle \leq r$ and $\langle c, a:m \rangle \leq r$.

Provided r is a sufficiently large in relation to the hyperbolicity constant, such an r -centroid will always exist. We will fix such an r and simply refer to m as a *centroid*. In fact, m is well defined up to bounded distance, and we write $\mu(a, b, c) = m$ for some choice of m . With this structure, (\mathcal{M}, ρ, μ) is a coarse median space of rank at most 1. Indeed, any rank-1 coarse median space arises in this way. (We note that rank-0 is trivially equivalent to having finite diameter.)

If $(\mathcal{M}_\xi)_\xi$ is a sequence of uniformly hyperbolic spaces, then (\mathcal{M}^*, μ^*) is a rank-1 median algebra (variously known in the literature as a “tree algebra”, “median pretree”, etc.). As already observed in Section 5, (\mathcal{M}^*, ρ^*) is an \mathbb{R}^* -tree, and $(\mathcal{M}^\infty, \rho^\infty)$ is an \mathbb{R} -tree.

It is shown in [Bowditch 2013] that if \mathcal{M} is a coarse median space of rank at most n , then there is no quasi-isometric embedding of \mathbb{R}^{n+1} into \mathcal{M} (since this would give rise to an injective map of \mathbb{R}^{n+1} into \mathcal{M}^∞ , contradicting the fact that \mathcal{M}^∞ has locally compact dimension at most n). In fact, the same argument can be applied to bound the radii of quasi-isometrically embedded balls. To state this more precisely, write B_r^n for the ball of radius r in the euclidean space \mathbb{R}^n .

Lemma 6.10. *Let \mathcal{M} be a coarse median space of rank at most n . Given parameters of \mathcal{M} and of quasi-isometry, there is some constant $r \geq 0$, such that there is no quasi-isometric embedding of B_r^{n+1} into \mathcal{M} with these parameters.*

Proof. Suppose that, for each $i \in \mathbb{N}$, the ball, B_i , of radius i admits a uniformly quasi-isometric embedding, $\phi_i : B_i \rightarrow \mathcal{M}$. Now pass to the asymptotic cone with indexing set, $\mathcal{Z} = \mathbb{N}$, and with scaling factors $1/i$. We end up with a continuous injective map, $\phi^\infty : B_1 \rightarrow \mathcal{M}^\infty$, contradicting the fact that \mathcal{M}^∞ has locally compact dimension at most n .

To see that only the parameters of \mathcal{M} are relevant to the value of r , we should allow \mathcal{M} also to vary among coarse median spaces with these parameters when taking the asymptotic cone. More precisely, suppose we have a sequence, $\phi_i : B_i \rightarrow \mathcal{M}_i$, of uniformly quasi-isometric maps, where the \mathcal{M}_i are uniform coarse median spaces. This time, we get a limiting map $\phi^\infty : B_1 \rightarrow \mathcal{M}^\infty$, where \mathcal{M}^∞ is the ultralimit of the spaces $(\mathcal{M}_i)_i$ again scaled by $1/i$. This leads to the same contradiction. In other words, there must be a bound on the diameter of a euclidean ball which we can quasi-isometrically embed, for any fixed parameters. \square

Note that it follows, for example, that \mathcal{M} admits no quasi-isometrically embedded euclidean half-space of dimension $n + 1$.

Remark. The last paragraph of the proof of Lemma 6.10 is a standard trick to obtain uniform constants and will be used again later. (See the remark after Lemma 14.5.)

7. A general construction of coarse medians

In this section, we give a general criterion for the existence of coarse medians on certain types of spaces associated to a surface. In particular, we will apply this to the marking graph in Section 8, to recover the result of [Behrstock and Minsky 2011]. The argument follows broadly as in that work using [Behrstock, Kleiner, Minsky and Mosher 2012]. In doing this, under our hypotheses, we give a version of the compatibility theorem for medians. To this end, we will list a set of axioms ((A1)–(A10) below) which relate to “projection maps” to spaces indexed by a set, \mathcal{X} , namely the collection of “subsurfaces” of a surface Σ . The main results of this section, namely Theorems 7.1 and 7.2, together give a more precise statement of Theorem 1.1.

In Section 8, we will explain how this applies, in particular, to the marking graph, and recover the result of [Behrstock and Minsky 2011].

As mentioned in Section 1, the main purpose of keeping the discussion general is that one can readily check that the hypotheses we give here apply in other situations — notably to Teichmüller space in either the Teichmüller or Weil–Petersson metrics; see [Bowditch 2016a; 2015]. This also applies to most of the discussion in Sections 9–12 here.

We remark that in [Behrstock, Hagen and Sisto 2017], the authors define the notion of a “hierarchically hyperbolic” space, based on a different (though related) set of axioms. These allow for more general indexing sets. However in the case where the indexing set is taken to be \mathcal{X} , as in the present paper, one can verify that hierarchically hyperbolic spaces satisfy our axioms, and are hence coarse median. See [Behrstock, Hagen and Sisto 2015, Section 7] for more discussion of this.

Let Σ be a compact orientable surface. Let $\xi(\Sigma)$ be its complexity, that is, $\xi(\Sigma) = 3g + p - 3$, where g is the genus, and p is the number of boundary components. If $\xi(\Sigma) = 0$ then Σ is a three-holed sphere. If $\xi(\Sigma) = 1$ then Σ is a four-holed sphere, or a one-holed torus. We will write $S_{g,p}$ to denote the topological type of surface of genus g and p boundary components.

Definition. By an *essential curve* in Σ , we mean a simple closed curve which homotopically nontrivial and nonperipheral (not homotopic into $\partial\Sigma$). By a *curve* we mean a free homotopy class of essential curves.

Definition. By an *essential subsurface* X we mean a compact connected subsurface, $X \subseteq \Sigma$, such that each boundary component of X is either a component of $\partial\Sigma$, or else an essential (and nonperipheral) simple closed curve in $\Sigma \setminus \partial\Sigma$, and such that X is not homeomorphic to a three-holed sphere.

Note that we are allowing Σ itself as a subsurface, as well as nonperipheral annuli.

Definition. A *subsurface* is a free homotopy class of essential subsurfaces.

We refer to an essential surface in the given homotopy class as a *realisation* of the subsurface. Note that there is a natural bijective correspondence between curves and annular subsurfaces.

Given $X, Y \in \mathcal{X}$, we distinguish five mutually exclusive possibilities denoted as follows:

- (1) $X = Y$: X and Y are homotopic.
- (2) $X < Y$: $X \neq Y$, and X can be homotoped into Y but not into ∂Y .
- (3) $Y < X$: $Y \neq X$, and Y can be homotoped into X but not into ∂X .
- (4) $X \wedge Y$: $X \neq Y$, and X, Y can be homotoped to be disjoint.
- (5) $X \pitchfork Y$: none of the above.

In (2)–(4) one can find realisations of X, Y in Σ such that $X \subseteq Y$, $Y \subseteq X$, $X \cap Y = \emptyset$, respectively. (Note that $X \wedge Y$ covers the case where X is an annulus homotopic to a boundary component of Y , or vice versa.) We can think of (5) as saying that the surfaces “overlap”. We write $X \preceq Y$ to mean $X < Y$ or $X = Y$. (Note that this excludes the case where Y is homotopic to an annular boundary component of a nonannular subsurface, X .)

We note that

$$\begin{aligned} X \wedge Y &\Leftrightarrow Y \wedge X, & X \pitchfork Y &\Leftrightarrow Y \pitchfork X, \\ X < Y < Z &\Rightarrow X < Z, & X \wedge Y \text{ and } Z < Y &\Rightarrow X \wedge Z. \end{aligned}$$

Given $X \in \mathcal{X}$, write $\mathcal{X}(X) = \{Y \in \mathcal{X} \mid Y \preceq X\}$.

We now introduce the hypotheses of the main result of this section.

We suppose that to each $X \in \mathcal{X}$, we have associated geodesic metric spaces $(\mathcal{M}(X), \rho_X)$ and $(\mathcal{G}(X), \sigma_X)$, as well as a map $\chi_X : \mathcal{M}(X) \rightarrow \mathcal{G}(X)$. We will generally abbreviate $\rho = \rho_X$ and $\sigma = \sigma_X$, where there is no confusion. Given $X, Y \in \mathcal{X}$ with $Y < X$, we suppose that we have a map $\psi_{YX} : \mathcal{M}(X) \rightarrow \mathcal{M}(Y)$. We write $\theta_{YX} = \chi_Y \circ \psi_{YX} : \mathcal{M}(X) \rightarrow \mathcal{G}(Y)$. We also assume that if $X, Y \in \mathcal{X}$ with $Y \pitchfork X$, or $Y < X$, then we have associated an element $\theta_X Y \in \mathcal{G}(X)$. If α is a curve, we will write $\theta_X \alpha = \theta_X Y$, where $Y = X(\alpha)$ is the annular neighbourhood of α . It will be seen that the hypotheses laid out below only really require these maps to be defined up to bounded distance.

(In Section 8, we will be setting $\mathcal{M}(X) = \mathbb{M}(X)$ and $\mathcal{G}(X) = \mathbb{G}(X)$, to be the intrinsic marking graphs and curve graph respectively, when $X \in \mathcal{X}_N$. The map χ_X is the natural projection, and ψ_{YX} is the usual subsurface projection. If $X \in \mathcal{X}_A$, then $\mathcal{G}(X) = \mathbb{G}(X)$ is the usual graph that measures twisting around the core curve. In this case, we set $\mathbb{M}(X) = \mathbb{G}(X)$ and χ_X to be the identity map.)

We will assume:

- (A1) “hyperbolic”: There exists $k \geq 0$ such that for all $X \in \mathcal{X}$, $\mathcal{G}(X)$ is k -hyperbolic.
- (A2) “ χ Lipschitz and cobounded”: there exist $k_1, k_2, k_3 \geq 0$ such that for all $X \in \mathcal{X}$ and for all $a, b \in \mathcal{M}(X)$, $\sigma(\chi_X a, \chi_X b) \leq \rho(a, b) + k_2$ and $\mathcal{G}(X) = N(\chi_X(\mathcal{M}(X)); k_3)$
- (A3) “ ψ Lipschitz”: there exist $k_1, k_2 \geq 0$ such that for all $X \in \mathcal{X}$ and for all $Y \in \mathcal{X}(X)$ and for all $a, b \in \mathcal{M}(X)$, $\rho(\psi_{YX} a, \psi_{YX} b) \leq \rho(a, b) + k_2$.
- (A4) “composition”: There is some $s_0 \geq 0$ such that if $X, Y, Z \in \mathcal{X}$ with $Z \prec Y \prec X$ and $a \in \mathcal{M}(X)$, then $\rho(\psi_{ZX} a, \psi_{ZY} \circ \psi_{YX} a) \leq s_0$.
- (A5) “disjoint projection”: there exists $s_1 \geq 0$ such that for all $X \in \mathcal{X}$ if $Y, Z \in \mathcal{X}$ with $Y \wedge Z$ or $Y \prec Z$, then $\sigma(\theta_X Y, \theta_X Z) \leq s_1$ whenever $\theta_X Y$ and $\theta_X Z$ are defined.

Thus, (A2) and (A3) tell us that our maps are all uniformly coarsely Lipschitz. In view of (A4) we will abbreviate θ_{YX} to θ_Y and ψ_{YX} to ψ_Y , whenever the domain of the map is clear. If $X = Y$, we set $\psi_X = \psi_{XX}$ to be the identity map on $\mathcal{M}(X)$. Note that, with these conventions, we can also write χ_X as θ_X . We will also abbreviate $\sigma_Y(a, b) = \sigma(\theta_Y a, \theta_Y b)$ and $\rho_Y(a, b) = \rho(\psi_Y a, \psi_Y b)$. To simplify the exposition, we will view θ_X as a map from $\mathcal{M}(X) \sqcup \mathcal{X}(X)$ to $\mathcal{G}(Y)$.

Given $a, b \in \mathcal{M}(X)$, write $R_X(a, b) = \max\{\sigma_Y(a, b) \mid Y \in \mathcal{X}(X)\}$. Similarly, if $a \in \mathcal{M}(X)$ and $Z \in \mathcal{X}(X) \setminus \{X\}$, write

$$R_X(a, Z) = \max\{\sigma_Y(a, Z)\}$$

as Y ranges over those elements of $\mathcal{X}(X)$ with either with $Y \prec Z$ or $Y \pitchfork Z$. (In the context of marking graphs, one can view R_X as measuring intersection numbers.)

We assume:

- (A6) “finiteness”: there exists $r_0 \geq 0$ such that for all $X \in \mathcal{X}$ and for all $a, b \in \mathcal{M}(X)$, the set of $Y \in \mathcal{X}(X)$ with $\sigma_Y(a, b) \geq r_0$ is finite.
- (A7) “distance bound”: for all $r \geq 0$ there exists $r' \geq 0$ such that for all $X \in \mathcal{X}$ and for all $a, b \in \mathcal{M}(X)$, if $R_X(a, b) \leq r$, then $\rho(a, b) \leq r'$.
- (A8) “bounded image”: there exists r_0 such that for all $X \in \mathcal{X}$ and for all $Y \in \mathcal{X}(X)$ and for all $a, b \in \mathcal{M}(X)$, if $\langle \theta_X a, \theta_X b; \theta_X Y \rangle \geq r_0$, then $\sigma_Y(a, b) \leq r_0$.
- (A9) “overlapping subsurfaces”: there exists r_0 such that for all $X \in \mathcal{X}$ and for all $Y, Z \in \mathcal{X}(X)$, if $Y \pitchfork Z$ and $x \in \mathcal{M}(X) \sqcup \mathcal{X}(X)$, then

$$\min\{\sigma_Y(x, Z), \sigma_Z(x, Y)\} \leq r_0.$$

- (A10) “realisation”: there exists r_0 such that for all $X \in \mathcal{X}$, 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 $a_Y \in \mathcal{M}(Y)$, then there is some $a \in \mathcal{M}(X)$ with $\rho(a_Y, \psi_Y a) \leq r_0$ and $R_X(a, Y) \leq r_0$ for all $Y \in \mathcal{X}$.

In fact, for (A10) it would be enough to take \mathcal{Y} to consist of an annular subsurface together with any non- $S_{0,3}$ complementary components — we can then keep cutting the surface into smaller and smaller pieces, and the general case follows by an inductive application of (A4) “composition”.

Note that, using by (A2) “ χ Lipschitz and cobounded” and (A3) “ ψ Lipschitz” we have a reverse inequality in (A7) “distance bound”, namely that $R_X(a, b)$ is (linearly) bounded above in terms of $\rho(a, b)$.

We note that (A6) “finiteness” and (A7) “distance bound” are both consequences of the following distance formula.

Given $r \geq 0$, we write

$$\mathcal{A}_X(a, b; r) = \{Y \in \mathcal{X}(X) \mid \sigma_Y(a, b) > r\}.$$

Given $a, b \in \mathbb{M}^0(\Sigma)$ and $r \geq 0$, let $D_X(a, b; r) = \sum_{Y \in \mathcal{A}_X(a, b; r)} \sigma_Y(a, b)$.

We suppose:

(B1) “distance formula”: there exists $r_0 \geq 0$ such that for all $r \geq r_0$, there exist $k_1 > 0$, $h_1, k_2, h_2 \geq 0$ such that for all $X \in \mathcal{X}$ and for all $a, b \in \mathcal{M}(X)$,

$$k_1 \rho(a, b) - h_1 \leq D_X(a, b; r) \leq k_2 \rho(a, b) + h_2.$$

We will sometimes abbreviate this statement to $D_X(a, b; r) \asymp \rho(a, b)$.

Less formally, this says that distances in $\mathcal{M}(X)$ agree to within linear bounds with the sum of all sufficiently large projected distances in $\mathcal{G}(Y)$ as Y ranges over subsurfaces of X . Here “sufficiently large” implies a lower threshold below which we ignore any contributions. The linear bounds will depend on the particular choice of threshold. For this to work, the threshold must be assumed sufficiently large.

In the case of markings, (B1) is the distance formula of [Masur and Minsky 2000], who also stated it for the pants complex. (We remark that for the Teichmüller metric, a similar formula has been proven by Rafi and by Durham, and is used in [Bowditch 2016a]. A more general version, which encompasses these cases is proven in [Behrstock, Hagen and Sisto 2015].) Again for markings, (A8) “bounded image” is a consequence of their bounded geodesic image theorem, (A9) “overlapping subsurfaces” is a consequence of Behrstock’s lemma, and (A10) “realisation” is a simple explicit construction. We will elaborate on this in Section 8.

Given $Y \in \mathcal{X}$, we write $\mu_Y : (\mathcal{G}(Y))^3 \rightarrow \mathcal{G}(Y)$ for the usual median (or “centroid”) operation on the uniformly hyperbolic space $\mathcal{G}(Y)$. (That is, $\mu(a, b, c)$ is a bounded distance from any geodesic connecting any two distinct points of $\{a, b, c\}$.)

We will show:

Theorem 7.1. *Under the hypotheses (A1)–(A10) above, there is some $t_0 \geq 0$ depending only on the parameters of the hypotheses such that if $X \in \mathcal{X}$ and*

$a, b, c \in \mathcal{M}(X)$, there is some $m \in \mathcal{M}(X)$ such that for all $Y \in \mathcal{X}(X)$ we have $\sigma(\theta_Y m, \mu_Y(\theta_Y a, \theta_Y b, \theta_Y c)) \leq t_0$.

By (A7) “distance bound”, m is well defined up to bounded distance. We set $\mu_X(a, b, c) = m$ for some such m to give us a ternary operation $\mu_X : (\mathcal{M}(X))^3 \rightarrow \mathcal{M}(X)$. Using (A4) “composition”, we see that if $Y \in \mathcal{X}(X)$, then $\psi_Y : \mathcal{M}(X) \rightarrow \mathcal{M}(Y)$ is a uniform quasimorphism, that is,

$$\rho_Y(\mu_Y(\psi_Y a, \psi_Y b, \psi_Y c), \psi_Y \mu_X(a, b, c)) \leq h \quad \text{for all } a, b, c \in \mathcal{M}(X),$$

where $h \geq 0$ depends only on the parameters of the hypotheses.

Theorem 7.2. *Under the hypotheses (A1)–(A10), there is a ternary operation, μ_X , defined on each space $\mathcal{M}(X)$ such that $(\mathcal{M}(X), \rho_X, \mu_X)$ is a coarse median space, and such that the maps $\theta_{YX} : \mathcal{M}(X) \rightarrow \mathcal{G}(Y)$ for $Y \leq X$ are all median quasimorphisms. The median μ_X is unique with this property, up to bounded distance. The maps $\psi_{YX} : \mathcal{M}(X) \rightarrow \mathcal{M}(Y)$ for $Y \leq X$ are also median quasimorphisms. The coarse median space $(\mathcal{M}(X), \rho_X, \mu_X)$ is finitely colourable, and has rank at most $\xi(X)$. Moreover, all constants of the conclusion (coarse median property, and quasimorphism) depend only on the constants of the hypotheses (A1)–(A10).*

Proof. Given Theorem 7.1, this follows directly from the results of [Bowditch 2013], in particular, Propositions 10.1 and 10.2 thereof. We just need to check that the respective hypotheses (P1)–(P4) and (P3′) are satisfied.

Here, (P1) is (A7) “distance bound” and (P2) is (A1) “hyperbolic”. (P3) is the statement that one can embed at most ξ subsurfaces disjointly in a surface of complexity ξ . Finally, (P4) follows exactly as in Lemma 11.7 there, which only uses properties (A8) “bounded image” and (A9) “overlapping subsurfaces”. Moreover, property (P3′) also holds here, for the same reason. \square

(In some cases, one can improve on the rank bound of Theorem 7.2, as in the case of the Weil–Petersson metric [Bowditch 2015].)

So far, we have made no reference to the action of $\text{Map}(\Sigma)$. In applications, the spaces and maps will be equivariant (up to bounded distance), and it follows that the medians we construct will necessarily also be equivariant up to bounded distance.

We now set about the proof of Theorem 7.1. To simplify the exposition, we will construct the median $\mu = \mu_\Sigma$ on Σ . The same arguments apply working intrinsically in any subsurface $X \in \mathcal{X}_N$.

We begin with some general observations about the treelike (rank-1 median) nature of hyperbolic spaces.

Definition. A *spanning tree* for a finite set A consists of a finite simplicial tree, Δ , and a map $\pi = \pi_\Delta : A \rightarrow V(\Delta)$ to the vertex set.

Recall that the vertex set of a finite simplicial tree is a rank-1 median algebra (and every finite rank-1 median algebra has this form). We can assume that every terminal (i.e., degree-1) vertex of Δ lies in πA . We say that Δ is *trivial* if it is a singleton.

Suppose that T is another spanning tree with an embedding of Δ in T . There is a natural retraction, ω , of T onto Δ , and hence of $V(T)$ to $V(\Delta)$. We say that the spanning tree T is an *enlargement* of Δ if $\pi_\Delta = \omega\pi_T$.

Suppose that $\{\Delta_i\}_{i \in \mathcal{J}}$ is a finite collection of spanning trees for A , indexed by some set \mathcal{J} . We say that a spanning tree T for A is a *common enlargement* of $\{\Delta_i\}_{i \in \mathcal{J}}$ if we can embed the Δ_i simultaneously in T so that their interiors are disjoint, and such that T is an enlargement of each Δ_i . Note that (after collapsing complementary trees), we may as well suppose that $T = \bigcup_{i \in \mathcal{J}} \Delta_i$. We write $T = T(\{\Delta_i\}_{i \in \mathcal{J}})$. (There may be some ambiguity, in that we may be able to swap to trees each consisting of single edge, and meeting at a vertex not in πA . However, this ambiguity will not matter to us.)

Definition. We say that a collection of spanning trees is *coherent* if it has a common enlargement.

We shall assume henceforth that all our spanning trees are nontrivial.

Lemma 7.3. *Two spanning trees Δ_0 and Δ_1 are coherent if and only if there are vertices, $v_{01} \in V(\Delta_0)$ and $v_{10} \in V(\Delta_1)$ such that $A = \pi_0^{-1}v_{01} \cup \pi_1^{-1}v_{10}$.*

Proof. If $T = T(\Delta_0, \Delta_1)$ is a common spanning tree for A , then T is obtained by taking $\Delta_0 \sqcup \Delta_1$ and identifying a vertex $v_{01} \in V(\Delta_0)$ with $v_{10} \in V(\Delta_1)$, to give a vertex $w \in V(T)$. Note that $\pi : A \rightarrow T$ is given by $\pi|(A \setminus \pi_0^{-1}v_{01}) = \pi_0$, $\pi|(A \setminus \pi_1^{-1}v_{10}) = \pi_1$ and $\pi(\pi_0^{-1}v_{01} \cap \pi_1^{-1}v_{10}) = \{w\}$. We can clearly invert the above process. \square

Suppose that $\{\Delta_0, \Delta_1, \Delta_2\}$ are coherent. Let $T = T(\Delta_0, \Delta_1, \Delta_2)$. Up to permutation of indices, there are two possibilities:

- (1) $\Delta_0, \Delta_1, \Delta_2$ meet at a common vertex $w = V(T)$. In this case, $v_{01} = v_{02}$, $v_{12} = v_{10}$ and $v_{20} = v_{21}$. Note that these vertices all get identified to w in T .
- (2) Δ_1 and Δ_2 do not meet in T . In this case, $v_{01} \neq v_{02}$, $v_{12} = v_{10}$ and $v_{20} = v_{21}$.

Note that the conditions on vertices above make sense if we assume only that Δ_0, Δ_1 and Δ_2 are pairwise coherent.

Lemma 7.4. *Let $\{\Delta_0, \Delta_1, \Delta_2\}$ be pairwise coherent. Then it is coherent if and only if at most one of the three equalities $v_{01} = v_{02}$, $v_{12} = v_{10}$ and $v_{20} = v_{21}$ does not hold.*

Proof. We have explained “only if”, so we prove “if”:

(1) Suppose all the equalities hold. Let $w_0 = v_{01} = v_{02}$, $w_1 = v_{12} = v_{10}$ and $w_2 = v_{20} = v_{21}$. Let T be obtained from $\Delta_0 \sqcup \Delta_1 \sqcup \Delta_2$ by identifying w_0, w_1 and w_2 to a single point $w \in V(T)$. We define $\pi : A \rightarrow V(T)$ by $\pi|_{(A \setminus \pi_i^{-1}w_i)} = \pi_i$ for $i = 0, 1, 2$ and setting $\pi(\pi_0^{-1}w_0 \cap \pi_1^{-1}w_1 \cap \pi_2^{-1}w_2) = \{w\}$.

(2) If not, then, without loss of generality, $v_{01} \neq v_{02}$. Let $w_1 = v_{12} = v_{10}$ and $w_2 = v_{20} = v_{21}$. We construct T from $\Delta_0 \sqcup \Delta_1 \sqcup \Delta_2$ by identifying v_{01} with w_1 to give $x_1 \in V(T)$ and v_{02} with w_2 to give $x_2 \in V(T)$. Note that A can be partitioned into five disjoint sets:

$$\begin{aligned} A_1 &= \pi_0^{-1}v_{01} \setminus \pi_1^{-1}w_1, & A_{01} &= \pi_0^{-1}v_{01} \cap \pi_1^{-1}w_1, & A_0 &= \pi_1^{-1}w_1 \cap \pi_2^{-1}w_2, \\ A_{02} &= \pi_0^{-1}v_{02} \cap \pi_2^{-1}w_2, & A_2 &= \pi_0^{-1}v_{02} \setminus \pi_2^{-1}w_2. \end{aligned}$$

We define $\pi : A \rightarrow V(T)$ by setting $\pi|_{A_i} = \pi_i$ for $i = 0, 1, 2$ and setting $\pi(A_{01}) = x_1$ and $\pi(A_{02}) = x_2$. \square

In fact, three trees are enough: a finite collection of spanning trees for A is coherent if and only if every subset of at most three elements is coherent. This is not hard to verify, but since we won't be needing it, we omit the proof.

We now move on to consider hyperbolic spaces. Recall that

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

for the Gromov product.

Lemma 7.5. *Suppose that (G, σ) is k -hyperbolic, $p \in \mathbb{N}$, and $t \geq 0$. Given a set $B \subseteq G$ with $|B| \leq p$, there is a simplicial tree, Δ , and maps $\pi : B \rightarrow V(\Delta)$ and $\lambda : V(\Delta) \rightarrow G$ such that for all $x, y, z \in V(\Delta)$, if $\langle \lambda x, \lambda y : \lambda z \rangle \leq t$, then $z \in [x, y]_{V(\Delta)}$. Moreover, λ is an h -quasimorphism and for all $x \in B$ we have $\sigma(x, \lambda \pi x) \leq h$, where h depends only on k, p and t .*

Proof. This is proven in [Bowditch 2013, Lemma 10.3]. It is a simple consequence of the fact that any finite set of points in a Gromov hyperbolic space can be approximated up to an additive constant by a finite tree (with vertex set B). The additive constant depends only on p and k . For the clause about Gromov products we need to collapse down ‘‘short’’ edges of the tree (hence the dependence of h on s). This can also be phrased in terms of Corollary 6.5 here. (In [Bowditch 2013] we had a stronger condition on the ‘‘crossratios’’ of four points of B , which is easily seen to imply the condition on Gromov products given here.) \square

We will apply this to the spaces $\mathcal{G}(X)$ featuring in the hypotheses of Theorem 7.1. By (A1) these are all k -hyperbolic. Recall that we have maps $\theta_X : \mathcal{M}(X) \rightarrow \mathcal{G}(X)$.

We fix some $A \subseteq \mathcal{M}(\Sigma)$ with $|A| = p < \infty$. (In our applications here we will have $p \leq 4$, but we can keep the discussion general for the moment.) We will choose universal $t \geq 0$ sufficiently large (depending only on p) as described

below. We apply Lemma 7.5 to $B = \theta_X(A) \subseteq \mathcal{G}(X)$ with t as above, to get a tree $\Delta(X)$ and maps $\pi : B \rightarrow V(\Delta(X))$ and $\lambda_X = \lambda : V(\Delta(X)) \rightarrow \mathcal{G}(X)$. We set $\pi_X = \pi \circ \theta_X : A \rightarrow V(\Delta(X))$.

All we require of this until Lemma 7.11, is:

(*) If $a, b, c \in A$ with $\langle \theta_X a, \theta_X b : \theta_X c \rangle \leq t$, then $\pi_X c \in [\pi_X a, \pi_X b]_{V(\Delta(X))}$.

In particular, if $\sigma(\theta_X a, \theta_X b) \leq t$, then $\pi_X a = \pi_X b$ (since $\langle \theta_X a, \theta_X a : \theta_X b \rangle \leq t$, so $\pi_X b \in [\pi_X a, \pi_X a] = \{\pi_X a\}$). It follows that if $\text{diam}(\theta_X A) \leq t(p)$, then $\Delta(X)$ is trivial (i.e., a singleton).

For future reference (see Lemma 7.11) we also note that λ is an l -quasimorphism, and that for all $a \in A$, $\sigma(\theta_X a, \lambda_X \pi_X a) \leq l$, where $l = h(p)$ depends only on p .

Lemma 7.6. *Let $X, Y \in \mathcal{X}$ with $X \pitchfork Y$, then there are points, $v_{XY} \in V(\Delta(X))$ and $v_{YX} \in V(\Delta(Y))$ such that $A = \pi_X^{-1} v_{XY} \cup \pi_Y^{-1} v_{YX}$.*

Proof. We can assume that neither $V(\Delta(X))$ nor $V(\Delta(Y))$ is trivial. Note that if $a \in A$, with $\sigma(\theta_X a, \theta_X Y) > r_0$, then $\sigma(\theta_Y a, \theta_Y X) \leq r_0$. If this were true for all $a \in A$, we would conclude that $\text{diam}(\theta_Y A) \leq 2r_0 < t(p)$ giving the contradiction that $V(\Delta(Y))$ is trivial. We can thus find $a_{XY} \in A$ with $\sigma(\theta_X a_{XY}, \theta_X Y) \leq r_0$. We set $v_{XY} = \pi_X a_{XY} \in V(\Delta(X))$. We similarly define $v_{YX} = \pi_Y a_{YX} \in V(\Delta(Y))$.

Now suppose that $b \in A \setminus (\pi_X^{-1} v_{XY} \cup \pi_Y^{-1} v_{YX})$. Then $\pi_X b \neq \pi_X a_{XY}$, and so $\sigma(\theta_X b, \theta_X a_{XY}) \geq t(p)$. Thus, $\sigma(\theta_X b, \theta_X Y) \geq t(p) - r_0 > r_0$. Similarly, $\sigma(\theta_Y b, \theta_Y X) > r_0$. This contradicts property (A9) ‘‘overlapping subsurfaces’’, proving that no such b exists. \square

Note that, by Lemma 7.3, we can naturally combine $\Delta(X)$ and $\Delta(Y)$ into a larger tree by identifying the vertices v_{XY} and v_{YX} . In other words, $\{\Delta(X), \Delta(Y)\}$ is coherent. We write $\Delta(X, Y)$ for the common enlargement.

Note that by construction if $\Delta(X)$ and $\Delta(Y)$ are nontrivial, then $\sigma(\theta_X a_{XY}, \theta_X Y) \leq r_0$, where $a_{XY} \in A$ is as in the proof of Lemma 7.6. By the same argument, if $Z \in \mathcal{X}$ with $\Delta(Z)$ nontrivial, we have $\sigma(\theta_X a_{XZ}, \theta_X Z) \leq r_0$, for some $a_{XZ} \in A$. If $\sigma(\theta_X Y, \theta_X Z) < t - 2r_0$, then $\sigma(\theta_X a_{XY}, \theta_X a_{XZ}) < s(p)$, so $v_{XY} = \pi_X a_{XY} = \pi_X a_{XZ} = v_{XZ}$. For future reference (Lemma 7.11) we also note that

$$\sigma(\theta_X a_{XY}, \lambda_X v_{XY}) = \sigma(\theta_X a_{XY}, \lambda_X \pi_X a_{XY}) \leq l,$$

so $\sigma(\theta_X Y, \lambda_X v_{XY}) \leq r_0 + l$.

We write \mathcal{X}_0 for the set of $X \in \mathcal{X}$ such that $\Delta(X)$ is nontrivial. It follows from property (A6) ‘‘finiteness’’, that \mathcal{X}_0 is finite.

Note that if $X, Y \in \mathcal{X}_0$ and $X \pitchfork Y$, then $\{\Delta(X), \Delta(Y)\}$ is coherent. This is an immediate consequence of Lemmas 7.4 and 7.6. Note that this determines vertices $v_{XY} \in \Delta(X)$ and $v_{YX} \in \Delta(Y)$ which get identified in $\Delta(X, Y)$.

Lemma 7.7. *Suppose that $X, Y, Z \in \mathcal{X}_0$ and that $X \pitchfork Y$ and $X \pitchfork Z$ and $v_{XY} \neq v_{XZ}$. Then $Y \pitchfork Z$.*

Proof. If not, then (since there must be boundary curves of Y and Z which are disjoint) and by (A5) “disjoint projection” we must have $\sigma(\theta_X Y, \theta_X Z) \leq r$, for some constant r depending only on that of (A5), which in turn depends only (or at most) on $\xi(\Sigma)$. Provided we have chosen $t > l + 2r_0$, this implies that $v_{XY} = v_{XZ}$. \square

Lemma 7.8. *Suppose that $X, Y, Z \in \mathcal{X}_0$ and that $X \pitchfork Y$, $X \pitchfork Z$ and $Y \pitchfork Z$. Then $\{\Delta(X), \Delta(Y), \Delta(Z)\}$ is coherent.*

Proof. By Lemma 7.7, it’s enough to show that at least two of $v_{XY} = v_{XZ}$, $v_{YZ} = v_{YX}$, $v_{ZX} = v_{ZY}$ must hold.

By property (A9) “overlapping subsurfaces”,

$$\min\{\sigma(\theta_X Y, \theta_X Z), \sigma(\theta_Y X, \theta_Y Z)\} \leq r_0.$$

Therefore, if $t \geq 3r_0$, we see that either $v_{XY} = v_{XZ}$ or $v_{YZ} = v_{YX}$. Similarly, we have $(v_{YZ} = v_{YX} \text{ or } v_{ZX} = v_{ZY})$ and $(v_{ZX} = v_{ZY} \text{ or } v_{XY} = v_{XZ})$, and so the statement follows. \square

We can now start on the proof of Theorem 7.1

Suppose $a, b, c \in \mathcal{M}(\Sigma)$. We want to find a median for a, b, c in $\mathcal{M}(\Sigma)$. First choose any $d \in \mathcal{M}(\Sigma)$ with $\sigma_\Sigma(\theta_\Sigma d, \mu_\Sigma(\theta_\Sigma a, \theta_\Sigma b, \theta_\Sigma c))$ bounded in $\mathcal{G}(\Sigma)$. (Such a d exists, since $\chi_\Sigma(\mathcal{M}(\Sigma))$ is cobounded in $\mathcal{G}(\Sigma)$ by (A2) “ χ Lipschitz an cobounded”.)

Now set $A = \{a, b, c, d\}$, and let $\pi_X : A \rightarrow \Delta(X)$ be as described in Lemma 7.5. Let $h = h(4)$. Write $d_X = \pi_X d$ and $e_X = \mu_X(\pi_X a, \pi_X b, \pi_X c)$. Recall that \mathcal{X}_0 is the (finite) set of $X \in \mathcal{X}$ such that $\Delta(X)$ is nontrivial. Let $\mathcal{X}_1 = \{X \in \mathcal{X}_0 \mid e_X \neq d_X\}$. By the choice of d , we see that $\Sigma \notin \mathcal{X}_1$.

Suppose that $X, Y \in \mathcal{X}_1$ with $X \pitchfork Y$. Recall that $T = \Delta(X, Y)$ is obtained by identifying $v_{XY} \in \Delta(X)$ with $v_{YX} \in \Delta(Y)$, to give $w \in T$. Note that $\pi_T d$ and $\mu_T(\pi_T a, \pi_T b, \pi_T c)$ must be distinct from w , and must lie in different subtrees $\Delta(X)$ and $\Delta(Y)$. It follows that exactly one of the following must hold:

- (1) $d_Y = v_{YX}$ and $e_X = v_{XY}$, or
- (2) $d_X = v_{XY}$ and $e_Y = v_{YX}$.

We write these cases respectively as $X \ll Y$ and $Y \ll X$ (which we take to imply that $X \pitchfork Y$).

(Intuitively, we think of these relations as follows. We imagine any finite set of elements of \mathcal{X} embedded disjointly as “horizontal” surfaces in $\Sigma \times \mathbb{R}$; that is, $X \in \mathcal{X}$ is identified with $X \times \{x\}$ for some $x \in \mathbb{R}$. The relations $=, <, \wedge$ and \pitchfork have their usual meaning on projecting to Σ , and $X \ll Y$ means that $X \pitchfork Y$ and X is “to the left” of Y in the sense that it has smaller \mathbb{R} -coordinate. The relations are well defined up

to isotopy, and satisfy the same properties as those laid out here. This picture ties in with the Minsky model for hyperbolic 3-manifolds homeomorphic to $\Sigma \times \mathbb{R}$.)

Lemma 7.9. *If $X, Y, Z \in \mathcal{X}_1$ and $X \ll Y$ and $Y \ll Z$, then $X \ll Z$.*

Proof. Since $X \ll Y$, $v_{XY} = d_Y$. Since $Y \ll Z$, $v_{YZ} = d_Z$. Since $Y \in \mathcal{X}_1$, $d_Y \neq e_Y$, so $v_{YX} \neq v_{YZ}$. By Lemmas 7.7 and 7.8, $X \pitchfork Z$, and $\{\Delta(X), \Delta(Y), \Delta(Z)\}$ is coherent. In particular, $e_X = v_{XY} = v_{XZ}$ and $d_Z = v_{YZ} = v_{ZX}$ so $X \ll Z$. \square

Recall that $X < Y$ implies that $X \neq Y$ and X is homotopic into Y . We therefore have two strict partial orders \ll and $<$ on \mathcal{X}_1 . Moreover, by hypothesis, $X \ll Y$ is incompatible with any of $X < Y$, $Y < X$, or $X \wedge Y$.

Lemma 7.10. *Given $X, Y, Z \in \mathcal{X}_1$ with $X \ll Y$ and $Y < Z$, then either $X \ll Z$ or $X < Z$.*

Proof. Recall that $X \pitchfork Z$ implies $X \ll Z$ or $Z \ll X$. Thus, if the conclusion of the lemma fails, the only alternatives would be $Z = X$, $Z < X$, $Z \ll X$ or $Z \wedge X$. Now $Z = X$ or $Z < X$ both give $Y < X$ contradicting $X \ll Y$; $Z \ll X$ gives $Z \ll Y$ contradicting $Y < Z$, and finally, $Z \wedge X$ gives $Y \wedge X$, contradicting $X \ll Y$. \square

Now write $X < Y$ to mean that either $X \ll Y$ or $X < Y$. This relation is antisymmetric on \mathcal{X}_1 . It is not in general transitive, but in view of Lemma 7.10, any relation of the form $X < Y < Z < W$ can be reduced to $X < V < W$ for $V \in \{Y, Z\}$. In particular, there are no cycles. It follows that \mathcal{X}_1 contains an element U which is maximal with respect to this relation. In other words, if $X \in \mathcal{X}_1$, then we have neither $U \ll X$ nor $U < X$. Note that $\Sigma \notin \mathcal{X}_1$, so $U \neq \Sigma$.

From this, we can deduce:

Lemma 7.11. *There is some universal $u_0 > 0$, such that if $a, b, c \in \mathcal{M}(\Sigma)$, there is some curve α such that if $X \in \mathcal{X}$, with $\alpha \pitchfork X$ or $\alpha < X$, then*

$$\sigma(\theta_X \alpha, \mu_X(\theta_X a, \theta_X b, \theta_X c)) \leq u_0.$$

Proof. Let $U \in \mathcal{X}_1$ be maximal with respect to $<$, as above. Let α be a component of the relative boundary of U in Σ . Suppose that $X \in \mathcal{X}$ with $\alpha < X$ or $\alpha \pitchfork X$. Then either $U < X$ or $U \pitchfork X$. According to the conventions described in Section 6, we use the notation \sim to mean ‘‘up to bounded distance’’. In all cases, $\theta_X \alpha$ is defined and $\theta_X \alpha \sim \theta_X U$, by (A5) ‘‘disjoint projection’’. Let $\lambda_X : V(\Delta_X) \rightarrow \mathcal{G}(X)$ be the quasimorphism described above (as given by Lemma 7.5). Now,

$$\begin{aligned} \lambda_X e_X &= \lambda_X \mu_{V(\Delta(X))}(\pi_X a, \pi_X b, \pi_X c) \\ &\sim \mu_X(\lambda_X \pi_X a, \lambda_X \pi_X b, \lambda_X \pi_X c) \\ &\sim \mu_X(\theta_X a, \theta_X b, \theta_X c). \end{aligned}$$

We therefore want to show that $\theta_X U \sim \lambda_X e_X$. Note that $\lambda_X d_X = \lambda_X \pi_X d \sim \theta_X d$.

Suppose first that $U \prec X$. Thus $X \notin \mathcal{X}_1$, so $d_X = e_X$. Now $\lambda_X e_X = \lambda_X d_X$, is a centroid for $\theta_X a, \theta_X b, \theta_X c$ in $\mathcal{G}(X)$, and so $\theta_X d$ is also a centroid. Therefore, if $\theta_X U$ were far enough away from $\theta_X d$, (depending only on the hyperbolicity constant), then we can assume that the Gromov products $\langle \theta_X a, \theta_X b; \theta_X U \rangle$ and $\langle \theta_X a, \theta_X d; \theta_X U \rangle$ are both greater than r_0 (after permuting a, b, c as necessary). By property (A8) “bounded image”, this implies that $\sigma_U(\theta_U a, \theta_U b) \leq r_0$ and $\sigma_U(\theta_U a, \theta_U d) \leq r_0$. It then follows that $\pi_U a = \pi_U b = \pi_U d \in V(\Delta(U))$, so $e_U = \mu_{V(\Delta(U))}(\pi_U a, \pi_U b, \pi_U c) = \pi_U d = d_U$, contradicting the fact that $U \in \mathcal{X}_1$. We have shown that if $U \prec X$, then $\theta_X U \sim \lambda_X e_X$ as required.

Suppose now that $U \pitchfork X$. In this case, by Lemma 7.6, the trees $\Delta(X)$ and $\Delta(U)$ are coherent. Moreover, since $\Delta(U)$ is nontrivial, we have $\theta_X U \sim \lambda_X v_{XU}$. If $X \in \mathcal{X}_1$, then $X \ll U$, so $e_X = v_{XU}$, thus $\theta_X U \sim \lambda_X v_{XU} = \lambda_X e_X$ as required. So we can suppose that $X \notin \mathcal{X}_1$ — in other words, $d_X = e_X$. If $X \notin \mathcal{X}_0$, then $\Delta(X)$ is trivial, so $e_X = d_X = v_{XU}$, and we are done, as above. If $X \in \mathcal{X}_0$, then again $d_X = v_{XU}$, otherwise we would get $e_U = d_U$ contradicting $U \in \mathcal{X}_1$.

In all cases, we have shown that $\theta_X U \sim \lambda_X e_X$, as required. \square

We can now prove the main result of this section:

Proof of Theorem 7.1. Uniqueness up to bounded distance is an immediate consequence of property (A7) “distance bound” here, so we prove existence.

Let $a, b, c \in \mathcal{M}(\Sigma)$. Let α be a curve as given by Lemma 7.11. Given $X \in \mathcal{X}$, write $\delta_X = \mu_X(\theta_X a, \theta_X b, \theta_X c) \in \mathcal{G}(X)$. So $\theta_X \alpha \sim \delta_X$ for all X with $\alpha \pitchfork X$ or $\alpha \prec X$. We consider only the case when α separates Σ . The nonseparating case is essentially the same.

Let $\Sigma = Y \cup Z$, where $Y \cap Z = \alpha$. Suppose first that neither Y nor Z is a $S_{0,3}$, so that $Y, Z \in \mathcal{X}$. By induction on the complexity of Σ , we can assume that Theorem 7.1 holds intrinsically to Y and Z . Thus, we can find $m_Y \in \mathcal{M}(Y)$, such that if $X = Y$ or $X \prec Y$, then $\sigma(\theta_X m_Y, \delta_X)$ is bounded. We have a similar element, $m_Z \in \mathcal{M}(Z)$. Let $\Omega \in \mathcal{X}_A$ be the annulus with core curve α . We apply property (A10) “realisation” with $\mathcal{Y} = \{X, Y, \Omega\}$ to give $m \in \mathcal{M}(\Sigma)$ such that $\rho(\psi_Y m, m_Y)$, $\rho(\psi_Z m, m_Z)$ and $\rho_\Omega(\psi_\Omega m, \delta_\Omega)$ are bounded. By (A4) “composition” and the construction of m_Y and m_Z , we have $\theta_X m \sim \delta_X$ for all $X \preceq Y$ all $X \preceq Z$.

Suppose that $X \in \mathcal{X}$. If $X \preceq Y$, $X \preceq Z$ or $X = \Omega$, then $\sigma(\theta_X m, \delta_X)$ is bounded by construction. If not, then either $\alpha \prec X$ or $\alpha \pitchfork X$. But then, by the choice of α , $\sigma(\theta_X \alpha, \delta_X)$ is bounded as already observed. But $\sigma(\theta_X m, \theta_X \alpha) \leq R_\Sigma(m, \alpha)$ is bounded by (A10) “realisation”, so we are done in this case.

If either Y or Z is an $S_{0,3}$, we just omit that subsurface from \mathcal{Y} , and proceed in the same way.

If α does not separate, we set \mathcal{Y} to consist of $X(\alpha)$ together with complement of α and proceed similarly. \square

8. The marking complex

In this section, we apply the results of Theorem 7.1 to the marking complex of Σ , to recover the result of [Behrstock and Minsky 2011], stated as Theorem 8.2 here. We first describe the curve graph associated to a compact surface, Σ .

For $\xi(\Sigma) \geq 1$, let $\mathbb{G} = \mathbb{G}(\Sigma)$ be the curve graph of Σ . Its vertex set, \mathbb{G}^0 , is the set of free homotopy classes of essential nonperipheral simple closed curves in Σ . As before, we refer to elements of \mathbb{G}^0 simply as *curves*. Two curves, $\alpha, \beta \in \mathbb{G}^0$ are adjacent if $\iota(\alpha, \beta)$ is equal to 2 if Σ is an $S_{0,4}$; equal to 1 if Σ is an $S_{1,1}$; or equal to 0 if $\xi(\Sigma) \geq 2$. Here $\iota(\alpha, \beta)$ denotes the geometric intersection number.

In all cases, $\mathbb{G}(\Sigma)$ is connected. A key result in the subject is:

Theorem 8.1. *There is a universal constant k such that for any compact surface Σ , $\mathbb{G}(\Sigma)$ is k -hyperbolic.*

The existence of such a k , depending on $\xi(\Sigma)$, was proven by Masur and Minsky [1999]. The fact that it is uniform (independent of $\xi(\Sigma)$) was proven independently in [Aougab 2013; Bowditch 2014b; Clay, Rafi and Schleimer 2014; Hensel, Przytycki and Webb 2015]. (The uniformity is not essential to the main results of this paper: we will only be dealing with finitely many topological types at any given time, namely subsurfaces of a given surface Σ . One can therefore simply assert dependence of constants on $\xi(\Sigma)$ at the relevant points.)

Given nonempty $a, b \subseteq \mathbb{G}^0$, let $\iota(a, b) = \max\{\iota(\alpha, \beta) \mid \alpha \in a, \beta \in b\}$. We write $\iota(a) = \iota(a, a)$.

Definition. If $\iota(a) = 0$, we refer to a as a *multicurve*.

Intuitively, we think of a multicurve in terms of its realisation as 1-manifold in M .

Definition. We say that $a \subseteq \mathbb{G}^0$ *fills* Σ if $\iota(a, \gamma) \neq 0$ for all $\gamma \in \mathbb{G}^0$.

If we realise a minimally, then this is the same as saying that all complementary components of $\bigcup a$ are disc or peripheral annuli.

Given $p, q \in \mathbb{N}$, define a graph $\mathbb{M} = \mathbb{M}(\Sigma, p, q)$ by taking the vertex set, \mathbb{M}^0 to be the set of $a \subseteq \mathbb{G}^0$ such that a fills Σ and $\iota(a) \leq p$, and by deeming $a, b \in \mathbb{M}^0$ to be adjacent if $\iota(a, b) \leq q$. This graph is always locally finite. Provided p is large enough and q is large enough in relation to p (independently of Σ) it will always be nonempty and connected. For definiteness, we can set $\mathbb{M}(\Sigma) = \mathbb{M}(\Sigma, 2, 4)$, though the actual choice will not matter. (The inclusion of $\mathbb{M}(\Sigma, 2, 4)$ into any larger $\mathbb{M}(\Sigma, p, q)$ is a quasi-isometry.)

Definition. We refer to $\mathbb{M}(\Sigma)$ as the *marking graph* of Σ .

(This is a slight variation on the marking complex of [Masur and Minsky 2000].)

Note that the mapping class group $\text{Map}(\Sigma)$ acts on $\mathbb{G}(\Sigma)$ and on $\mathbb{M}(\Sigma)$ with finite quotient. In particular, we see that $\text{Map}(\Sigma)$ is quasi-isometric to $\mathbb{M}(\Sigma)$. Note

also that bounding distance in the marking complex is equivalent to bounding intersection numbers between markings.

Recall that $\mathcal{X} = \mathcal{X}_A \sqcup \mathcal{X}_N$ is the set of (non- $S_{0,3}$) subsurfaces of Σ , partitioned into annular and nonannular subsurfaces.

If $X \in \mathcal{X}_N$, we can define $\mathbb{G}(X)$ and $\mathbb{M}(X)$ intrinsically to X as above. If $X \in \mathcal{X}_A$, one needs to define $\mathbb{G}(X)$ as an arc complex in the annular cover of Σ corresponding to X ; see [Masur and Minsky 2000, Section 2.4]. This is quasi-isometric to the real line. In this case, we set $\mathbb{M}(X) = \mathbb{G}(X)$. (One could give a unified description in terms of covers of Σ corresponding to subsurfaces, though we will omit discussion of that here.) We will write $\mathbb{G}(\gamma) = \mathbb{G}(X)$ and $\mathbb{M}(\gamma) = \mathbb{M}(X)$, when $\gamma \in \mathbb{G}^0$, where $X = X(\gamma)$ is the annular neighbourhood of γ .

We will write $\sigma = \sigma_X$ and $\rho = \rho_X$ respectively for the combinatorial metrics on $\mathbb{G}(X)$ and $\mathbb{M}(X)$.

Given $X \in \mathcal{X}$ we have a map $\chi_X : \mathbb{M}(X) \rightarrow \mathbb{G}(X)$. If $X \in \mathcal{X}_A$, this is the identity. If $X \in \mathcal{X}_N$, it just chooses some curve from the marking. Up to bounded distance, the map χ_X is determined by the fact that $\iota(a, \chi_X a)$ is bounded for all $a \in \mathbb{M}(X)$.

Given $X, Y \in \mathcal{X}$ with $Y \preceq X$, we have a subsurface projection $\psi_{YX} : \mathbb{M}(X) \rightarrow \mathbb{M}(Y)$. This is the same construction as in [Masur and Minsky 2000]. We realise a and Y in minimal general position (so that $a \cap Y$ has a minimal number of components). Now $a \cap Y$ consists of a collection of arcs and curves. We say two arcs are “parallel” if they are homotopic, sliding the endpoints in the boundary components of Y . For each parallel class of arcs we get a disjoint curve (namely the boundary component of a regular neighbourhood of the arc union the boundary components it meets). The collection of such curves, together with the curves of a already lying in Y , give us a collection of curves of Y of bounded self-intersection, and hence give rise to a marking of Y . We write this as $\psi_{YX} a$. Up to bounded distance, the map ψ_{YX} is determined by the fact that the intersection of $\psi_{YX} a$ with every component of $a \cap X$ is bounded.

We set $\theta_Y = \chi_Y \circ \psi_{YX} : \mathbb{M}(X) \rightarrow \mathbb{G}(Y)$.

One can also define subsurface projection for curves. Suppose $\gamma \in \mathbb{G}^0(\Sigma)$ and $X \in \mathcal{X}$ with $\gamma \pitchfork X$ or $\gamma \prec X$, then we can define $\theta_X(\gamma) \in \mathbb{G}(X)$. This is consistent with that already defined, in that if $\gamma \in a \in \mathbb{M}^0(X)$, then $\theta_X(\gamma) \sim \theta_X(a)$. In particular, $\theta_X \circ \chi_X(a) \sim \theta_X(a)$ when this is defined. Similarly, if $X, Y \in \mathcal{X}$ with $Y \pitchfork X$ or $Y \prec X$ we can define $\theta_X(Y) \in \mathbb{G}(X)$. This can be defined by setting $\theta_X(Y) = \theta_X(\gamma)$ for some boundary curve, γ , of Y .

We can now deduce the following result [Behrstock and Minsky 2011]:

Theorem 8.2. *There is a constant t_0 , depending only on $\xi(\Sigma)$, such that if $a, b, c \in \mathbb{M}(\Sigma)$, then there is some $m \in \mathbb{M}(\Sigma)$ such that for all $X \in \mathcal{X}(\Sigma)$,*

$$\sigma(\theta_X m, \mu_X(\theta_X a, \theta_X b, \theta_X c)) \leq t_0.$$

Moreover, if $m' \in \mathbb{M}(\Sigma)$ is another such element, then $\rho(m, m') \leq t_1$, where t_1 is a constant depending only on $\xi(\Sigma)$.

We can therefore define a median map $\mu : (\mathbb{M}(\Sigma))^3 \rightarrow \mathbb{M}(\Sigma)$ by $\mu(a, b, c) = m$. Of course, it is enough to define $\mu(a, b, c)$ for a, b, c in the vertex set, $\mathbb{M}^0(\Sigma)$, of $\mathbb{M}(\Sigma)$.

To prove Theorem 8.2, we set $\mathcal{M}(X) = \mathbb{M}(X)$ and $\mathcal{G}(X) = \mathbb{G}(X)$. We verify (A1)–(A10) of Section 7 for these spaces and the maps χ_X and ψ_{YX} defined above. In fact, for (A6) “finiteness” and (A7) “distance bound” we will verify (B1) “distance formula”.

Note that (A1) “hyperbolicity” is an immediate consequence of Theorem 8.1 above. Properties (A2) “ χ Lipschitz and cobounded”, (A3) “ ψ Lipschitz” and (A4) “composition” are elementary properties of subsurface projection, and (A5) “disjoint projection” holds with $s_1 = 1$ in this case.

Property (B1) “distance formula” is an immediate consequence of the following due to Masur and Minsky [2000] (applied intrinsically to subsurfaces).

Theorem 8.3 [Masur and Minsky 2000]. *There is some $r_0 \geq 0$ depending only on $\xi(\Sigma)$ such that for all $r \geq r_0$, there are constants, $k_1 > 0$, $h_1, k_2, h_2 \geq 0$ depending only on r and $\xi(\Sigma)$ such that if $a, b \in \mathbb{M}^0(\Sigma)$, then $k_1 \rho(a, b) - h_1 \leq D_\Sigma(a, b; r) \leq k_2 \rho(a, b) + h_2$.*

This implies (A6) and (A7). Property (A8) “bounded image” is an immediate consequence of the bounded geodesic image theorem; [Masur and Minsky 2000, Theorem 3.1]. (Note that the Gromov product $\langle \alpha, \beta; \gamma \rangle_X$ is, up to an additive constant, the same as the distance from γ to any geodesic from α to β .) A simpler proof of the bounded geodesic image theorem (with uniform constants, independent of $\xi(\Sigma)$) is given in [Webb 2015].

Property (A9) “overlapping subsurfaces” is an immediate consequence of Behrstock’s lemma:

Lemma 8.4. *There is some universal r_0 such that if $X, Y \in \mathcal{X}$ and $\gamma \in \mathbb{G}^0(\Sigma)$ with $X \pitchfork Y$, $\gamma \pitchfork X$ and $\gamma \pitchfork Y$, then $\min\{\sigma(\theta_X(\gamma), \theta_X(Y)), \sigma(\theta_Y(\gamma), \theta_Y(X))\} \leq r_0$.*

This is Theorem 4.3 of [Behrstock 2006] (where r_0 may depend on $\xi(\Sigma)$). A simpler proof, which gives explicit universal constants can be found in [Mangahas 2010].

Property (A10) “realisation” is a simple explicit construction. We can assume that $X = \Sigma$. Let τ be the multicurve consisting of the union of the $\partial_\Sigma Y$, as Y ranges over \mathcal{Y} . Each marking m_Y for $Y \in \mathcal{Y} \cap \mathcal{X}_N$ gives us a marking on some component of $\Sigma \setminus \tau$. We now take the union of τ with the union of all these markings. We add in curves transverse to each of the elements of τ to give us a set of curves which fill Σ with bounded self-intersection. We can arrange (after applying suitable Dehn

twists about the elements of τ) that the marking has the correct projection to the elements of $\mathcal{Y} \cap \mathcal{X}_A$. Note that this construction automatically gives us a marking, m , with $\iota(m, \tau)$ bounded. By construction, $\psi_Y \sim m_Y$ for all $Y \in \mathcal{Y}$. Also if $Z \in \mathcal{X}$ and $Y \in \mathcal{Y}$ with $Y \prec Z$ or $Y \pitchfork Z$, then $\theta_Z m \sim \theta_Z \tau \sim \theta_Z Y$, so $\sigma_Z(m, Y) \sim 0$, and it follows that $R_\Sigma(m, Y) \sim 0$ as required for (A10).

Finally note that bounding the distance, $\rho(a, b)$ between two markings, $a, b \in \mathbb{M}^0(\Sigma)$ is equivalent to bounding their intersection number, $\iota(a, b)$, which in turn, is equivalent to bounding the quantity, $R_\Sigma(a, b)$ featuring in (A7) “distance bounds”. One can find explicit estimates in the references cited, though we will not need them here.

9. Multicurves

Again, in this section, Σ will be a compact surface with $\xi(\Sigma) \geq 1$. We will again assume the hypotheses of Section 7, as we recall below.

Let $\tau \subseteq \mathbb{G}^0(\Sigma)$ be a multicurve in Σ . As usual, we will often identify τ with its realisation as a 1-manifold in Σ . Let $\mathcal{X}_A(\tau) = \{X(\gamma) \in \mathcal{X}_A \mid \gamma \in \tau\}$ be the set of annular surfaces corresponding to the components of τ . Let $\mathcal{X}_N(\tau) \subseteq \mathcal{X}_N$ be the set of components of $\Sigma \setminus \tau$ which are not the $S_{0,3}$. We write $\mathcal{X}(\tau) = \mathcal{X}_A(\tau) \sqcup \mathcal{X}_N(\tau)$.

Given $Y \in \mathcal{X}$, we write $\tau \pitchfork Y$ to mean that $\gamma \pitchfork Y$ or $\gamma \prec Y$ for some $\gamma \in \tau$. Let $\mathcal{X}_T(\tau) = \{Y \in \mathcal{X} \mid Y \pitchfork \tau\}$. Let $\mathcal{X}_I(\tau) = \mathcal{X} \setminus \mathcal{X}_T(\tau)$. It is easily seen that $Y \in \mathcal{X}_I(\tau)$ if and only if $Y \preceq X$ for some $X \in \mathcal{X}(\tau)$. In other words, Y can be homotoped into some component of $\Sigma \setminus \tau$. (This includes the possibility that Y is homotopic to a component of τ .)

Now suppose we have spaces $\mathcal{G}(X)$, $\mathcal{M}(X)$ and maps ψ_{YX} , χ_X , θ_X , etc., satisfying the hypotheses (A1)–(A10) of Section 7. We refer to the constants featuring in these axioms as the *parameters* of $\mathcal{M}(\Sigma)$.

Given $X \in \mathcal{X}_T(\tau)$, we set $\theta_X(\tau) = \theta_X(\gamma)$ for some $\gamma \in \tau$. By (A5) “disjoint projection”, we have $\sigma(\theta_X \tau, \theta_X Y) \leq s_1$ for all $Y \in \mathcal{X}_I(\tau)$. In particular, $\theta_X(\tau)$ is well defined up to bounded distance. As usual, we will abbreviate $\sigma_X(\tau, a) = \sigma(\theta_X \tau, \theta_X a)$ for $a \in \mathcal{M}(\Sigma)$ etc.

Given $a \in \mathcal{M}(\Sigma)$, write

$$R(a, \tau) = \max\{\sigma_X(a, \tau) \mid X \in \mathcal{X}_T(\tau)\}.$$

Thus $R(a, \tau) = \max\{R(a, \gamma) \mid \gamma \in \tau\}$ (cf., the definition of R_Σ in Section 7). (In the case of markings, one can think of this as measuring intersection numbers; see Lemma 9.7.) Given $r \geq 0$, let

$$T(\tau; r) = \{a \in \mathcal{M}(\Sigma) \mid R(a, \tau) \leq r\}.$$

Note that if $a \in T(\tau; r)$ and $Y \in \mathcal{X}_I(\tau)$, then $\sigma(a, Y) \leq r + s_1$ (by (A10) “realisation”). Also, if for each $X \in \mathcal{X}(\tau)$, we have associated some $a_X \in \mathcal{M}(X)$; then by (A10), there is some $a \in T(\tau; r'_0)$ with $\rho(a_X, \psi_X a) \leq r_0$ for all $X \in \mathcal{X}_T(\tau)$, where $r'_0 = r_0 + s_1$. Note that a is well defined up to bounded distance. In fact:

Lemma 9.1. *If $a, b \in T(\tau; r)$ with $\rho_X(a, b) \leq r'$ for all $X \in \mathcal{X}(\tau)$, then $\rho(a, b)$ is bounded above in terms of r and r' .*

Proof. Suppose that $Y \in \mathcal{X}$. If $Y \in \mathcal{X}_I(\tau)$, then $\sigma_Y(a, b)$ is bounded using (A5) “disjoint projection”. If $Y \in \mathcal{X}_T(\tau)$, the $\sigma_Y(a, \tau)$ and $\sigma_Y(b, \tau)$ are both bounded above by hypothesis, so $\sigma_Y(a, b)$ is again bounded. The statement now follows by (A7) “distance bound”. \square

We will abbreviate $T(\tau) = T(\tau; r'_0)$.

Let $\mathcal{T}(\tau) = \prod_{X \in \mathcal{X}(\tau)} \mathcal{M}(X)$. We give $\mathcal{T}(\tau)$ the l^1 metric (though any quasi-isometrically equivalent metric would serve for our purposes). Note that $\mathcal{T}(\tau)$ is a coarse median space, with the median defined coordinatewise. We can combine the maps $\psi_X : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(X)$ for $X \in \mathcal{X}(\tau)$, to give a uniformly coarsely Lipschitz quasimorphism, $\psi_\tau : \mathcal{M}(\Sigma) \rightarrow \mathcal{T}(\tau)$.

By Lemma 9.1 and the subsequent remark, we get a map $\nu_\tau : \mathcal{T}(\tau) \rightarrow T(\tau) \subseteq \mathcal{M}(\Sigma)$, such that $\psi_\tau \circ \nu_\tau : \mathcal{T}(\tau) \rightarrow \mathcal{T}(\tau)$ is the identity up to bounded distance. Note that ν_τ is also a uniformly coarsely Lipschitz quasimorphism, whose image is a uniformly bounded Hausdorff distance from $T(\tau)$.

This in turn gives rise to a quasimorphism, $\omega_\tau = \nu_\tau \circ \psi_\tau : \mathcal{M}(\Sigma) \rightarrow T(\tau)$. It is characterised by the property that $\psi_X \circ \omega_\tau \sim \psi_X$ for all $X \in \mathcal{X}(\tau)$, or equivalently, that $\theta_Y \circ \omega_\tau \sim \theta_Y$ for all $Y \in \mathcal{X}_I(\tau)$. We note:

Lemma 9.2. *Given $r \geq 0$, there is some r' depending only on r and the parameters of the hypotheses, such that for any multicurve, τ , $T(\tau; r) \subseteq N(T(\tau); r')$.*

Proof. Let $b = \omega_\tau(a) \in T(\tau)$. By the above, we have $\theta_Y a \sim \theta_Y b$ for all $Y \in \mathcal{X}_I(\tau)$. Also $\theta_Y a \sim \theta_Y \tau \sim \theta_Y b$ for all $Y \in \mathcal{X}_T(\tau)$. Since $\mathcal{X} = \mathcal{X}_I(\tau) \cup \mathcal{X}_T(\tau)$, we see that $R_\Sigma(a, b)$ is bounded. Property (A7) “distance bound” now tells us that $a \sim b$. \square

This shows that $T(\tau; r)$ is well defined up to bounded Hausdorff distance for all $r \geq r'_0$, and can be described as the set of $a \in \mathcal{M}(\Sigma)$ such that $\theta_Y a \sim \theta_Y \tau$ for all $Y \in \mathcal{X}_T(\tau)$.

Lemma 9.3. *$T(\tau)$ is uniformly quasiconvex in $\mathcal{M}(\Sigma)$.*

Proof. Suppose $a, b \in T(\tau)$ and $c \in \mathcal{M}(\Sigma)$. If $X \in \mathcal{X}_T(\tau)$, then

$$\theta_X \mu_\Sigma(a, b, c) \sim \mu_X(\theta_X a, \theta_X b, \theta_X c) \sim \mu_X(\theta_X \tau, \theta_X \tau, \theta_X c) \sim \theta_X \tau,$$

and so by Lemma 9.2, $\mu_\Sigma(a, b, c)$ is a bounded distance from $T(\tau)$. \square

Lemma 9.4. *The map $\omega_\tau : \mathcal{M}(\Sigma) \rightarrow T(\tau)$ is a coarse gate map.*

Proof. Let $x \in \mathcal{M}(X)$ and $c \in T(\tau)$. If $X \in \mathcal{X}_T(\tau)$, then $\omega_\tau x \in T(\tau)$, and

$$\theta_X \mu_\Sigma(x, \omega_\tau x, c) \sim \mu_X(\theta_X x, \theta_X \omega_\tau x, \theta_X c) \sim \mu_X(\theta_X x, \theta_X \tau, \theta_X \tau) \sim \theta_X \tau \sim \theta_X \omega_\tau x.$$

If $X \in \mathcal{X}_I(\tau)$, then

$$\theta_X \mu_\Sigma(x, \omega_\tau x, c) \sim \mu_X(\theta_X x, \theta_X \omega_\tau x, \theta_X c) \sim \mu_X(\theta_X x, \theta_X x, \theta_X c) \sim \theta_X x \sim \theta_X \omega_\tau x.$$

Since $\mathcal{X} = \mathcal{X}_I(\tau) \cup \mathcal{X}_T(\tau)$, property (A7) “distance bound” tells us $\mu_\Sigma(x, \omega_\tau x, c) \sim \omega_\tau x$, as required. \square

Note that, if $a, b \in \mathcal{M}(\Sigma)$, then $\omega_\tau a$ lies in a coarse interval from a to b , and $\omega_\tau b$ lies on a coarse interval from a to $\omega_\tau a$. By Lemma 6.6, ω_τ is a uniform quasimorphism (depending on $\xi(\Sigma)$). Note that the proof of Lemma 9.4 shows that $\mu_X(\theta_X a, \theta_X \omega_\tau a, \theta_X \omega_\tau b) \sim \theta_X \omega_\tau a$ (putting $x = a$ and $c = b$). Similarly, $\mu_X(\theta_X b, \theta_X \omega_\tau b, \theta_X \omega_\tau a) \sim \theta_X \omega_\tau b$. It follows that $\rho(\omega_\tau a, \omega_\tau b)$ is bounded above by a linear function of $\rho(a, b)$. We see that $\sigma_X(\omega_\tau a, \omega_\tau b)$ is bounded above by a linear function of $\sigma_X(a, b)$. In particular, if $\theta_X a \sim \theta_X b$, then $\theta_X(\omega_\tau a) \sim \theta_X(\omega_\tau b)$.

Lemma 9.5. *If τ and τ' are two multicurves which together fill Σ , then the diameter of $\omega_\tau(T(\tau'))$ in $\mathcal{M}(\Sigma)$ is finite and bounded above in terms of $\xi(\Sigma)$.*

(Here, we are assuming that we have fixed, once and for all, the constant r'_0 used in defining $T(\tau)$, in terms of $\xi(\Sigma)$.)

Proof. Note that if $X \in \mathcal{X}$, then either $X \in \mathcal{X}_T(\tau)$ or $X \in \mathcal{X}_T(\tau')$. Let $a, b \in T(\tau')$. If $X \in \mathcal{X}_T(\tau)$, then $\theta_X(\omega_\tau a) \sim \theta_X \tau \sim \theta_X(\omega_\tau b)$. If $X \in \mathcal{X}_T(\tau')$ then $\theta_X a \sim \theta_X b$, and by the observation preceding the lemma, $\theta_X(\omega_\tau a) \sim \theta_X(\omega_\tau b)$. It now follows by (A7) “distance bound” that $a \sim b$ as required. \square

If τ, τ' fill Σ , we choose elements $\omega_\tau(\tau') \in \omega_\tau(T(\tau'))$ and $\omega_{\tau'}(\tau) \in \omega_{\tau'}(T(\tau))$. These are well defined up to bounded distance.

Now if $a \in T(\tau)$ and $b \in T(\tau')$, then $\omega_\tau(\tau)$ lies in a coarse interval from a to b , and $\omega_\tau(\tau')$ lies in a coarse interval from a to $\omega_\tau(\tau)$. It follows that $\rho(a, b) \asymp \rho(a, \omega_\tau(\tau')) + \rho(\omega_\tau(\tau'), \omega_\tau(\tau)) + \rho(\omega_\tau(\tau), b)$.

We can use this observation to prove the following.

Lemma 9.6. *Suppose that τ, τ' fill Σ , and that any pair of points of $T(\tau') \subseteq \mathcal{M}(\Sigma)$ lie a bounded distance from some uniform bi-infinite quasigeodesic in $T(\tau')$. Then there are constants $k, t \geq 0$, depending only on the constants of the hypotheses, such that if $x \in T(\tau')$ and $r \geq 0$, then there is some $y \in T(\tau')$ with $\rho(y, T(\tau)) \geq r$ and $\rho(x, y) \leq kr + t$.*

Proof. From the hypotheses, there is a uniformly quasigeodesic ray with basepoint $\omega_{\tau'}(\tau)$ and containing x . Now choose y to be a suitable point on this ray beyond x , and apply the above observation. \square

For the remainder of this section, we explore these statements further in the specific case where $\mathcal{M}(\Sigma) = \mathbb{M}(\Sigma)$ is the marking graph of Σ . (Note that in this case all the parameters of $\mathbb{M}(\Sigma)$ depend only on $\xi(\Sigma)$.)

Given $k \geq 0$, let $\hat{T}(\tau; k) = \{a \in \mathbb{M}^0(\Sigma) \mid \iota(a, \tau) \leq k\}$.

Lemma 9.7. (1) *For all $k \geq 0$, there is some $r \geq 0$, depending on k and $\xi(\Sigma)$, such that $\hat{T}(\tau; k) \subseteq N(T(\tau), r)$.*

(2) *There is some $k_0 \geq 0$, depending only on $\xi(\Sigma)$, such that $T(\tau) \cap \mathbb{M}^0(\Sigma) \subseteq \hat{T}(\tau; k_0)$.*

Proof. (1) It is an elementary property of subsurface projection that if $a \in \mathbb{M}^0(\Sigma)$, $\gamma \in \mathbb{G}^0(\Sigma)$ and $X \in \mathcal{X}$ with $\gamma \prec X$ or $\gamma \pitchfork X$, then $\sigma_X(\gamma, a)$ is bounded above in terms of $\iota(\gamma, a)$. It follows that $\hat{T}(\tau; k) \subseteq T(\tau; r'')$ for some r'' depending only on r and k . We now apply Lemma 9.2.

(2) We have observed that, in the case of markings, the verification of (A10) “realisation” gives us a marking which has bounded intersection with τ (where τ is the union of all boundary curves of the set of surfaces). This was used in the construction of ν_τ and hence of ω_τ . In particular, it follows that if $a \in \mathbb{M}(\Sigma)$, then $\iota(\omega_\tau a, \tau)$ is bounded for any $a \in \mathbb{M}(\Sigma)$. Now if $a \in T(\tau) \cap \mathbb{M}^0(\Sigma)$ then (since ω_τ is a coarse gate map) $\rho(a, \omega_\tau a)$ is bounded. It follows that $\iota(a, \omega_\tau a)$ is bounded, and so $\iota(a, \tau)$ is bounded. This bound depends only on $\xi(\Sigma)$. \square

Definition. A *complete multicurve* is a multicurve, τ , such that each component of $\Sigma \setminus \tau$ is an $S_{0,3}$.

In other words, $\mathcal{X}_N(\tau) = \emptyset$, so $\mathcal{X}(\tau) = \mathcal{X}_A(\tau)$. It is equivalent to saying that τ has exactly $\xi(\Sigma)$ components. (It is essentially the same thing as a “pants decomposition” in other terminology.)

Suppose that τ is a complete multicurve. In this case,

$$\mathcal{X}(\tau) = \mathcal{X}_A(\tau) = \{X(\gamma) \mid \gamma \in \tau\}.$$

If $X \in \mathcal{X}(\tau)$, then $\mathbb{G}(X) = \mathbb{M}(X)$ is quasi-isometric to the real line, and so $\mathcal{T}(\tau)$ is quasi-isometric to \mathbb{R}^ξ . Thus, ν_τ gives rise to a quasi-isometric embedding of \mathbb{R}^ξ into $\mathbb{M}(\Sigma)$, whose image is a bounded Hausdorff distance from $T(\tau)$.

We can also view this in terms of the action of $\text{Map}(\Sigma)$. Let $G(\tau) \leq \text{Map}(\Sigma)$ be the group generated by Dehn twists about the elements of τ . Thus, $G(\tau) \cong \mathbb{Z}^\xi$. We put the standard word metric on $G(\tau)$. Now $G(\tau)$ acts coboundedly on $\mathcal{T}(\tau)$ hence also on $T(\tau)$.

The following result, proven in [Farb, Lubotzky and Minsky 2001], is an immediate consequence, though it also follows directly from the distance formula [Masur and Minsky 2000] (given as Theorem 8.3 here).

Lemma 9.8. *Given any multicurve τ there is some $a \in \mathbb{M}(\Sigma)$ such that the map $a \mapsto ga : G(\tau) \rightarrow \mathbb{M}(\Sigma)$ is a uniform quasi-isometric embedding.*

In fact, we can take any $a \in T(\tau)$, and the orbit, $G(\tau)a$, is a uniformly bounded Hausdorff distance from $T(\tau)$. (The uniformity is somewhat spurious here, since there are only finitely many orbits of multicurves under the action of $\text{Map}(\Sigma)$; though our arguments give explicit bounds.)

We will refer to a set of the form $T(\tau)$ for a complete multicurve, τ , as a *coarse Dehn twist flat* (generally regarded as defined up to a uniformly bounded Hausdorff distance).

Lemma 9.9. *There are uniform constants $k, t \geq 0$ such that if τ, τ' are complete multicurves, with $\tau \neq \tau'$, $x \in T(\tau')$ and $r \geq 0$, then there is some $y \in T(\tau')$ with $\rho(y, T(\tau)) \geq r$ and $\rho(x, y) \leq kr + t$.*

Proof. If $\tau \cap \tau' = \emptyset$, then τ, τ' fill Σ , so the result follows immediately from Lemma 9.6. (Note that any path in a Dehn twist flat lies in a uniform bi-infinite quasigeodesic.)

For the general case, let $\tau_0 = \tau \cap \tau'$. Now $T(\tau_0)$ is, up to quasi-isometry, a direct product of a euclidean space (given by Dehn twists about the elements of τ_0) and copies of $\mathbb{M}(X)$ as X ranges over the elements of $\mathcal{X}_N(\tau_0)$. Applying the above intrinsically to the restrictions of τ and τ' to any such X we deduce the general case. \square

10. Quasicubes

Throughout this section, we again suppose that $\mathcal{M}(\Sigma)$ satisfies the axioms (A1)–(A10) of Section 7. We refer the constants involved as the “parameters of $\mathcal{M}(\Sigma)$ ”.

Definition. A *quasicube* in $\mathcal{M}(\Sigma)$ is an l -quasimorphism $\phi : Q \rightarrow \mathcal{M}(\Sigma)$, where Q is an n -cube.

We refer to it as an l -quasi- n -cube, if we want to specify the parameters.

In this section, we give a description of “nondegenerate” quasicubes of maximal rank. In Section 11, we will apply this to the (extended) asymptotic cone $\mathcal{M}^*(\Sigma)$.

We begin by recalling the following fact:

Lemma 10.1. *There is some $l_0 \geq 0$, depending only $\xi(\Sigma)$ such that if $X, Y \in \mathcal{X}$ and there exist $a, b, c, d \in \mathcal{M}(\Sigma)$ with $(a, b : c, d)_X \geq l_0$ and $(a, c : b, d)_Y \geq l_0$, then $X \wedge Y$.*

Here $(a, b : c, d)_X$ denotes the “crossratio”

$$(a, b : c, d)_X = \frac{1}{2}(\max\{\sigma_X(a, c) + \sigma_X(b, d), \sigma_X(a, d) + \sigma_X(b, c)\} - (\sigma_X(a, b) + \sigma_X(c, d)))$$

in $\mathcal{G}(X)$. Similarly for $(a, b : c, d)_Y$ in $\mathcal{G}(Y)$.

Proof. This is property (P4) in [Bowditch 2013], and was verified for the marking graph by Lemma 11.7 of that paper. As already observed (in Section 7 here) the proof there only made use of properties (A8) “bounded image” and (A9) “overlapping subsurfaces”. \square

Recall, from Section 7, that $\mathcal{A}_X(a, b; r) = \{Y \in \mathcal{X}(X) \mid \sigma_Y(a, b) > r\}$.

Definition. Given $a, b \in \mathcal{M}(X)$ and $r \geq 0$, we say that a, b are *weakly (X, r) -related* if for all $Y \in \mathcal{A}(a, b; r)$, we have $Y \preceq X$.

Intuitively, we can think of a, b as being “close outside X ”. More specifically, if $Z \in \mathcal{X}$, with $Z \wedge X$, then $\mathcal{A}(a, b; r) \cap \mathcal{X}(Z) = \emptyset$, so by (A7) “distance bound”, we see that $\rho_Z(a, b)$ is bounded.

Definition. We say that a, b are *(X, r) -related* if they are weakly (X, r) -related and $\rho(a, T(\partial_\Sigma X)) \leq r$ and $\rho(b, T(\partial_\Sigma X)) \leq r$.

We will often suppress mention of r where a choice (ultimately depending on the parameters of $\mathcal{M}(\Sigma)$) is clear from context, and simply refer to a, b as being “(weakly) X -related”.

Note that this property is “median convex” in the sense that if $c \in \mathcal{M}(\Sigma)$, with $\rho(\mu(a, b, c), c) \leq l$, then c is (X, r') -related to a and to b , where r' depends only on r, l and the parameters of the hypotheses.

Lemma 10.2. *If $a, b \in \mathcal{M}(\Sigma)$ are (X, r) -related, then $\rho(a, b)$ agrees with $\rho(a, b)$ up to linear bounds depending only on r and $\xi(\Sigma)$.*

Proof. The fact that $\psi_X : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(X)$ is uniformly coarsely Lipschitz immediately gives a linear upper bound for $\rho_X(a, b)$. For the other direction, set $\tau = \partial_\Sigma X$. Then, up to bounded distance, $\psi_\tau a$ and $\psi_\tau b$ differ only in the X -coordinate. Since ν_τ is uniformly coarsely Lipschitz, we have that $\rho(\omega_\tau a, \omega_\tau b) = \rho(\nu_\tau \psi_\tau a, \nu_\tau \psi_\tau b)$ is linearly bounded above by $\rho_X(\psi_X a, \psi_X b) = \rho_X(a, b)$. By assumption $\rho(a, T(\tau))$ and $\rho(b, T(\tau))$ are bounded. By Lemma 9.4, since ω_τ is a coarse gate map to $T(\tau)$, it follows that $\rho(a, \omega_\tau a)$ and $\rho(b, \omega_\tau b)$ are bounded. This bounds $\rho(a, b)$, as required. \square

Note that by median convexity, we see also that $\rho_X(x, y) \asymp \rho(x, y)$ for any $x, y \in \mathcal{M}(\Sigma)$ with $\mu(a, b, x) \sim x$ and $\mu(a, b, y) \sim y$.

Here is a criterion which implies that two elements of $\mathcal{M}(\Sigma)$ are X -related:

Lemma 10.3. *There is a constant $r_1 \geq 0$ depending only on $\xi(\Sigma)$ with the following property. Suppose $r_2 \geq r_1$, and that $a, b \in \mathcal{M}(\Sigma)$ and $X \in \mathcal{X}$. Suppose that for all $Z \in \mathcal{A}(a, b; r_2)$ we have $Z \preceq X$. Suppose moreover that whenever γ is a curve with $\gamma \pitchfork X$, there is some $Y \in \mathcal{A}_\Sigma(a, b; r_1)$ with $\gamma \pitchfork Y$ and $Y \preceq X$. Then a, b are (X, r) -related for some r depending only on r_2 and the parameters of $\mathcal{M}(\Sigma)$.*

Proof. By assumption, a, b are weakly (X, r) -related, so we just need to check that $\rho(a, T(\tau))$ and $\rho(b, T(\tau))$ are bounded, where $\tau = \partial_\Sigma(X)$.

By Lemma 9.2, it is enough to check that $\sigma_U(a, \tau)$ and $\sigma_U(b, \tau)$ are bounded for all $U \in \mathcal{X}_T(\tau)$. Now if $U \in \mathcal{X}_T(\tau)$, then U contains or crosses some boundary curve of X , and so $U \pitchfork X$ or $X \prec U$. Either way, U will contain a curve, γ , with $\gamma \pitchfork X$. By hypothesis, there is some $Y \preceq X$, with $Y \pitchfork \gamma$ and $Y \in \mathcal{A}_\Sigma(a, b; r_1)$. Note that either $Y \prec U$ or $Y \pitchfork U$.

Now $\sigma_Y(a, b) > r_1$. Also, since U is not contained in X , it does not lie in $\mathcal{A}_\Sigma(a, b; r_2)$, i.e., $\sigma_U(a, b) \leq r_2$. Suppose first that $Y \prec U$. If r_1 is greater than r_0 , the constant of (A8) “bounded image”, then it follows $\langle \theta_U a, \theta_U b; \theta_U Y \rangle \leq r_0$, and so, by the definition of Gromov product, $\sigma_U(a, Y) + \sigma_U(b, Y) \leq 2(r_2 + r_0)$. Suppose instead that $Y \pitchfork U$. If r_1 is bigger than twice the constant, r_0 , of (A9) “overlapping subsurfaces”, then without loss of generality (swapping a and b), we have $\sigma_Y(a, U) > r_0$, so by (A9) we must have $\sigma_U(a, Y) \leq r_0$. Since $\sigma_U(a, b) \leq r_2$, it follows that $\sigma_U(b, Y) \leq r_2 + r_0$. Thus, in all cases, we have shown that $\sigma_U(a, Y)$ and $\sigma_U(b, Y)$ are bounded.

But now, $Y \preceq X$, so $Y \wedge \tau$. Thus, by (A5) “disjoint projection”, we have that $\sigma_U(Y, \tau)$ is bounded. We deduce that $\sigma_U(a, \tau)$ and $\sigma_U(b, \tau)$ are bounded for all $U \in \mathcal{X}_T(\tau)$ as claimed. \square

We now move on to consider quasicubes.

Suppose that $Q = \{-1, 1\}^n$ is an n -cube. By an i -th side of Q , we mean an unordered pair, $c, d \in Q$, which differ precisely in their i -th coordinates. Note that any two i -th sides are parallel in the median sense. If $a, b \in Q$, we can speak of the sides of Q crossed by a, b , that is those which (up to parallelism) correspond to the coordinates for which a, b differ. (Note that the walls of Q are in bijective correspondence with the parallel classes of sides.)

Suppose that $\phi : Q \rightarrow \mathcal{M}(\Sigma)$ is an l -quasimorphism. If c, d and c', d' are both i -th sides of Q , then $\rho(\phi c, \phi d) \asymp \rho(\phi c', \phi d')$. (Since $\mu(\phi c, \phi d, \phi c') \sim \phi c$ and $\mu(\phi c, \phi d, \phi d') \sim \phi d$, we get a linear upper bound for $\rho(\phi c, \phi d)$, and the lower bound follows symmetrically.) We will write $s_i = \min \rho(\phi c, \phi d)$ as c, d ranges over all i -th sides. Thus $\rho(\phi c', \phi d') \asymp s_i$ for any other i -th sides, c', d' . We also note that for all $X \in \mathcal{X}$, we have $\sigma_X(\phi c, \phi d) \asymp \sigma_X(\phi c', \phi d')$ and $\rho_X(\phi c, \phi d) \asymp \rho_X(\phi c', \phi d')$ (similarly, since $\theta_X \circ \phi$ and $\psi_X \circ \phi$ are quasimorphisms to $\mathcal{G}(X)$ and $\mathcal{M}(X)$ respectively). Here, the linear bounds depend implicitly on l .

If $a, b \in Q$, then a repeated application of Lemma 6.1 shows $\rho(\phi a, \phi b) \asymp \sum_i s_i$, where the sum is taken over all sides of Q crossed by a, b .

Lemma 10.4. *Let $\phi : Q \rightarrow \mathcal{M}$ be an l -quasicube. There is some $k_0 \geq 0$, depending only on h and the parameters of the hypotheses, such that if $X, Y \in \mathcal{X}$ with*

$\sigma_X(\phi c, \phi d) \geq k_0$ and $\sigma_Y(\phi c', \phi d') \geq k_0$, where c, d and c', d' are respectively i -th and j -th sides of Q , then either $i = j$ or $X \wedge Y$.

Proof. This is an immediate consequence of Lemma 10.1 above. \square

It follows from Lemma 10.4 that for each i , there is a (possibly empty, possibly disconnected) subsurface, Y_i , of Σ , which contains all $X \in \mathcal{X}$, for which $\sigma_X(\phi c, \phi d) \geq k_0$ for any i -th side, c, d , of Q . (Here, we are using the term “subsurface” to mean a disjoint union of essential subsurfaces as defined in Section 7.) We can also take Y_i to be minimal with this property. To ensure that each Y_i must be nonempty, we will assume that ϕ is *nondegenerate*, that is, for each side, c, d , $\mathcal{A}_\Sigma(c, d; r) \neq \emptyset$, for a fixed $r \geq k_0$. Note that, by (A7) “distance bound”, this is implied by placing a suitable lower bound on $\min\{s_i \mid 1 \leq i \leq n\}$. We will also take r to be at least the constant r_1 featuring in Lemma 10.3.

Recall that $\xi(\Sigma)$ is the maximal number of disjoint and distinct (non- $S_{0,3}$) subsurfaces we can embed in Σ . Thus, if $n = \xi(\Sigma)$, we see that if ϕ is nondegenerate, then each Y_i is connected and is either an annulus, or has complexity-1 (that is a $S_{0,4}$ or $S_{1,1}$).

Definition. We say that a multicurve τ is *big* if each component of $\Sigma \setminus \tau$ is a $S_{0,3}$, $S_{0,4}$ or $S_{1,1}$.

In this case, the set of all relative boundary components of all the Y_i is a big multicurve, τ , such that $\mathcal{X}(\tau)$ is precisely the set of Y_i .

Lemma 10.5. *Suppose that Q is a ξ -cube, and that $\phi : Q \rightarrow \mathcal{M}(\Sigma)$ is a nondegenerate quasimorphism. Then there is a big multicurve, $\tau \subseteq \Sigma$, such that we can write $\mathcal{X}(\tau) = \{Y_1, \dots, Y_\xi\}$, so that if c, d is any i -th side of Q , then $\phi c, \phi d$ are Y_i -related.*

(We remark that it follows that $\phi(Q)$ lies in a bounded neighbourhood of $T(\tau)$.)

Proof. We construct disjoint surfaces Y_i as above, and as already observed, the set of these Y_i is precisely $\mathcal{X}(\tau)$ for a big multicurve, τ . Recall that for all $X \in \mathcal{X}$, if c, d and c', d' are i -th sides of Q , the $\sigma_X(\phi c, \phi d) \asymp \sigma_X(\phi c', \phi d')$. Let r_1 be the constant of Lemma 10.3. If the nondegeneracy constant is sufficiently large, then $\mathcal{A}(\phi c, \phi d; r) \neq \emptyset$. So if $r \geq \max\{r_1, k_0\}$, the subsurfaces of $\mathcal{A}(\phi c, \phi d; r)$ fill Y_i , so $\phi c, \phi d, Y_i$ satisfy the hypotheses of Lemma 10.3. We see that $\phi c, \phi d$ are Y_i -related as claimed \square

Note that a big multicurve τ satisfying the conclusion of Lemma 10.5 might not be unique. For example, if $\gamma \in \tau$ bounds an $S_{0,3}$ component of $\Sigma \setminus \tau$ on both sides (perhaps the same $S_{0,3}$), then we can remove it, and the conclusion will still hold. However, this is essentially the only ambiguity that can arise.

We will also need:

Lemma 10.6. *Suppose that Q is a ξ -cube, and that $\phi : Q \rightarrow \mathcal{M}(\Sigma)$ is a non-degenerate quasimorphism, and that c, d is an i -th side of Q . Let Y_i be as given by Lemma 10.5. Suppose that $x, y \in \mathcal{M}(\Sigma)$ with $\mu(\phi c, \phi d, x) \sim x$ and $\mu(\phi c, \phi d, y) \sim y$. Then $\rho(x, y) \asymp \rho_{Y_i}(x, y)$, where the additive bounds depend only on l, n and the parameters of $\mathcal{M}(\Sigma)$.*

Proof. By Lemma 10.4, and $\phi c, \phi d$ are Y_i -related, and so therefore are x, y (with suitable constants). The statement then follows by Lemma 10.2 and the subsequent observation. \square

11. The asymptotic cone of $\mathcal{M}(\Sigma)$

As in Section 5, let $\mathcal{M}^*(\Sigma)$ be the extended asymptotic cone of a space $\mathcal{M}(\Sigma)$ satisfying the hypotheses (A1)–(A10).

Let \mathcal{Z} be a countable set with a nonprincipal ultrafilter, as in Section 5. Let $\mathcal{U}\mathbb{G} = \mathcal{U}\mathbb{G}(\Sigma)$ be the ultrapower of $\mathbb{G}(\Sigma)$. This is a graph with vertex set $\mathcal{U}\mathbb{G}^0$. Note that the intersection number, ι , extends to a map $\mathcal{U}\iota : (\mathcal{U}\mathbb{G}^0)^2 \rightarrow \mathcal{U}\mathbb{N}$. We also have an ultrapower, $\mathcal{U}\mathcal{X} = \mathcal{U}\mathcal{X}_A \sqcup \mathcal{U}\mathcal{X}_N$. There is a natural bijection between $\mathcal{U}\mathcal{X}_A$ and $\mathcal{U}\mathbb{G}^0$. (Here, \mathbb{G}^0 is playing its role as an indexing set.)

We can extend the notation introduced in Section 5. For example, if $X, Y \in \mathcal{U}\mathcal{X}$, we write $X \wedge Y$ to mean that $X_\zeta \wedge Y_\zeta$ almost always. We similarly define $X \prec Y$ and $X \pitchfork Y$. Since there are only finitely many possibilities (in fact, five), we have the following pentachotomy: if $X, Y \in \mathcal{X}$, exactly one of $X = Y$, $X \wedge Y$, $X \prec Y$, $Y \prec X$ or $X \pitchfork Y$ must hold (exactly as in Section 7).

Note that $\mathcal{U}\text{Map}(\Sigma)$ acts on both $\mathcal{U}\mathbb{G}$ and $\mathcal{U}\mathcal{X}$ with finite quotient.

Terminology. In this section, we refer to an element of $\mathcal{U}\mathbb{G}^0$ as a *curve* and an element of $\mathbb{G}^0 \subseteq \mathcal{U}\mathbb{G}^0$ as *standard curve*. We similarly refer to “subsurfaces”, “standard subsurfaces” etc.

As observed in Section 5, two standard objects lie in the same $\mathcal{U}\text{Map}(\Sigma)$ -orbit, then they lie in the same $\text{Map}(\Sigma)$ -orbit.

Moreover, any configuration of curves and surfaces of bounded complexity can be assumed standard up to the action of the mapping class group. One way to express this is as follows.

Lemma 11.1. *Suppose that $a \subseteq \mathcal{U}\mathbb{G}^0$ and $\mathcal{U}\iota(a) \in \mathbb{N}$, then there is some $g \in \mathcal{U}\text{Map}(\Sigma)$ with $ga \subseteq \mathbb{G}^0$.*

Proof. By hypothesis, $\iota(a_\zeta)$ is almost always constant. Therefore, we can find $g_\zeta \in \text{Map}(\Sigma)$ such that $g_\zeta a_\zeta \subseteq \mathbb{G}^0(\Sigma)$ lies in one of only finitely many possible subsets of $\mathbb{G}^0(\Sigma)$. Therefore, $g_\zeta a_\zeta$ is almost always constant, that is, ga is standard, where g is the limit of $(g_\zeta)_\zeta$. \square

Note that this applies, for example, to multicurves, or to collections of pairwise disjoint subsurfaces of Σ . In particular, it makes sense to refer to the topological type of a subsurface; for example, that it is an $S_{1,1}$ or an $S_{0,4}$ (up to the action of $\mathcal{U}\text{Map}(\Sigma)$). We can also refer to boundary curves of a surface, or say that a collection of curves fill a subsurface, etc.

If τ is a multicurve, we can define $\mathcal{UX}(\tau) \subseteq \mathcal{UX}$ as in Section 9. (It is the limit of the sets $\mathcal{X}(\tau_\zeta)$.)

In what follows we deal mostly with extended asymptotic cones. This seems more natural in this context than restricting to the asymptotic cone, though most of the discussion would apply equally well in both situations.

We assume that $\mathcal{M}(\Sigma)$ and $\mathcal{G}(\Sigma)$ are spaces satisfying properties (A1)–(A10) of Section 7.

Let $t \in \mathcal{UR}$ be a positive infinitesimal. Rescaling as in Section 5 we get extended asymptotic cones, $\mathcal{M}^* = \mathcal{M}^*(\Sigma)$ and $\mathcal{G}^* = \mathcal{G}^*(\Sigma)$ of $\mathcal{M}(\Sigma)$ and $\mathcal{G}(\Sigma)$, respectively. (We topologise them as the disjoint union of their components.) We write ρ^* , σ^* , respectively, for the limiting nonstandard metrics. The asymptotic cones, $(\mathcal{M}^\infty(\Sigma), \rho^\infty)$ and $(\mathcal{G}^\infty(\Sigma), \sigma^\infty)$ are complete metric spaces. In fact, since $\mathcal{G}(\Sigma)$ is hyperbolic, $\mathcal{G}^*(\Sigma)$ is an \mathbb{R}^* -tree, and $\mathcal{G}^\infty(\Sigma)$ is an \mathbb{R} -tree. The map $\chi : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$ is coarsely Lipschitz, and (after the rescaling) gives rise to a Lipschitz map $\chi^* : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$.

The coarse median μ on $\mathcal{M}(\Sigma)$ gives rise to a median μ^* on \mathcal{M}^* , and restricts to a median, μ^∞ on the asymptotic cone, \mathcal{M}^∞ . By Theorems 6.8 and 6.9, ρ^∞ is bi-Lipschitz equivalent to a median metric inducing the same median structure. The construction is not canonical, but we will write ρ_M^∞ for some choice of such median metric. We similarly define ρ_M^* bi-Lipschitz equivalent to ρ^* on \mathcal{M}^* .

Suppose that $X \in \mathcal{UX}$. The spaces $\mathcal{G}(X_\zeta)$ give rise to an extended asymptotic cone, denoted $\mathcal{G}^*(X)$ which is an \mathbb{R}^* -tree. The maps θ_{X_ζ} are uniformly coarsely Lipschitz, and give rise to a Lipschitz homomorphism, $\theta_X^* : \mathcal{M}^*(X) \rightarrow \mathcal{G}^*(X)$.

Similarly, we have a limit $\mathcal{M}^*(X)$ of the spaces $\mathcal{M}(X_\zeta)$. This has a median μ^* arising from the coarse medians μ_ζ and we again get a topological median algebra. We similarly have limiting Lipschitz homomorphisms $\psi_X^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{M}^*(X)$. In fact, as observed above, up to the action of $\mathcal{U}\text{Map}(\Sigma)$, we could take X to be standard, and so $\mathcal{M}^*(X)$ is isomorphic to the space defined intrinsically on a surface of this topological type.

If $X, Y \in \mathcal{UX}$, with $X \preceq Y$, then we have a limiting map, $\psi_{YX}^* : \mathcal{M}^*(X) \rightarrow \mathcal{M}^*(Y)$, with $\psi_{YX}^* \circ \psi_X^* = \psi_Y^*$. We will generally abbreviate ψ_{XY}^* to ψ_X^* , when the domain is clear from context.

Note that if $\gamma \in \mathcal{UG}^0$ and $X \in \mathcal{UX}$ with $\gamma \preceq X$ or $\gamma \pitchfork X$, we have a well defined subsurface projection, $\theta_X^*(\gamma) \in \mathcal{G}^*(X)$. Similarly, if $X, Y \in \mathcal{UX}$, with $Y \prec X$ or $Y \pitchfork X$, we can define $\theta_X^*(Y) \in \mathcal{G}^*(X)$.

We also note that if $\gamma \in \mathcal{UG}^0$, then we can write $\mathcal{M}^*(\gamma) = \mathcal{G}^*(\gamma) = \mathcal{G}^*(X)$ where $X = X(\gamma)$ is an annular neighbourhood of γ .

Suppose that $\tau \subseteq \mathcal{UG}^0$ is a multicurve. Let $T^*(\tau) \subseteq \mathcal{M}^*(\Sigma)$ be the limit of the subsets $T(\tau_\zeta) \subseteq \mathcal{M}(\Sigma)$. This is a closed subset of $\mathcal{M}^*(\Sigma)$. (Note that it is also the limit of the sets $T(\tau_\zeta; r)$ for any sufficiently large $r \in [0, \infty)$.)

We can describe the structure of $T^*(\tau)$ as follows.

Let $\mathcal{UX}(\tau) \subseteq \mathcal{UX}$ be the ultraproduct of the $\mathcal{X}(\tau_\zeta)$. (By Lemma 11.1, this is finite, and standard up to the action of $\mathcal{U} \text{Map}(\Sigma)$.) Let $\mathcal{T}^*(\tau)$ be the direct product of the spaces $\mathcal{M}^*(X)$ for $X \in \mathcal{UX}(\tau)$ in the l^1 extended metric. This is the same as the extended asymptotic cone of the spaces $\mathcal{T}(\tau_\zeta)$.

Recall that in Section 9, we defined maps $\psi_{\tau_\zeta} : \mathcal{M}(\Sigma) \rightarrow \mathcal{T}(\tau_\zeta)$, $\nu_{\tau_\zeta} : \mathcal{T}(\tau_\zeta) \rightarrow T(\tau_\zeta)$ and $\omega_{\tau_\zeta} = \nu_{\tau_\zeta} \circ \psi_{\tau_\zeta} : \mathcal{M}(\Sigma) \rightarrow \mathcal{T}(\tau_\zeta)$. These are all uniformly coarsely Lipschitz quasimorphisms, and so give rise to maps, $\psi_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{T}^*(\tau)$, $\nu_\tau^* : \mathcal{T}^*(\tau) \rightarrow T^*(\tau)$ and $\omega_\tau^* = \nu_\tau^* \circ \psi_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow T^*(\tau)$. In fact, from Lemma 9.4, we see that $\omega_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow T^*(\tau)$ is a gate map.

It follows that $T^*(\tau)$ is convex, and that ω_τ^* is the unique gate map to $T^*(\tau)$. Note also that if $\gamma \in \mathcal{UG}^0(\Sigma)$ with $\tau \pitchfork \gamma$ and $a \in T^*(\tau)$, then $\theta_\gamma^*(a) = \theta_\gamma^*(\tau)$.

Given a subset, $S \subseteq \mathcal{M}^*(\Sigma)$, write

$$C(S) = \{X \in \mathcal{UX} \mid \theta_X^*|_S \text{ is injective}\},$$

and

$$D(S) = \{X \in \mathcal{UX} \mid \psi_X^*|_S \text{ is injective}\}.$$

Clearly $C(S) \subseteq D(S)$. We write $D^0(S) = C(S) \cap \mathcal{UX}_A = D(S) \cap \mathcal{UX}_A$, which we can identify as a subset of \mathcal{UG}^0 .

Given $a, b \in \mathcal{M}^*(\Sigma)$, write $C(a, b) = C([a, b])$, $D(a, b) = D([a, b])$ and $D^0(a, b) = D^0([a, b])$.

We also write $A(a, b) = C(\{a, b\}) = \{X \in \mathcal{UX} \mid \theta_X^*a \neq \theta_X^*b\}$. Clearly $C(a, b) \subseteq A(a, b)$.

Lemma 11.2. *Suppose that $a, b, c, d \in \mathcal{M}^*(\Sigma)$ are all distinct. Suppose $X, Y \in \mathcal{UX}$ $c \in [a, d]$, $b \in [a, c]$, $X \in A(a, b) \cap A(c, d)$ and $Y \in A(b, c)$, then either $X = Y$ or $X \wedge Y$.*

Proof. Suppose first, for contradiction, that $X \pitchfork Y$. (This is the case that is actually of interest to us.) Let $a_\zeta, b_\zeta, c_\zeta, d_\zeta \in \mathcal{M}(X_\zeta)$ be sequences converging to $a, b, c, d \in \mathcal{M}^*(\Sigma)$. We can suppose that $\rho_{X_\zeta}(c_\zeta, \mu(a_\zeta, c_\zeta, d_\zeta))$ and $\rho_{X_\zeta}(b_\zeta, \mu(a_\zeta, b_\zeta, c_\zeta))$ are bounded (after replacing c_ζ by $\mu(a_\zeta, c_\zeta, d_\zeta)$ and then b_ζ by $\mu(a_\zeta, b_\zeta, c_\zeta)$). Now $\sigma_{X_\zeta}(a_\zeta, b_\zeta) \rightarrow \infty$ (since $\theta_{X_\zeta}a_\zeta \rightarrow \theta_X^*a$ and $\theta_{X_\zeta}b_\zeta \rightarrow \theta_X^*b$, which by hypothesis are distinct). In particular, $\sigma_{X_\zeta}(a_\zeta, b_\zeta)$ is almost always greater than $2r_0$, where r_0 is the constant of property (A9) ‘‘overlapping subsurfaces’’. Thus, $\theta_{X_\zeta}Y_\zeta$ must be at distance greater than r_0 from either $\theta_{X_\zeta}a_\zeta$ or $\theta_{X_\zeta}b_\zeta$, and so by (A9), $\theta_{Y_\zeta}X_\zeta$ is within

a distance r_0 from either $\theta_{Y_\zeta} a_\zeta$ or $\theta_{Y_\zeta} b_\zeta$. Similarly, $\theta_{Y_\zeta} X_\zeta$ is also almost always within distance r_0 of either $\theta_{Y_\zeta} c_\zeta$ or $\theta_{Y_\zeta} d_\zeta$. But $\mathcal{G}(Y_\zeta)$ is uniformly hyperbolic, and θ_{X_ζ} is a median quasimorphism. Therefore, up to bounded distance, $\theta_{Y_\zeta} b_\zeta$ and $\theta_{Y_\zeta} c_\zeta$ lie on a geodesic from $\theta_{Y_\zeta} a_\zeta$ to $\theta_{Y_\zeta} d_\zeta$, and occur in this order. Therefore, whichever of the above possibilities arises, we see that $\sigma_{Y_\zeta}(b_\zeta, c_\zeta)$ is bounded, and so $\sigma_Y^*(b, c) = 0$. That is, $\theta_X^* b = \theta_X^* c$, so $Y \notin A(b, c)$.

After swapping the roles of X and Y if necessary, we also need to rule out the possibility that $X < Y$. If that were the case, we could derive a similar contradiction using property (A8) ‘‘bounded image’’. Briefly, if $\sigma_{X_\zeta}(a_\zeta, b_\zeta)$ is large then $\theta_{Y_\zeta} X_\zeta$ must lie close to any geodesic in $\mathcal{G}(Y_\zeta)$ from $\theta_{Y_\zeta} a_\zeta$ to $\theta_{Y_\zeta} b_\zeta$. Similarly, $\theta_{Y_\zeta} X_\zeta$ lies close to any geodesic from $\theta_{Y_\zeta} c_\zeta$ to $\theta_{Y_\zeta} d_\zeta$. We again get that $\sigma_{Y_\zeta}(b_\zeta, c_\zeta)$ is bounded, and so derive a contradiction. (We omit details, since we will not need this case in this paper.) \square

We say that a subset O of $\mathcal{M}^*(\Sigma)$ is *monotone* if it admits a total order $<$ such that if $x < y < z$ in O then $y \in [x, z]$.

Recall that, $C(O)$ is the set of $X \in \mathcal{UX}$ such that $\theta_X^*|_O : O \rightarrow \mathcal{G}^*(X)$ is injective. The following is an immediate corollary of Lemma 11.2.

Corollary 11.3. *If $O \subseteq \mathcal{M}^*(\Sigma)$ is monotone, $|O| \geq 4$, and $X, Y \in C(O)$, then either $X = Y$, or $X \wedge Y$.*

(In fact, the only information we really need from Lemma 11.2 and Corollary 11.3 is that X and Y cannot cross.)

Note, in particular, this applies if $O \subseteq \mathcal{M}^*(\Sigma)$ is a convex subset order-isomorphic to a totally ordered set.

In particular, Corollary 11.3 tells us that, if $|O| \geq 4$, then $C^0(O)$ is a multicurve.

Note, in particular, this applies if $O \subseteq \mathcal{M}^*(\Sigma)$ is a nontrivial convex subset order-isomorphic to a totally ordered set.

12. Cubes in $\mathcal{M}^*(\Sigma)$

Let $Q \subseteq \mathcal{M}^*(\Sigma)$ be an n -cube. If c, d and c', d' are both i -th sides of Q , 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 $x, y \in [c, d]$, let $x' = \mu(c', d', x)$ and $y' = \mu(c', d', y)$. Given $X \in \mathcal{UX}$, and $\theta_X^* x = \theta_X^* y$, then $\theta_X^* x' = \theta_X^* y'$. We see that if $\theta_X^*|_{[c', d']}$ is injective, then so is $\theta_X^*|_{[c, d]}$ and conversely by symmetry. The same applies to ψ_X^* . Thus $C([c, d]) = C([c', d'])$, so we can write this as $C_i(Q)$. We similarly write $D_i(Q) = D([c, d]) = D([c', d'])$. We write $D_i^0(Q) = D_i(Q) \cap \mathcal{UG}^0(\Sigma) = C_i(Q) \cap \mathcal{UG}^0(\Sigma)$, which we identify with the set of curves $\gamma \in \mathcal{UG}^0(\Sigma)$ such that $\theta_\gamma^*|_{[c, d]}$ is injective.

Suppose now that $n = \xi$. In this case, if c, d is any side of Q , then $[c, d]$ is a rank-1 median algebra (a totally ordered set). We refer to $[c, d]$ as a *face* of the

convex hull, $\text{hull}(Q)$, of Q . In fact, $\text{hull}(Q)$ is a median direct product of its faces. If $n = \xi$, then each of the faces has rank 1. Such a face is linearly ordered, and so it is isometric in ρ_M^* to an interval in \mathbb{R}^* (via the map $x \mapsto \rho^*(c, x)$).

Applying Corollary 11.3, we immediately get:

Lemma 12.1. *If $X, Y \in C_i(Q)$ for any i , then either $X = Y$ or $X \wedge Y$.*

Recall that by Lemma 6.3, there are uniform quasimorphisms, $\phi_\zeta : Q \rightarrow \mathcal{M}(\Sigma)$, such that $\phi_\zeta x \rightarrow x$ for all $x \in Q$. Note that, necessarily, we have that ϕ_ζ is nondegenerate for almost all ζ .

Lemma 12.2. *If $X \in C_i(Q)$ and $Y \in C_j(Q)$, then either $i = j$ or $X \wedge Y$.*

Proof. Let $\phi_\zeta : Q \rightarrow \mathcal{M}(\Sigma)$ be \mathcal{Z} -sequence of uniform quasimorphisms as given by Lemma 6.7, with $\phi_\zeta x \rightarrow x$ for all $x \in Q$. Let c, d and c', d' be i -th and j -th sides of Q , respectively. Then $\sigma_{X_\zeta}(\phi_\zeta c, \phi_\zeta d) \rightarrow \infty$ and $\sigma_{Y_\zeta}(\phi_\zeta c', \phi_\zeta d') \rightarrow \infty$, and so for almost all ζ , $\sigma_{X_\zeta}(\phi_\zeta c, \phi_\zeta d) \geq k_0$ and $\sigma_{Y_\zeta}(\phi_\zeta c', \phi_\zeta d') \geq k_0$, where k_0 is the constant of Lemma 10.4. If $i \neq j$, then by Lemma 10.4, $X_\zeta \wedge Y_\zeta$, so it follows that $X \wedge Y$. \square

Given that $D_i^0(Q) \subseteq C_i(Q)$, we see that if $\alpha \in D_i^0(Q)$ and $\beta \in D_j^0(Q)$, then either $\alpha = \beta$ or $\alpha \wedge \beta$, and in the former case, $i = j$.

Lemma 12.3. *Suppose that $\gamma \in D_i^0(Q)$ and that $X \in D_j(Q)$ is a complexity-1 subsurface (i.e., an $S_{0,4}$ or $S_{1,1}$). If $\gamma \prec X$, then $i = j$.*

Proof. Let k_0 be the constant of Lemma 10.4. Let ϕ_ζ, c, d, c', d' be as in the proof of Lemma 12.2. Now $\sigma_{Y_\zeta}(\phi_\zeta c, \phi_\zeta d) \rightarrow \infty$, and $\rho_{X_\zeta}(\phi_\zeta c', \phi_\zeta d') \rightarrow \infty$. Thus we have the following for almost all ζ . First, $\sigma_{Y_\zeta}(\phi_\zeta c, \phi_\zeta d) \geq k_0$. Second, using (A7) “distance bound”, there is some $Y_\zeta \preceq X_\zeta$ with $\sigma_{Y_\zeta}(\phi_\zeta c, \phi_\zeta d) \geq k_0$. Third, $\gamma_\zeta \prec X_\zeta$ and X_ζ has complexity 1. It follows that either $\gamma_\zeta \preceq Y_\zeta$ or $\gamma_\zeta \pitchfork Y_\zeta$. But then, by Lemma 10.4, we must have $i = j$. \square

Lemma 12.4. *Let $Q \subseteq \mathcal{M}^*(\Sigma)$ be a ξ -cube. Then there is a big multicurve τ such that we can write $\mathcal{UX}(\tau) = \{Y_1, \dots, Y_\xi\}$ with the Y_i all distinct, and with $Y_i \in D_i(Q)$.*

Proof. Let $\phi_\zeta : Q \rightarrow \mathcal{M}^*(\Sigma)$ be uniform quasimorphisms as in the two previous proofs. Let τ_ζ be the standard big multicurve given by Lemma 10.5, and write $\mathcal{X}(\tau_\zeta) = \{Y_{1,\zeta}, \dots, Y_{\xi,\zeta}\}$. Let c, d be an i -th side of Q . Then $\phi_\zeta c$ and $\phi_\zeta d$ are $Y_{i,\zeta}$ -related. Write $c_\zeta = \phi_\zeta c$ and $d_\zeta = \phi_\zeta d$. Let τ be the limit of $(\tau_\zeta)_\zeta$, and let Y_i be the limit of $(Y_{i,\zeta})_\zeta$. Thus $\mathcal{UX}(\tau) = \{Y_1, \dots, Y_\xi\}$.

It remains to show that $Y_i \in D_i(Q)$. Suppose that $x, y \in [c, d]$. Let $x_\zeta, y_\zeta \in \mathcal{M}(\Sigma)$ with $x_\zeta \rightarrow x$ and $y_\zeta \rightarrow y$. After replacing x_ζ by $\mu(c_\zeta, d_\zeta, x_\zeta)$ and y_ζ by $\mu(c_\zeta, d_\zeta, y_\zeta)$, we can assume that $\mu(c_\zeta, d_\zeta, x_\zeta) \sim x_\zeta$ and $\mu(c_\zeta, d_\zeta, y_\zeta) \sim y_\zeta$. By Lemmas 10.5, $\phi_\zeta c, \phi_\zeta d$ are X -related. By Lemma 10.2 and subsequent remarks, it follows that $\rho(x_\zeta, y_\zeta) \asymp \rho_{Y_{i,\zeta}}(x_\zeta, y_\zeta)$. But $\rho(x_\zeta, y_\zeta) \rightarrow \rho(x, y)$ and $\rho_{Y_{i,\zeta}}(x_\zeta, y_\zeta) \rightarrow$

$\rho_{Y_i}(x, y)$, and so $\rho(x, y)$ and $\rho_{Y_i}(x, y)$ are bi-Lipschitz related. In particular, if $\psi_{Y_i}^*x = \psi_{Y_i}^*y$, then $\rho_{Y_i}(x, y) = 0$, so $\rho(x, y) = 0$, so $x = y$. In other words, this shows that $\psi_{Y_i}^*[[c, d]]$ is injective, so $Y_i \in D_i(Q)$ as claimed. \square

Let I be the set of i such that $Y_i = X(\gamma_i)$ is an annulus. Thus, $\gamma_i \in D_i^0(Q)$, and $\tau = \{\gamma_i \mid i \in I\}$. By Lemmas 12.2 and 12.4, we see that, in fact, $D_i^0(Q) = \{\gamma_i\}$. Moreover, if $D_j^0(Q) \neq \emptyset$ for some $j \notin I$, then using Corollary 11.3 and Lemmas 12.2 and 12.3, we again see that $D_j^0(Q)$ consists of a single curve γ_j , with $\gamma_j \prec Y_j$. Now write $I(Q) = \{i \mid D_i^0(Q) \neq \emptyset\}$, and let $\tau(Q) = \{\gamma_i \mid i \in I(Q)\}$. We see that $\tau(Q)$ is a big multicurve containing τ and that it also satisfies the conclusion of Lemma 12.4 since $\mathcal{UX}_N(\tau(Q)) \subseteq \mathcal{UX}_N(\tau)$. (Therefore retrospectively, we could have taken $\tau = \tau(Q)$ in Lemma 12.4.)

We have shown that each of the maps θ_{γ_i} or ψ_{Y_i} restricted to the i -th face of $\text{hull}(Q)$ is injective. It follows that the map $\psi_\tau^* : \text{hull}(Q) \rightarrow T^*(\tau)$, and hence $\omega_\tau^* : \text{hull}(Q) \rightarrow T^*(\tau)$, is injective.

Up until now, we have just assumed (A1)–(A10). To proceed, we will need also to assume the distance formula (B1) given in Section 7. As a consequence, we can weaken the hypotheses of Lemma 10.2 as follows:

Lemma 12.5. *If $a, b \in \mathcal{M}(\Sigma)$ are weakly (X, r) -related, then $\rho(a, b)$ agrees with $\rho_X(a, b)$ up to linear bounds depending only on r and on $\xi(\Sigma)$.*

Proof. The only contribution to the distance formula (property (B1)) in $\mathcal{M}(\Sigma)$ comes from subsurfaces of X , and so gives the same answer in $\mathcal{M}(X)$ up to linear bounds. \square

Lemma 12.6. *Let $Q \subseteq \mathcal{M}^*(\Sigma)$ be a ξ -cube. Then $Q \subseteq T^*(\tau(Q))$.*

Proof. Suppose, for contradiction, that $a \in Q \setminus T^*(\tau)$. We can write $\text{hull}(Q)$ as a direct product of the intervals $[a, a_i]$, where $\{a, a_i\}$ is the i -th side of Q containing a . By Lemma 2.1, $\mathcal{M}^*(\Sigma)$ is locally convex (since every component is), so there is a convex neighbourhood C of a with $C \cap T^*(\tau) = \emptyset$. Since $\mathcal{M}^*(\Sigma)$ has no isolated points, we can find some $b_i \in ([a, a_i] \cap C) \setminus \{a\}$. Now, since ω_τ^* is injective, $\omega_\tau^*b_i \neq \omega_\tau^*a$. Let W_i be a wall of $\mathcal{M}^*(\Sigma)$ separating $\omega_\tau^*b_i$ from ω_τ^*a . Since ω_τ^* is a gate map to $T^*(\tau)$, this also separates b_i from a . Also, since C and $T^*(\tau)$ are convex, there is another wall W separating C from $T^*(\tau)$. (Any two disjoint convex subsets of a median algebra are separated by a wall.) We see that the walls, W, W_1, \dots, W_ξ , all pairwise cross. We deduce that $\mathcal{M}^*(\Sigma)$ has rank at least $\xi + 1$, contradicting Lemma 6.6. \square

Lemma 12.7. *Suppose that $\gamma \in D_i^0(Q)$ and that $X \in D_j(Q)$ is a complexity-1 subsurface. If $i \neq j$, then either $\gamma \wedge X$, or there is some (unique) $\beta \in D_j^0(Q)$ with $\beta \prec X$.*

Proof. If γ is not disjoint from X , then we must have $\gamma \prec X$ or $\gamma \pitchfork X$, but the first possibility is ruled out by Lemma 12.3. Since X has complexity 1, either $\gamma \pitchfork Y$ for all $Y \preceq X$, or else there is a unique $\beta \prec X$ such that $\gamma \pitchfork Y$ whenever $Y \prec X$ and $Y \neq X(\beta)$. (In the former case, $\gamma \cap X$ cuts X into a collection of discs, and in the later it cuts X into discs together with one annulus, and we take β to be its core curve.) We claim that the latter case holds, and that $\beta \in D_j(Q)$.

Let ϕ_ζ, c, d, c', d' be as in the proof of Lemma 12.2. In the first case above, we follow the argument of Lemma 12.3 to derive a contradiction. In the second case, let $\beta_\zeta \rightarrow \beta$. For almost all ζ , we have $\beta_\zeta \prec X_\zeta$ and $\gamma_\zeta \pitchfork X_\zeta$. Given such ζ , suppose that $Y_\zeta \preceq X_\zeta$ and $Y \neq X(\beta_\zeta)$. Then almost always, $\sigma_\zeta(\phi_\zeta c', \phi_\zeta d')$ is bounded. (Otherwise, since $\gamma_\zeta \pitchfork Y_\zeta$ we derive a contradiction, as in the proof of Lemma 12.2.) Given any $e, f \in [\phi c, \phi d]$, we see that $\sigma_{Y_\zeta}(e_\zeta, f_\zeta)$ is (almost always) bounded. Thus, intrinsically to $\mathcal{M}(X_\zeta)$, we see that $\psi_{X_\zeta} e_\zeta$ and $\psi_{X_\zeta} f_\zeta$ are weakly β_ζ -related. It follows that $\sigma_{\beta_\zeta}(e_\zeta, f_\zeta) \asymp \rho_{X_\zeta}(e_\zeta, f_\zeta)$. Thus, if $e \neq f$, then since $X \in D_j(Q)$, we have $\rho_X(e, f) \neq 0$, so $\sigma_\beta(e, f) \neq 0$. It follows that $\theta_\beta^*[[\phi c', \phi d']]$ is injective. In other words, $\beta \in D_j^0(Q)$ as claimed. \square

We can summarise what we have shown as follows. Recall that $D_i(Q)$ is the set of subsurfaces, X , for which $\psi_X^*|[c, d]$ is injective for some (or equivalently any) i -th side, $[c, d]$, of Q .

Proposition 12.8. *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 $Y_i \in D_i(Q)$. The set of 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 $\tau(Q)$.*

We note, in particular, that γ_i or Y_i is completely determined intrinsically by any i -th side of Q , without reference to Q itself.

13. Flats in $\mathbb{M}^*(\Sigma)$

In this section, we restrict to the case where $\mathcal{M}(\Sigma) = \mathbb{M}(\Sigma)$ is the marking graph, and consider flats in the (extended) asymptotic cone. The parameters now depend only on $\xi(\Sigma)$.

First, we consider a particular case arising from complete multicurves. Suppose that $\tau \subseteq \mathcal{UG}^0$ is a (nonstandard) complete multicurve. In other words, τ has ξ components and cuts Σ into $S_{0,3}$'s. In this case, each factor is a copy of \mathbb{R}^* , so $T^*(\tau)$ is isomorphic to $(\mathbb{R}^*)^\xi$. We refer to $T^*(\tau)$ as an *extended Dehn twist flat*.

More generally, if τ is big (that is each component of the complement is an $S_{0,3}$, $S_{0,4}$, or $S_{1,1}$), then again $\mathcal{UX}(\tau)$ has ξ elements, and $T^*(\tau)$ is a direct product of ξ \mathbb{R}^* -trees.

If $X \in \mathcal{UX}$, then $\mathbb{M}^*(X)$ and $\mathbb{G}^*(X)$ have preferred basepoints. These are defined as follows. Fix any $a \in \mathbb{M}(\Sigma)$ and let $e_X \in \mathbb{M}^*(X)$ be the limit of the points

$\psi_{X_\zeta}(a) \in \mathbb{M}^*(X_\zeta)$. This limit is independent of a . We similarly define $f_X \in \mathbb{G}^*(\Sigma)$ as the limit of $\theta_{X_\zeta}(a)$ (or equivalently, as $f_X = \chi^* e_X$). Let $\mathbb{M}^\infty(X)$ and $\mathbb{G}^\infty(X)$ be the components containing e_X and f_X respectively. Using Lemma 11.1, one sees that these are isomorphic to the asymptotic cones defined intrinsically on a standard surface of the topological type of X (unless X is an annulus, in which case, they are both isometric copies of \mathbb{R}). Note that $\theta_X^*(\mathbb{M}^\infty(\Sigma)) \subseteq \mathbb{G}^\infty(\Sigma)$. We will denote the restriction of θ_X^* to $\mathbb{M}^\infty(\Sigma)$ by θ_X^∞ .

If τ is a complete multicurve, we write $T^\infty(\tau) = T^*(\tau) \cap \mathbb{M}^\infty(\Sigma)$. This is either empty or isomorphic to \mathbb{R}^ξ . In the latter case, this is naturally identified with $\mathcal{T}^\infty(\tau)$ — the direct product of $\mathbb{M}^\infty(X)$ for $X \in \mathcal{UX}_T(\tau)$.

Definition. A *Dehn twist flat* in $\mathbb{M}^\infty(\Sigma)$ is a nonempty set of the form $T^*(\tau) \cap \mathbb{M}^\infty(\Sigma)$, where $\tau \subseteq \mathcal{UG}^0(\tau)$ is a complete multicurve.

(We will explain the general term “flat” in this context below.)

By Lemma 11.1, up to the action of $\mathcal{U}\text{Map}(\Sigma)$, we can take τ to be standard. One way to construct $T^\infty(\tau)$ in this case is as follows. Recall that $G(\tau) \cong \mathbb{Z}^\xi$ is the subgroup of $\text{Map}(\Sigma)$ generated by Dehn twists about the components of τ . Let a be any element of $\mathbb{M}(\Sigma)$. The orbit, Ga , is a bounded Hausdorff distance from $T(\tau)$, and so $T^\infty(\tau)$ is the limit of Ga in the asymptotic cone $\mathbb{M}^\infty(\Sigma)$. The natural map from $\mathbb{Z}^\xi \cong G$ to Ga limits on an isomorphism from \mathbb{R}^ξ to $T^\infty(\tau)$, where we view \mathbb{R}^ξ as the asymptotic cone of \mathbb{Z}^ξ .

More generally, if τ is a big multicurve, then $T^\infty(\tau) = T^*(\tau) \cap \mathbb{M}^\infty(\Sigma)$ is either empty or a direct product of ξ \mathbb{R} -trees. In the latter case, it will contain many flats. We aim to show that every (maximal dimensional) flat in $\mathbb{M}^\infty(\Sigma)$ has this form. First, we consider the case of an $S_{1,1}$ or $S_{0,4}$.

Suppose that Σ is an $S_{1,1}$ or $S_{0,4}$. In this case, $\mathbb{G}(\Sigma)$ is a Farey graph, and (up to quasi-isometry) $\mathbb{M}(\Sigma)$ is the dual 3-valent tree (that is the dual to the Farey complex obtained by attaching a 2-simplex to every 3-cycle in $\mathbb{G}(\Sigma)$). To each $\gamma \in \mathcal{G}^0(\Sigma)$ we can associate a bi-infinite geodesic, or *axis*, in $\mathbb{M}(\Sigma)$. Up to bounded Hausdorff distance, we can identify this axis with the space $T(\gamma) = T(\{\gamma\})$ defined in Section 9, which we can, in turn, identify up to quasi-isometry, with $\mathbb{M}(\gamma) = \mathbb{G}(\gamma)$. Any two distinct axes meet in at most a single edge of $\mathbb{M}(\Sigma)$.

As noted before, in this case, $\mathbb{M}^*(\Sigma)$ and $\mathbb{G}^*(\Sigma)$ are both \mathbb{R}^* -trees. If $\gamma \in \mathcal{UG}^0(\Sigma)$, we get a closed convex subset $T^*(\gamma) \subseteq \mathbb{M}^*(\Sigma)$, which can be identified with $\mathbb{M}^*(\gamma) = \mathbb{G}^*(\gamma) \cong \mathbb{R}^*$. If $\alpha, \beta \in \mathcal{UG}^0(\Sigma)$ are distinct, then $T^*(\alpha) \cap T^*(\beta)$ consists of at most one point. The gate map $\omega_\gamma^* : \mathbb{M}^*(\Sigma) \rightarrow T^*(\gamma)$ is the limit of subsurface projection.

We now want to describe flats more generally. In this context, we make the following definition:

Definition. A *flat* in $\mathbb{M}^*(\Sigma)$ is a closed convex subset median isomorphic to \mathbb{R}^ξ .

Note that, in the case of a median metric space, this notion is equivalent to the notion of a flat as defined in Section 3. (In particular, “flats” are always assumed to have maximal rank.) In fact, we know that $\mathbb{M}^*(\Sigma)$ is bi-Lipschitz equivalent to a median metric space, so with respect to this median metric, the two notions coincide.

Let $\Phi \subseteq \mathbb{M}^\infty(\Sigma)$ be a flat. We identify Φ with \mathbb{R}^ξ via a median isomorphism. Given $i \in \{1, \dots, \xi\}$, let $L_i \subseteq \Phi$ be an i -th coordinate line. (Note that two such are parallel. Moreover, they are determined up to permutation of the indices i .) Let $D_i(\Phi) = D(L_i)$, that is, the set of $X \in \mathcal{UX}(\Sigma)$ such that $\psi_X^\infty|_{L_i}$ is injective. (This is independent of the choice of L_i .) We similarly define $D_i^0(\Phi) \subseteq \mathcal{UG}^0(\Sigma)$, which we can identify as a subset of $D_i(\Phi)$.

We now bring Proposition 12.8 into play. Note that if Q is any ξ -cube in Φ , then $D_i(Q) \supseteq D_i(\Phi)$. In fact, there is some ξ -cube $Q_0 \subseteq \Phi$, with $D_i(Q_0) = D_i(\Phi)$ for all i , and so $D_i(Q) = D_i(\Phi)$ for any cube in Φ bigger than Q_0 (that is, with $Q_0 \subseteq \text{hull}(Q)$). In particular, $|D_i^0(\Phi)| \leq 1$. Let $I(\Phi) = \{i \mid D_i^0(\Phi) \neq \emptyset\}$. If $i \in I(\Phi)$, write $D_i^0(\Phi) = \{\gamma_i\}$, and let $\tau(\Phi) = \{\gamma_i \mid i \in I(\Phi)\}$. Thus, $\tau = \tau(\Phi) = \tau(Q_0)$ is a big multicurve. If Q is any bigger cube, then Lemma 12.6 tells us that $Q \subseteq T^\infty(\tau)$. Since $\text{hull}(Q)$ is exhausted by such hulls, we conclude:

Proposition 13.1. *If $\Phi \subseteq \mathbb{M}^\infty(\Sigma)$ is a flat, then $\tau(\Phi)$ is a big multicurve, and $\Phi \subseteq T^\infty(\tau(\Phi))$. Moreover, if $Y \in \mathcal{UX}_N(\tau(\Phi))$, then $Y \in D_i(\Phi)$ for some $i \in \{1, \dots, \xi\} \setminus I(\Phi)$.*

Note also, as in Proposition 12.8, that each $Y \in \mathcal{UX}(\tau)$ lies in $D_i(\Phi)$ for some unique $i \notin I(\Phi)$.

Note that, applying Lemma 12.7 to a large cube in Φ , we see that if $\gamma \in D_i^0(\Phi)$, then for all $j \in \{1, \dots, \xi\} \setminus I(\Phi)$ if $X \in D_j(\Phi)$ is an $S_{1,1}$ or $S_{0,4}$, and $i \neq j$, then $\gamma \wedge X$ (since the second possibility is ruled out by the fact that $D_j^0(\Phi) = \emptyset$).

Next, we aim to describe when two flats meet in a codimension-1 plane (necessarily a coordinate subspace).

Lemma 13.2. *Let Φ_0, Φ_1 , be two flats with $\Phi_0 \cap \Phi_1$ a codimension-1 coordinate plane. Then $\tau = \tau(\Phi_0) \cap \tau(\Phi_1)$ is a big multicurve. Moreover, $|\tau(\Phi_i) \setminus \tau| \leq 1$. If $\beta_0 \in \tau(\Phi_0) \setminus \tau$ and $\beta_1 \in \tau(\Phi_1) \setminus \tau$ then $\beta_0 \neq \beta_1$ and β_0 and β_1 lie in the same complementary component of τ .*

Proof. Choose coordinates on Φ_0 and Φ_1 so that $\Phi_0 \cap \Phi_1$ is a plane orthogonal to the 1st axis, and so that the other coordinates agree on $\Phi_0 \cap \Phi_1$. Write $I_i = I(\Phi_i)$ and $\tau_i = \tau(\Phi_i)$. Let $\tau = \tau_0 \cap \tau_1$. Now $I_0 \setminus \{1\} = I_1 \setminus \{1\}$ (since these sets are determined by lines in $\Phi_0 \cap \Phi_1$). The only case we need to consider is where $1 \in I_0 \cap I_1$ (otherwise, at least one of τ_0 or τ_1 agrees with τ and the statement follows). We aim to show that τ_0 and τ_1 differ only inside a complexity-1 component of $\Sigma \setminus \tau$, and it will follow that τ is big.

So suppose that $1 \in I_0 \cap I_1$. Then $\tau_0 = \tau \cup \{\beta_0\}$ and $\tau_1 = \tau \cup \{\beta_1\}$. Let $Y_i \in \mathcal{UX}_N(\tau)$ be the component containing β_i .

If $Y_0 \neq Y_1$, then $Y_0 \in \mathcal{UX}_N(\tau_1)$, so $Y_0 \in D_i(\Phi_1)$ for some $i \neq 1$ (as observed after Proposition 13.1). But $D_i(\Phi_0) = D_i(\Phi_1)$. In other words, we have $\beta_0 \prec Y_0$, $\beta_0 \in D_1(\Phi_0)$, $Y_0 \in D_i(\Phi_0)$ and Y_0 is an $S_{1,1}$ or $S_{0,4}$. It now follows that $\beta_0 \wedge Y_0$, giving a contradiction.

Thus, $Y_0 = Y_1 = Y$, say. Since $\Phi_0 \neq \Phi_1$, we must have $\beta_0 \neq \beta_1$. We claim that Y is an $S_{1,1}$ or an $S_{0,4}$. For suppose not. We use the fact that τ_0 and τ_1 are big. Either $\beta_0 \pitchfork \beta_1$ or $\beta_0 \wedge \beta_1$. In the former case, we have $\beta_0 \pitchfork Z$ for some $Z \in \mathcal{UX}(\tau_1)$ and we get a contradiction as before. In the latter case, we have $\beta_0 \prec W$ for some $W \in \mathcal{UX}(\tau_1)$ and we derive a similar contradiction.

Thus, Y is an $S_{1,1}$ or $S_{0,4}$. Since τ_0 and τ_1 are big, and differ only in the curves β_0, β_1 , it follows that τ is big. \square

Elaborating on the above proof, we see that there are essentially three possibilities (up to swapping Φ_0 and Φ_1). Let us suppose that Φ_0 and Φ_1 differ in the first coordinate. We have one of the following:

- (1) $\tau(\Phi_0) = \tau(\Phi_1) = \tau$. In this case, there is some $Y \in \mathcal{UX}_N(\tau)$ corresponding to the first factor of both $T(\tau_0)$ and $T(\tau_1)$, so that Φ_0 and Φ_1 project to lines meeting in a single point in the \mathbb{R} -tree $\mathbb{M}^\infty(Y)$.
- (2) $\tau(\Phi_0) = \tau$ and $\tau(\Phi_1) = \tau \cup \{\beta\}$. Let $Y \in \mathcal{UX}_N(\tau)$ be the component containing β . In the \mathbb{R} -tree $\mathbb{M}^\infty(Y)$, Φ_1 projects to the axis corresponding to β , and Φ_0 projects to a line meeting this axis in a single point.
- (3) $\tau(\Phi_0) = \tau \cup \{\beta_0\}$ and $\tau(\Phi_1) = \tau \cup \{\beta_1\}$. Let $Y \in \mathcal{UX}(\tau)$ be the component containing β_0 and β_1 . Then Φ_0 and Φ_1 respectively project to the axes in $\mathbb{M}^\infty(Y)$ corresponding to β_0 and β_1 . These axes intersect in a single point.

We next want to characterise Dehn twist flats.

Lemma 13.3. *Suppose that $\Phi \subseteq \mathbb{M}^\infty(\Sigma)$ is a flat. Suppose that for each i there is another flat $\Phi_i \subseteq \mathbb{M}^\infty(\Sigma)$ with $\Phi \cap \Phi_i$ a codimension-1 coordinate plane orthogonal to the i -th axis. Then Φ is a Dehn twist flat.*

In fact, it is enough to assume the hypothesis for those $i \in I(\Phi)$.

Proof. Suppose $i \in I(\Phi)$. Let $\gamma_i \in \tau(\Phi)$ be the corresponding curve. By Lemma 13.2 and subsequent discussion, we see that $\tau(\Phi_i)$ is obtained from $\tau(\Phi)$ by deleting γ_i and possibly replacing it by another curve in the complementary component of $\tau(\Phi) \setminus \{\gamma_i\}$ that contained γ_i . But $\tau(\Phi_i)$ is big, so either way, it follows that γ_i must lie in an $S_{1,1}$ or $S_{0,4}$ component of the complement of $\tau(\Phi) \setminus \{\gamma_i\}$. Put another way, γ_i bounds an $S_{0,3}$ component of $\Sigma \setminus \tau(\Phi)$ (possibly the same $S_{0,3}$) on each side. Since this holds for all $i \in I(\Phi)$ (that is for all components of $\tau(\Phi)$) it follows that each component of $\Sigma \setminus \tau(\Phi)$ is an $S_{0,3}$. In other words, $\tau(\Phi)$ is complete. \square

For the converse, suppose that Φ is a Dehn twist flat. For simplicity, we can assume that $\tau = \tau(\Phi)$ is standard. Let $G = G(\tau) \subseteq \text{Map}(\Sigma)$ be the subgroup generated by Dehn twists about the components of τ . Thus $G \cong \mathbb{Z}^\xi$. Let $\mathcal{U}G \leq \mathcal{U}\text{Map}(\Sigma)$ be its ultraproduct, and let $\mathcal{U}^0G = \mathcal{U}G \cap \mathcal{U}^0\text{Map}(\Sigma)$. (Recall, from Section 5, that $\mathcal{U}^0\text{Map}(\Sigma)$ is defined to be the setwise stabiliser of $\mathbb{M}^\infty(\Sigma)$.) Then \mathcal{U}^0G acts transitively on Φ , preserving the coordinate directions.

Lemma 13.4. *Suppose that Φ is a Dehn twist flat. Then if Θ is any codimension-1 coordinate subspace in Φ , then there is some Dehn twist flat Ψ with $\Theta = \Phi \cap \Psi$.*

Proof. For simplicity, we can assume $\tau = \tau(\Phi)$ to be standard. Let $\gamma \in \tau$ be the curve corresponding to the coordinate direction perpendicular to Θ . Let $Y \in \mathcal{UX}(\tau \setminus \{\gamma\})$ be the component containing γ . Let $\beta \in \mathbb{G}^0(\Sigma)$ be any other standard curve in Y . Now the axes of β and γ in $\mathbb{G}^\infty(Y)$ meet in a single point. Let $\tau' = (\tau \setminus \{\gamma\}) \cup \{\beta\}$, and let $\Psi = T(\tau')$. Then Ψ is a Dehn twist flat meeting Φ in a codimension-1 plane parallel to Θ . By the homogeneity of Φ described before the statement of the lemma, this is sufficient to prove the result. \square

Putting the above together with Proposition 4.6, we get:

Proposition 13.5. *Suppose that $\Phi \subseteq \mathbb{M}^\infty(\Sigma)$ is a closed subset and that there is a homeomorphism $f : \mathbb{R}^\xi \rightarrow \Phi$ with the following property. For each codimension-1 coordinate plane $H \subseteq \mathbb{R}^\xi$ there is a closed subset $\Psi \subseteq \mathbb{M}^\infty(\Sigma)$ homeomorphic to \mathbb{R}^ξ such that $f(H) = \Phi \cap \Psi$. Then Φ is a Dehn twist flat, and f is a median isomorphism. Moreover, every Dehn twist flat arises in this way.*

In particular, we see that the collection of Dehn twist flats is determined by the topology of $\mathbb{M}^\infty(\Sigma)$, as shown in [Behrstock, Kleiner, Minsky and Mosher 2012]. In fact, we only need an injective map. Moreover, we can take two different surfaces with the same complexity. In summary, we conclude:

Theorem 13.6. *Suppose that Σ and Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 2$. Suppose that we have a continuous injective map $f : \mathbb{M}^\infty(\Sigma) \rightarrow \mathbb{M}^\infty(\Sigma')$ with closed image. If Φ is a Dehn twist flat in $\mathbb{M}^\infty(\Sigma)$, then $f(\Phi)$ is a Dehn twist flat in $\mathbb{M}^\infty(\Sigma')$.*

Note that this applies equally well to any components of $\mathbb{M}^*(\Sigma)$ and $\mathbb{M}^*(\Sigma')$, since they are all respectively isomorphic to $\mathbb{M}^\infty(\Sigma)$ and to $\mathbb{M}^\infty(\Sigma')$.

14. Controlling Hausdorff distance

We begin a general statement, which generalises a construction of [Behrstock, Kleiner, Minsky and Mosher 2012].

Let (M, ρ) be a metric space. Given subsets, $A, B, D \subseteq M$, we say that A, B are r -close on D if $A \cap D \subseteq N(B; r)$ and $B \cap D \subseteq N(A; r)$. (Thus r -close on M

means that the Hausdorff distance, $\text{hd}(A, B)$, from A to B is at most r .) Let t be a positive infinitesimal, and let M^* be the extended asymptotic cone determined by t . Given $e \in M^*$, let M_e^∞ be the component of M^* containing e . Let $r = 1/t$.

Let $\mathcal{UP}(M)$ be the ultrapower of the power set, $\mathcal{P}(M)$, of M . Given $A \in \mathcal{UP}(M)$, let \mathcal{UA} and $A^* \subseteq M^*$ be the images of A under the natural maps $\mathcal{UP}(M) \rightarrow \mathcal{P}(M) \rightarrow \mathcal{P}(M^*)$ (as discussed in Section 5).

The following is a simple observation (a similar statement is used in [Behrstock, Kleiner, Minsky and Mosher 2012]).

Lemma 14.1. *Suppose that $A, B \in \mathcal{UP}(M)$, and $e \in \mathcal{UA}$ (that is $e_\zeta \in A_\zeta$ for almost all ζ). Let $e \in M^*$ be the image of e in M^* (so that $e \in A^*$). Suppose that $\epsilon, R > 0$ are positive real numbers. Then A^*, B^* are ϵ -close on $N(e; R)$ if and only if, for all $R' > R$ and all $\epsilon' > \epsilon$, the sets A_ζ, B_ζ are $\epsilon' r_\zeta$ -close on $N(e_\zeta; R' r_\zeta)$ for almost all ζ .*

In particular, if $A^* \cap M_e^\infty = B^* \cap M_e^\infty$, then for all $R > \epsilon > 0$, the sets A_ζ, B_ζ are almost always ϵr_ζ -close on $N(e_\zeta; R r_\zeta)$. (Here ‘‘almost’’ may depend on ϵ and R .) Note that, in the above, only the component, M_e^∞ , of M^* containing e is relevant.

Lemma 14.2. *Suppose that for all $R > \epsilon > 0$ there is some $e \in A^*$ such that A^*, B^* are ϵ -close on $N(e; R)$. Then, there is some component M^0 of M^* such that $A^* \cap M^0 = B^* \cap M^0 \neq \emptyset$.*

Proof. Given any $n \in \mathbb{N}$, there is some e_n such that A^*, B^* are $1/(2n)$ -close on $N(e_n; 2n)$. Write $e_n = (e_{n,\zeta})_\zeta$. Let \mathcal{Z}_n be the set of $\zeta \in \mathcal{Z}$ such that A_ζ, B_ζ are r_ζ/n -close on $N(e_{n,\zeta}; n r_\zeta)$. Thus, for all n , \mathcal{Z}_n has measure 1. Given $\zeta \in \mathcal{Z}$, let $m(\zeta) = \max(\{n \mid \zeta \in \mathcal{Z}_n\} \cup \{0\}) \in \mathbb{N} \cup \{\infty\}$. Let $p: \mathcal{Z} \rightarrow \mathbb{N}$ be any map with $p(\zeta) \rightarrow \infty$ (for example, any injective map from \mathcal{Z} to \mathbb{N}). Let $n(\zeta) = \min\{m(\zeta), p(\zeta)\} \in \mathbb{N}$. Note that $n(\zeta) \rightarrow \infty$ (since for any $n \in \mathbb{N}$, $p(\zeta) > n$ almost always, and $\zeta \in \mathcal{Z}_n$ so that $m(\zeta) > n$ almost always). Let $e_\zeta = e_{n(\zeta),\zeta}$, and let e be the image of $(e_\zeta)_\zeta$ in A^* . Now, for all n , A_ζ, B_ζ are almost always r_ζ/n -close on $N(e_\zeta; n r_\zeta)$, so A^*, B^* are $1/n$ -close on $N(e; n)$. Since this holds for all n , we have $A^* \cap M_e^\infty = B^* \cap M_e^\infty \neq \emptyset$, as required. \square

Suppose now that \mathcal{E} and \mathcal{F} are collections of subsets of M . We write \mathcal{UE} and \mathcal{UF} for their respective ultrapowers.

We suppose:

- (S1) E is (coarsely) connected for all $E \in \mathcal{E}$.
- (S2) If $F, F' \in \mathcal{UF}$ and there is some component M^0 of M^* such that $F^* \cap M^0 = (F')^* \cap M^0 \neq \emptyset$, then $F = F'$.
- (S3) For all $E \in \mathcal{UE}$, and for all components M^0 of M^* , there is some $F \in \mathcal{UF}$ such that $E^* \cap M^0 = F^* \cap M^0$.

In fact, we only really require (S3) if $E^* \cap M^0 \neq \emptyset$.

(In (S1), “coarsely connected” can be taken to mean that $N(E; s)$ is connected for some fixed $s \in [0, \infty) \subseteq \mathbb{R}$.)

Lemma 14.3. *If \mathcal{E}, \mathcal{F} satisfy (S1)–(S3) above, then there is some $k > 0$ such that for all $E \in \mathcal{E}$, there is some $F \in \mathcal{F}$, such that $\text{hd}(E, F) \leq k$.*

Proof. Suppose not. Let $\epsilon > 0$. Given any $\zeta \in \mathcal{Z}$, there is some $E_\zeta \in \mathcal{E}$ such that for all $F \in \mathcal{F}$, $\text{hd}(E_\zeta, F) > \epsilon r_\zeta$. Let $\mathbf{E} = (E_\zeta)_\zeta \in \mathcal{UE}$. Let e_ζ be any element of E_ζ (so that $e \in E^*$). By (S3), there is some $\mathbf{F} \in \mathcal{UF}$ such that $E^* \cap M_e^\infty = F^* \cap M_e^\infty$. In particular, for all $R > 4\epsilon$, we have that E_ζ, F_ζ are almost always $\epsilon r_\zeta/2$ -close on $N(e_\zeta; 2Rr_\zeta)$. But $\text{hd}(E_\zeta, F_\zeta) > \epsilon r_\zeta$, so there is some $e'_\zeta \in E_\zeta$ such that E_ζ, F_ζ are not $\epsilon r_\zeta/2$ -close on $N(e'_\zeta; 2Rr_\zeta)$. By (S1), we can find $q_\zeta, q'_\zeta \in E_\zeta$ with $\rho(q_\zeta, q'_\zeta)$ bounded such that E_ζ, F_ζ are $\epsilon r_\zeta/2$ -close on $N(q_\zeta; 2Rr_\zeta)$ but not on $N(q'_\zeta; 2Rr_\zeta)$. But by (S3) again, there is almost always some $F'_\zeta \in \mathcal{F}$ such that E_ζ, F'_ζ are $\epsilon r_\zeta/2$ -close on $N(q'_\zeta; 2Rr_\zeta)$. Clearly $F'_\zeta \neq F_\zeta$. It follows that F_ζ, F'_ζ are ϵr_ζ -close on $N(q_\zeta; Rr_\zeta) \subseteq N(q_\zeta; 2Rr_\zeta) \cap N(q'_\zeta; 2Rr_\zeta)$. (Almost always, $\rho(q_\zeta, q'_\zeta) < Rr_\zeta$.) Let $\mathbf{F}' = (F'_\zeta)_\zeta$. We see that $F^*, (F')^*$ are ϵ -close on $N(q; R)$. Since $R > 4\epsilon > 0$ were arbitrary, it follows from Lemma 14.2 that there is some component, M^0 , of M^* such that $F^* \cap M^0 = (F')^* \cap M^0 \neq \emptyset$. By (S2), we have $\mathbf{F} = \mathbf{F}'$. But $F'_\zeta \neq F_\zeta$ almost always, giving a contradiction. \square

We have the following criterion to verify (S2).

Given $A, B \subseteq M$, we say that B *linearly diverges* from A if there are constants $k, t \geq 0$ such that for all $r \geq 0$ and all $x \in B$, there is some $y \in B$ with $\rho(y, A) \geq r$ and $\rho(x, y) \leq kr + t$. We say that a collection \mathcal{F} of subsets of M *linearly diverges* if given any distinct $A, B \in \mathcal{F}$, B linearly diverges from A , with k, t uniform over \mathcal{F} .

Lemma 14.4. *If a family \mathcal{F} of subsets linearly diverges, then it satisfies (S2) above.*

Proof. Suppose that $\mathbf{A}, \mathbf{B} \in \mathcal{UF}$ and $A^* \cap M^0 = B^* \cap M^0 \neq \emptyset$, for some component M^0 of M^* . If $e \in B^* \cap M^0$, then we have $e_\zeta \in B_\zeta$ with $e_\zeta \rightarrow e$. Setting $\epsilon = 1$ and $R > 3k$, we have that A_ζ and B_ζ are almost always r_ζ -close on $N(e; Rr_\zeta)$. If $A_\zeta \neq B_\zeta$, then there is some $y \in B_\zeta$, with $\rho(y, A_\zeta) \geq 2r_\zeta$ and $\rho(e_\zeta, y) \leq 2kr_\zeta + t < 3kr_\zeta$ almost always. Thus, $y \in N(e; Rr_\zeta)$, so we get the contradiction that $\rho(y, A_\zeta) \leq r_\zeta$. Thus $A_\zeta = B_\zeta$ almost always, that is, $\mathbf{A} = \mathbf{B}$. \square

Finally, we apply this to the marking complexes to show that coarse Dehn twist flats get sent (close) to coarse Dehn twist flats under a quasi-isometric embedding.

Suppose that Σ and Σ' are compact surfaces with $\xi = \xi(\Sigma) = \xi(\Sigma')$. Suppose that $\phi: \mathbb{M}(\Sigma) \rightarrow \mathbb{M}(\Sigma')$. This gives rise to a continuous map $\phi^*: \mathbb{M}^*(\Sigma) \rightarrow \mathbb{M}^*(\Sigma')$ with closed image. In fact, each component $\mathbb{M}_e^*(\Sigma)$ of $\mathbb{M}^*(\Sigma)$ gets sent into the component $\mathbb{M}_{\phi^*(e)}^*(\Sigma')$ of $\mathbb{M}^*(\Sigma')$. Moreover, distinct components get sent into distinct components.

Let $\mathcal{F}(\Sigma)$ be the set of coarse twist flats, $T(\tau)$, as τ ranges over all complete multicurves, τ . This satisfies (S1). Also, it is linearly divergent, by Lemma 9.9, and so therefore satisfies (S2) by Lemma 14.4. Note that a Dehn twist flat in a component M^0 of $\mathbb{M}^*(\Sigma)$ is by definition a nonempty set of the form $F^* \cap M^0$ for some $F \in \mathcal{U}\mathcal{F}(\Sigma)$. The same discussion applies to $\mathcal{F}(\Sigma')$.

Let $\mathcal{E} = \{\phi(F) \mid F \in \mathcal{F}(\Sigma)\}$. We claim that $\mathcal{E}, \mathcal{F}(\Sigma')$ satisfies (S3) with $M = \mathbb{M}(\Sigma')$.

Suppose $E \in \mathcal{U}\mathcal{E}$. Then $E = (\phi W_\zeta)_\zeta$, where $W_\zeta \in \mathcal{F}(\Sigma)$. Thus $E^* = \phi^* W^*$, where $W = (W_\zeta)_\zeta$. Suppose that M^0 is a component of $\mathbb{M}^*(\Sigma')$ with $E^* \cap M^0 \neq \emptyset$. Choose any $e \in W^*$ with $\phi^* e \in M^0$. Thus, $M^0 = \mathbb{M}_{\phi^*(e)}^*(\Sigma')$. We see that $\phi^*(\mathbb{M}_e^*(\Sigma)) = M^0 \cap \phi^*(\mathbb{M}^*(\Sigma))$. Now $W^* \cap \mathbb{M}_e^*(\Sigma)$ is a Dehn twist flat in M^0 , so by Theorem 13.6, $S^* \cap M^0 = \phi(W^*) \cap M^0 = \phi(W^*) \cap \mathbb{M}_e^*(\Sigma)$ is a Dehn twist flat in M^0 . In other words, there is some $F \in \mathcal{U}\mathcal{F}(\Sigma')$ with $F^* \cap M^0 = E^* \cap M^0$. This verifies property (S3) for $\mathcal{E}, \mathcal{F}(\Sigma')$.

Lemma 14.5. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 2$, and that $\phi : \mathbb{M}(\Sigma) \rightarrow \mathbb{M}(\Sigma')$ is a quasi-isometric embedding. Then there is some $k \geq 0$ such that if τ is a complete multicurve in Σ , then there is a complete multicurve τ' in Σ' such that $\text{hd}(T(\tau'), \phi T(\tau)) \leq k$.*

Proof. We apply Lemma 14.3 to the sets $\mathcal{E} = \{\phi(F) \mid F \in \mathcal{F}(\Sigma)\}$ and $\mathcal{F} = \mathcal{F}(\Sigma')$. We have verified that \mathcal{E} and \mathcal{F} satisfy (S1)–(S3). \square

As we have stated it (to keep the logic of the argument simpler) the bound k might depend on the particular map ϕ . In fact, it can be seen to depend only on ξ and the parameters of ϕ . For this, fix some parameters of quasi-isometry, and now take \mathcal{E} to the set of all images $\phi(F)$, both as F ranges over the set of coarse Dehn twist flats, $\mathcal{F}(\Sigma)$, and as ϕ ranges over all quasi-isometric embeddings from $\mathbb{M}(\Sigma)$ to $\mathbb{M}(\Sigma')$ with these parameters. To verify (S3) we take $E = (\phi_\zeta W_\zeta)_\zeta$ and apply Theorem 13.6, to the limiting map ϕ^* of $(\phi_\zeta)_\zeta$. The same argument now gives us a uniform constant k independent of any particular ϕ . (See the remark at the end of Section 6.)

15. Rigidity of the marking graph

In this section, we show that, modulo a few exceptional cases, a quasi-isometric embedding between mapping class groups is a bounded distance from a left multiplication (hence a quasi-isometry). This strengthens the result of [Hamenstädt 2005; Behrstock, Kleiner, Minsky and Mosher 2012].

Let (X, ρ) be a geodesic space. Given $A, B \subseteq X$, write $A \sim B$ to mean that $\text{hd}(A, B) < \infty$. Clearly, this is an equivalence relation, and we write $\mathcal{B}(X)$ for the set of \sim -classes. Let $\mathcal{Q}(X) \subseteq \mathcal{B}(X)$ denote the set of \sim -classes of images of bi-infinite quasigeodesics.

If $A, B \in \mathcal{B}(X)$, we write $A \leq B$ to mean that some representative of A is contained in some representative of B . This “coarse inclusion” defines a partial order on $\mathcal{B}(X)$.

We say that two sets $A, B \subseteq X$ have *coarse intersection* if there is some $r \geq 0$ such that for all $s \geq r$, $N(A; r) \cap N(B; r) \sim N(A; s) \cap N(B; s)$ (cf., [Behrstock, Kleiner, Minsky and Mosher 2012]). Clearly, this depends only on the \sim -classes of A and B , and determines an element of $\mathcal{B}(X)$, denoted $A \wedge B$.

Note that if $\phi : X \rightarrow Y$ is a quasi-isometric embedding of X into another geodesic space, Y , then ϕ induces an injective map from $\mathcal{B}(X)$ to $\mathcal{B}(Y)$. Note that this respects inclusion and coarse intersection.

Suppose now that Γ is a group acting by isometry on X . We say that Γ acts *discretely* if for some (or equivalently any) $a \in X$ and any $r \geq 0$, the set $\{g \in \Gamma \mid \rho(a, ga) \leq r\}$ is finite. (In other words, a has finite stabiliser and locally finite orbit.) We will assume the action to be discrete here.

Any subgroup, $G \leq \Gamma$ determines an element $B(G)$ of $\mathcal{B}(X)$, namely the \sim -class of any G -orbit. If $G \leq H \leq \Gamma$, then $B(G) \leq B(H)$, with equality if and only if G has finite index in H . In fact, if $G, H \leq \Gamma$, then $B(G) = B(H)$ if and only if G, H are commensurable in Γ (i.e., $G \cap H$ has finite index in both G and H). More generally, for any $G, H \leq \Gamma$, $B(G)$ and $B(H)$ have coarse intersection, and $B(G \cap H) = B(G) \wedge B(H)$. Note that $B(G)$ is the class of bounded sets if and only if G is finite. Also, the class $B(G)$ contains a bi-infinite quasigeodesic if and only if G is two-ended (virtually \mathbb{Z}) and undistorted in X .

Now, let Σ be a compact surface. Note that $\text{Map}(\Sigma)$ acts discretely on $\mathbb{M}(\Sigma)$. If $\tau \subseteq \Sigma$ is a multicurve, let $G(\tau) \subseteq \text{Map}(\Sigma)$ be the group generated by twists about the elements of τ . Thus, $G(\tau) \cong \mathbb{Z}^{|\tau|}$. Write $B(\tau) = B(G(\tau))$. Note that $B(\tau)$ determines τ uniquely. If τ, τ' are multicurves, then $G(\tau \cap \tau') = G(\tau) \cap G(\tau')$, and so $B(\tau \cap \tau') = B(\tau) \wedge B(\tau')$. Note that if τ is a complete multicurve, then $B(\tau)$ is the class of the coarse Dehn twist flat, $T(\tau)$.

Now if $\gamma \in \mathbb{G}^0(\Sigma)$, then we can always find complete multicurves τ, τ' with $\tau \cap \tau' = \{\gamma\}$. (In fact, we can choose τ, τ' with $\iota(\tau, \tau')$ uniformly bounded.) If $\gamma, \delta \in \mathbb{G}^0(\Sigma)$, then γ, δ are equal or adjacent in $\mathbb{G}(\Sigma)$ if and only if there is a complete multicurve, τ containing both γ and δ . Thus, $B(\gamma), B(\delta) \leq B(\tau)$.

Suppose now that Σ, Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 2$. Suppose that $\phi : \mathbb{M}(\Sigma) \rightarrow \mathbb{M}(\Sigma')$ is a quasi-isometric embedding.

Suppose that $\tau \subseteq \Sigma$ is a complete multicurve. Now Lemma 14.5 gives us a complete multicurve $\tau' \subseteq \Sigma'$ with $\text{hd}(T(\tau'), \phi T(\tau))$ bounded and, in particular, finite. Thus, $\phi(B(\tau)) = B(\tau')$. Moreover, this determines τ' uniquely, and we denote it by $\theta\tau$. Note that, from the remark following Lemma 14.5, we see that the bound depends only on the complexity of the surfaces and the parameters of quasi-isometry.

Suppose that $\gamma \in \mathbb{G}^0(\Sigma)$. Choose complete multicurves τ, τ' with $\tau \cap \tau' = \{\gamma\}$. Thus $B(\tau) \wedge B(\tau') = B(\gamma) \in \mathcal{Q}(\mathbb{M}(\Sigma))$, and so $B(\theta\tau) \wedge B(\theta\tau') \in \mathcal{Q}(\mathbb{M}(\Sigma'))$. It follows that $\theta\tau \cap \theta\tau'$ consists of a single curve, $\delta \in \mathbb{G}^0(\Sigma')$. Note that $B(\delta) = \phi(B(\gamma))$, and we see that δ is determined by γ . We write it as $\theta\gamma$. We have shown that there is a unique map, $\theta : \mathbb{G}^0(\Sigma) \rightarrow \mathbb{G}^0(\Sigma')$ such that $B(\theta\gamma) = \phi B(\gamma)$ for all $\gamma \in \mathbb{G}^0(\Sigma)$. Since $\phi : \mathcal{B}(\mathbb{M}(\Sigma)) \rightarrow \mathcal{B}(\mathbb{M}(\Sigma'))$ is injective, it follows that θ is injective.

Moreover, if γ, δ are equal or adjacent in $\mathbb{G}(\Sigma)$, then $\gamma, \delta \in \tau$ for some complete multicurve τ . So $B(\gamma), B(\delta) \leq B(\tau)$, so $B(\theta\gamma), B(\theta\delta) \leq B(\theta\tau)$, and so $\theta\gamma, \theta\delta$ are equal or adjacent in $\mathbb{G}(\Sigma')$. In other words, θ gives an injective embedding of $\mathbb{G}(\Sigma)$ into $\mathbb{G}(\Sigma')$.

We now use the following fact from [Shackleton 2007]:

Theorem 15.1. *Suppose that Σ and Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 4$. If $\theta : \mathbb{G}(\Sigma) \rightarrow \mathbb{G}(\Sigma')$ is an injective embedding, then $\Sigma = \Sigma'$ and there is some $g \in \text{Map}(\Sigma)$ such that $\theta\gamma = g\gamma$ for all $\gamma \in \mathbb{G}^0(\Sigma)$. The same conclusion holds if Σ, Σ' are both an $S_{2,0}$; if both an $S_{0,6}$; if both an $S_{0,5}$; or if at least one is an $S_{1,3}$, and the other has complexity $\xi = 3$.*

Applying this to our situation, we see that $\Sigma = \Sigma'$, and that there is some $g \in \text{Map}(\Sigma)$ with $\theta\gamma = g\gamma$ for all $\gamma \in \mathbb{G}^0(\Sigma)$. After postcomposing with g^{-1} , we may as well assume that g is the identity. In particular, it follows that $B(\tau) = \phi(B(\tau))$ for all complete multicurves τ in Σ . Now Lemma 14.5 gives us a uniform k such that $\text{hd}(T(\tau'), \phi T(\tau)) \leq k$ for some multicurve τ' in Σ . But we now know that $\tau' = \tau$, and so we deduce that $\text{hd}(T(\tau), \phi T(\tau)) \leq k$ for all multicurves, τ .

Now if $x \in \mathbb{M}(\Sigma)$, we can always find τ, τ' with $\tau \cap \tau' = \emptyset$, and with $\iota(\tau, \tau')$, $\rho(x, T(\tau))$ and $\rho(x, T(\tau'))$ all uniformly bounded. It follows that ϕx is a bounded distance from both $\phi T(\tau)$ and $\phi T(\tau')$ and so $\rho(\phi x, T(\tau))$ and $\rho(\phi x, T(\tau'))$ are also uniformly bounded. But $T(\tau)$ and $T(\tau')$ coarsely intersect in the class of bounded sets. Since there are only finitely many possibilities for the pair τ, τ' up to the action of $\text{Map}(\Sigma)$ we can take the various constants to be uniform. This shows that $\rho(x, \phi x)$ is bounded.

We have shown:

Theorem 15.2. *Suppose that Σ and Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 4$, and that $\phi : \mathbb{M}(\Sigma) \rightarrow \mathbb{M}(\Sigma')$ is a quasi-isometric embedding. Then $\Sigma = \Sigma'$ and there is some $g \in \text{Map}(\Sigma)$ such that for all $a \in \mathbb{M}(\Sigma)$, we have $\rho(\phi a, ga) \leq k$, where k depends only on $\xi(\Sigma) = \xi(\Sigma')$ and the parameters of quasi-isometry of ϕ .*

(Note that if Σ, Σ' are compact surfaces and there is a quasi-isometric embedding of $\mathbb{M}(\Sigma)$ into $\mathbb{M}(\Sigma')$, then certainly $\xi(\Sigma) \leq \xi(\Sigma')$, since the complexity, $\xi = \xi(\Sigma)$, is the maximal dimension of a quasi-isometrically embedded copy of \mathbb{R}^ξ in $\mathbb{M}(\Sigma)$. It is not clear when a quasi-isometric embedding exists if $\xi(\Sigma) < \xi(\Sigma')$.)

One can also describe the lower complexity cases. Note that complexity $\xi = 3$ corresponds to one of $S_{2,0}$, $S_{1,3}$ and $S_{0,6}$. Suppose that $\xi(\Sigma) = \xi(\Sigma') = 3$. Then the result of [Shackleton 2007], quoted as Theorem 15.1 here, tells us if $S_{1,3} \in \{\Sigma, \Sigma'\}$, then again $\Sigma = \Sigma'$ in which case, the conclusion of the theorem holds. Otherwise, it is necessary to assume that $\Sigma = \Sigma'$, and then the conclusion holds. Note that, in fact, the centre of $\text{Map}(S_{2,0})$ is \mathbb{Z}_2 , generated by the hyperelliptic involution. The quotient $\text{Map}(S_{2,0})/\mathbb{Z}_2$ is isomorphic to $\text{Map}(S_{0,6})$. Thus, $\mathbb{M}(S_{2,0})$ and $\mathbb{M}(S_{0,6})$ are quasi-isometric. Of course, the above allows us to describe the quasi-isometric embeddings between them up to bounded distance, as compositions of maps of the above type.

Suppose that $\xi(\Sigma) = \xi(\Sigma') = 2$. In this case $\Sigma \in \{S_{1,2}, S_{0,5}\}$. If $\Sigma = \Sigma' = S_{0,5}$ then the result again holds (using Theorem 15.1). However, if $\Sigma = \Sigma' = S_{1,2}$, then the conclusion of Theorem 15.1 fails without further hypotheses; see [Shackleton 2007]. Note however, that the centre of $\text{Map}(S_{1,2})$ is \mathbb{Z}_2 , and the quotient is isomorphic to the index-5 subgroup of $\text{Map}(S_{0,5})$ which fixes a boundary curve. Therefore $\mathbb{M}(S_{1,2})$ is quasi-isometric to $\mathbb{M}(S_{0,5})$, and this fact allows us again to describe all quasi-isometric embeddings between the marking complexes of surfaces of complexity 2 up to bounded distance. In particular, they are again all quasi-isometries.

The complexity-1 case corresponds to $S_{1,1}$ or $S_{0,4}$. In these cases the marking complexes are quasitrees, and there are uncountably many classes of quasi-isometries between them up to bounded distance. Finally, the mapping class groups of $S_{0,3}$, $S_{0,2}$, $S_{0,1}$ and $S_{0,0}$ are all finite.

Note that this gives a complete quasi-isometry classification of the groups $\text{Map}(\Sigma)$ — they are all different apart from the classes $\{S_{2,0}, S_{0,6}\}$, $\{S_{1,2}, S_{0,5}\}$, $\{S_{1,1}, S_{1,0}, S_{0,4}\}$ and $\{S_{0,3}, S_{0,2}, S_{0,1}, S_{0,0}\}$.

Acknowledgements

Much of the work for this paper was carried out while visiting the Universities of Toulouse, Grenoble, ETH Zürich, Orsay and the Tokyo Institute of Technology. I am grateful to each of these institutions for their hospitality.

References

- [Aougab 2013] T. Aougab, “Uniform hyperbolicity of the graphs of curves”, *Geom. Topol.* **17**:5 (2013), 2855–2875. MR Zbl
- [Bandelt and Hedlíková 1983] H.-J. Bandelt and J. Hedlíková, “Median algebras”, *Discrete Math.* **45**:1 (1983), 1–30. MR Zbl
- [Behrstock 2006] J. Behrstock, “Asymptotic geometry of the mapping class group and Teichmüller space”, *Geom. Topol.* **10** (2006), 1523–1578. MR Zbl
- [Behrstock and Minsky 2008] J. Behrstock and Y. Minsky, “Dimension and rank for mapping class groups”, *Ann. of Math. (2)* **167**:3 (2008), 1055–1077. MR Zbl

- [Behrstock and Minsky 2011] J. Behrstock and Y. Minsky, “Centroids and the rapid decay property in mapping class groups”, *J. Lond. Math. Soc.* (2) **84**:3 (2011), 765–784. MR Zbl
- [Behrstock, Druţu and Sapir 2011] J. Behrstock, C. Druţu, and M. Sapir, “Median structures on asymptotic cones and homomorphisms into mapping class groups”, *Proc. Lond. Math. Soc.* (3) **102**:3 (2011), 503–554. MR Zbl
- [Behrstock, Hagen and Sisto 2015] J. Behrstock, M. F. Hagen, and A. Sisto, “Hierarchically hyperbolic spaces, II: Combination theorems and the distance formula”, preprint, 2015. arXiv 1509.00632
- [Behrstock, Hagen and Sisto 2017] J. Behrstock, M. F. Hagen, and A. Sisto, “Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups”, *Geom. Topol.* **21**:3 (2017), 1731–1804. MR Zbl
- [Behrstock, Kleiner, Minsky and Mosher 2012] J. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, “Geometry and rigidity of mapping class groups”, *Geom. Topol.* **16**:2 (2012), 781–888. MR Zbl
- [Bowditch 2013] B. H. Bowditch, “Coarse median spaces and groups”, *Pacific J. Math.* **261**:1 (2013), 53–93. MR Zbl
- [Bowditch 2014a] B. H. Bowditch, “Embedding median algebras in products of trees”, *Geom. Dedicata* **170** (2014), 157–176. MR Zbl
- [Bowditch 2014b] B. H. Bowditch, “Uniform hyperbolicity of the curve graphs”, *Pacific J. Math.* **269**:2 (2014), 269–280. MR Zbl
- [Bowditch 2015] B. H. Bowditch, “Large-scale rank and rigidity of the Weil–Petersson metric”, preprint, University of Warwick, 2015, Available at <http://homepages.warwick.ac.uk/~masgak/preprints.html>.
- [Bowditch 2016a] B. H. Bowditch, “Large-scale rank and rigidity of the Teichmüller metric”, *J. Topol.* **9**:4 (2016), 985–1020. MR Zbl
- [Bowditch 2016b] B. H. Bowditch, “Quasi-isometric maps between direct products of hyperbolic spaces”, *Internat. J. Algebra Comput.* **26**:4 (2016), 619–633. MR Zbl
- [Bowditch 2016c] B. H. Bowditch, “Some properties of median metric spaces”, *Groups Geom. Dyn.* **10**:1 (2016), 279–317. MR Zbl
- [Chatterji, Druţu and Haglund 2010] I. Chatterji, C. Druţu, and F. Haglund, “Kazhdan and Haagerup properties from the median viewpoint”, *Adv. Math.* **225**:2 (2010), 882–921. MR Zbl
- [Chepoi 2000] V. Chepoi, “Graphs of some CAT(0) complexes”, *Adv. in Appl. Math.* **24**:2 (2000), 125–179. MR Zbl
- [Clay, Rafi and Schleimer 2014] M. Clay, K. Rafi, and S. Schleimer, “Uniform hyperbolicity of the curve graph via surgery sequences”, *Algebr. Geom. Topol.* **14**:6 (2014), 3325–3344. MR Zbl
- [van den Dries and Wilkie 1984] L. van den Dries and A. J. Wilkie, “Gromov’s theorem on groups of polynomial growth and elementary logic”, *J. Algebra* **89**:2 (1984), 349–374. MR Zbl
- [Eilenberg and Steenrod 1952] S. Eilenberg and N. Steenrod, *Foundations of algebraic topology*, Princeton Univ. Press, 1952. MR Zbl
- [Eskin, Masur and Rafi 2017] A. Eskin, H. Masur, and K. Rafi, “Large-scale rank of Teichmüller space”, *Duke Math. J.* **166**:8 (2017), 1517–1572. MR Zbl
- [Farb, Lubotzky and Minsky 2001] B. Farb, A. Lubotzky, and Y. Minsky, “Rank-1 phenomena for mapping class groups”, *Duke Math. J.* **106**:3 (2001), 581–597. MR Zbl
- [Gromov 1993] M. Gromov, “Asymptotic invariants of infinite groups”, pp. 1–295 in *Geometric group theory, II* (Sussex, 1991), edited by G. A. Niblo and M. A. Roller, London Math. Soc. Lecture Note Ser. **182**, Cambridge Univ. Press, 1993. MR Zbl

- [Hamenstädt 2005] U. Hamenstädt, “Geometry of the mapping class groups, III: Quasi-isometric rigidity”, preprint, 2005. arXiv math/0512429
- [Hensel, Przytycki and Webb 2015] S. Hensel, P. Przytycki, and R. C. H. Webb, “1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs”, *J. Eur. Math. Soc.* **17**:4 (2015), 755–762. MR Zbl
- [Isbell 1980] J. R. Isbell, “Median algebra”, *Trans. Amer. Math. Soc.* **260**:2 (1980), 319–362. MR Zbl
- [Kapovich, Kleiner and Leeb 1998] M. Kapovich, B. Kleiner, and B. Leeb, “Quasi-isometries and the de Rham decomposition”, *Topology* **37**:6 (1998), 1193–1211. MR Zbl
- [Kleiner and Leeb 1997] B. Kleiner and B. Leeb, “Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings”, *Inst. Hautes Études Sci. Publ. Math.* **86** (1997), 115–197. MR Zbl
- [Mangahas 2010] J. Mangahas, “Uniform uniform exponential growth of subgroups of the mapping class group”, *Geom. Funct. Anal.* **19**:5 (2010), 1468–1480. MR Zbl
- [Masur and Minsky 1999] H. Masur and Y. Minsky, “Geometry of the complex of curves, I: Hyperbolicity”, *Invent. Math.* **138**:1 (1999), 103–149. MR Zbl
- [Masur and Minsky 2000] H. Masur and Y. Minsky, “Geometry of the complex of curves, II: Hierarchical structure”, *Geom. Funct. Anal.* **10**:4 (2000), 902–974. MR Zbl
- [Rafi and Schleimer 2011] K. Rafi and S. Schleimer, “Curve complexes are rigid”, *Duke Math. J.* **158**:2 (2011), 225–246. MR Zbl
- [Roller 1998] M. Roller, *Poc sets, median algebras and group actions: an extended study of Dunwoody’s construction and Sageev’s theorem*, Habilitationsschrift, Universität Regensburg, 1998. arXiv 1607.07747
- [Shackleton 2007] K. J. Shackleton, “Combinatorial rigidity in curve complexes and mapping class groups”, *Pacific J. Math.* **230**:1 (2007), 217–232. MR Zbl
- [Sholander 1954] M. Sholander, “Medians and betweenness”, *Proc. Amer. Math. Soc.* **5** (1954), 801–807. MR Zbl
- [Verheul 1993] E. R. Verheul, *Multimedians in metric and normed spaces*, CWI Tract **91**, Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam, 1993. MR Zbl
- [Webb 2015] R. C. H. Webb, “Uniform bounds for bounded geodesic image theorems”, *J. Reine Angew. Math.* **709** (2015), 219–228. MR Zbl

Received April 8, 2016. Revised January 4, 2017.

BRIAN H. BOWDITCH
MATHEMATICS INSTITUTE
UNIVERSITY OF WARWICK
COVENTRY, CV4 7AL
UNITED KINGDOM

PACIFIC JOURNAL OF MATHEMATICS

Founded in 1951 by E. F. Beckenbach (1906–1982) and F. Wolf (1904–1989)

msp.org/pjm

EDITORS

Don Blasius (Managing Editor)
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
blasius@math.ucla.edu

Paul Balmer
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
balmer@math.ucla.edu

Kefeng Liu
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
liu@math.ucla.edu

Jie Qing
Department of Mathematics
University of California
Santa Cruz, CA 95064
qing@cats.ucsc.edu

Vyjayanthi Chari
Department of Mathematics
University of California
Riverside, CA 92521-0135
chari@math.ucr.edu

Jiang-Hua Lu
Department of Mathematics
The University of Hong Kong
Pokfulam Rd., Hong Kong
jhlu@maths.hku.hk

Daryl Cooper
Department of Mathematics
University of California
Santa Barbara, CA 93106-3080
cooper@math.ucsb.edu

Sorin Popa
Department of Mathematics
University of California
Los Angeles, CA 90095-1555
popa@math.ucla.edu

Paul Yang
Department of Mathematics
Princeton University
Princeton NJ 08544-1000
yang@math.princeton.edu

PRODUCTION

Silvio Levy, Scientific Editor, production@msp.org

SUPPORTING INSTITUTIONS

ACADEMIA SINICA, TAIPEI
CALIFORNIA INST. OF TECHNOLOGY
INST. DE MATEMÁTICA PURA E APLICADA
KEIO UNIVERSITY
MATH. SCIENCES RESEARCH INSTITUTE
NEW MEXICO STATE UNIV.
OREGON STATE UNIV.

STANFORD UNIVERSITY
UNIV. OF BRITISH COLUMBIA
UNIV. OF CALIFORNIA, BERKELEY
UNIV. OF CALIFORNIA, DAVIS
UNIV. OF CALIFORNIA, LOS ANGELES
UNIV. OF CALIFORNIA, RIVERSIDE
UNIV. OF CALIFORNIA, SAN DIEGO
UNIV. OF CALIF., SANTA BARBARA

UNIV. OF CALIF., SANTA CRUZ
UNIV. OF MONTANA
UNIV. OF OREGON
UNIV. OF SOUTHERN CALIFORNIA
UNIV. OF UTAH
UNIV. OF WASHINGTON
WASHINGTON STATE UNIVERSITY

These supporting institutions contribute to the cost of publication of this Journal, but they are not owners or publishers and have no responsibility for its contents or policies.

See inside back cover or msp.org/pjm for submission instructions.

The subscription price for 2018 is US \$475/year for the electronic version, and \$640/year for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to Pacific Journal of Mathematics, P.O. Box 4163, Berkeley, CA 94704-0163, U.S.A. The Pacific Journal of Mathematics is indexed by Mathematical Reviews, Zentralblatt MATH, PASCAL CNRS Index, Referativnyi Zhurnal, Current Mathematical Publications and Web of Knowledge (Science Citation Index).

The Pacific Journal of Mathematics (ISSN 0030-8730) at the University of California, c/o Department of Mathematics, 798 Evans Hall #3840, Berkeley, CA 94720-3840, is published twelve times a year. Periodical rate postage paid at Berkeley, CA 94704, and additional mailing offices. POSTMASTER: send address changes to Pacific Journal of Mathematics, P.O. Box 4163, Berkeley, CA 94704-0163.

PJM peer review and production are managed by EditFlow® from Mathematical Sciences Publishers.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2018 Mathematical Sciences Publishers

PACIFIC JOURNAL OF MATHEMATICS

Volume 293 No. 1 March 2018

Large-scale rigidity properties of the mapping class groups	1
BRIAN H. BOWDITCH	
Bach-flat isotropic gradient Ricci solitons	75
ESTEBAN CALVIÑO-LOUZAO, EDUARDO GARCÍA-RÍO, IXCHEL GUTIÉRREZ-RODRÍGUEZ and RAMÓN VÁZQUEZ-LORENZO	
Contact stationary Legendrian surfaces in \mathbb{S}^5	101
YONG LUO	
Irreducibility of the moduli space of stable vector bundles of rank two and odd degree on a very general quintic surface	121
NICOLE MESTRANO and CARLOS SIMPSON	
A capillary surface with no radial limits	173
COLM PATRIC MITCHELL	
Initial-seed recursions and dualities for d -vectors	179
NATHAN READING and SALVATORE STELLA	
Codimensions of the spaces of cusp forms for Siegel congruence subgroups in degree two	207
ALOK SHUKLA	
Nonexistence results for systems of elliptic and parabolic differential inequalities in exterior domains of \mathbb{R}^n	245
YUHUA SUN	



0030-8730(201803)293:1;1-W