

RANK AND RIGIDITY PROPERTIES OF SPACES ASSOCIATED TO A SURFACE

BRIAN H. BOWDITCH

ABSTRACT. We describe the large scale geometry of the mapping class group and of the pants graph (or equivalently the Teichmüller space in the Weil-Petersson metric) of a compact orientable surface, from the point of view of coarse median spaces. We derive various results about coarse rank and quasi-isometric rigidity of such spaces. In particular, we show that a quasi-isometric embedding of a mapping class group into itself is a bounded distance from a left multiplication, generalising the result of Hamenstädt and of Behrstock, Kleiner, Minsky and Mosher. We show that, apart from finitely many cases, the pants graphs of different surfaces are quasi-isometrically distinct. We also show that the pants graphs of all but finitely many surfaces are quasi-isometrically rigid.

1. INTRODUCTION

In this paper, we give a number of results regarding the large scale geometry of spaces associated to a compact orientable surface, Σ . We will be mainly interested in the coarse geometry of the mapping class group, $\text{Map}(\Sigma)$, and the pants graph of Σ (the latter being equivalent to studying the Weil-Petersson geometry). In particular, we are aiming at rank and rigidity statements about such spaces. We recover several results that have appeared elsewhere (for example in [Ha, BehM1, BehKMM, EsMR]). We also get some strengthenings of these, and some new results, (see for example, Theorems 1.1, 1.2 and 1.3 below). In [Bo7], we study the Teichmüller metric from a similar perspective.

The central notion for our approach is that of a coarse median. Such a notion can be found in [BehM2] and [BehDS]. Related constructions can be found in [BesBF, BehKMM] and medians are discussed more generally in [ChaDH]. The key observation is that the mapping class group is a coarse median space as defined in [Bo1]. This was proven in [Bo1], based on the construction in [BehM2]. (We give a more self-contained account of this in Section 15 here.) Here we phrase this in

Date: First draft: 3rd April 2014. Revised: 20th October 2014.

terms of the “marking graph” of the surface, which is quasi-isometric to $\text{Map}(\Sigma)$. One can use this to deduce that the pants graph of Σ is also a coarse median space. Recall that the result of [Bro] tells us that the pants graph is quasi-isometric to the Teichmüller space with the Weil-Petersson metric, so for our purposes, this is equivalent.

One can define a notion of “rank” for a coarse median space. In the cases of interest to us here, this turns out to be equal to the maximal dimension of a quasi-isometrically embedded euclidean space (or “quasiflat”). One can compute the rank of $\text{Map}(\Sigma)$ to be the same as the “complexity”, $\xi(\Sigma)$, of the surface: that is $\xi(\Sigma) = 3g + p - 3$, where g is the genus of Σ , and p is the number of boundary components. This was done in [Bo1], thereby recovering the rank theorem in [Ha] and [BehM1]. Here, we also compute the rank of the pants graph of Σ viewed as a coarse median space. This turns out to be equal to $\xi_0(\Sigma) = \lfloor (\xi(\Sigma) + 1) / 2 \rfloor$. We therefore recover another result in [BehM1] (see also [EsMR]) that the maximal dimension of a euclidean space quasi-isometrically embedded in the Weil-Petersson geometry equals $\xi_0(\Sigma)$. It also gives another way of seeing that the Weil-Petersson geometry for a complexity-2 surface is hyperbolic, as shown in [BroF] (see also [Ar]).

Another goal of the paper is to study rigidity properties. For the purposes of this introduction, we will say that a geodesic space, X , admitting an isometric action of $\text{Map}(\Sigma)$ is “quasi-isometrically rigid” if any quasi-isometry of X is a bounded distance from an isometry given by some element of $\text{Map}(\Sigma)$. It was shown in [Ha] and independently in [BehKMM] that for “most” surfaces Σ , $\text{Map}(\Sigma)$ is quasi-isometrically rigid (i.e. any Cayley graph of $\text{Map}(\Sigma)$ is quasi-isometrically rigid). As a consequence one gets a complete quasi-isometric classification of the mapping class groups — they are all different apart from a few low-complexity cases. Here, we give another proof quasi-isometric rigidity, along broadly similar lines to [BehKMM], though we use our median formulation. In fact, we generalise the above statement to quasi-isometric embeddings: As a consequence, we show:

Theorem 1.1. *Suppose that Σ, Σ' are compact orientable surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 4$, and that ϕ is a quasi-isometric embedding of $\text{Map}(\Sigma)$ into $\text{Map}(\Sigma')$. Then $\Sigma = \Sigma'$ and ϕ is a bounded distance from a left multiplication in $\text{Map}(\Sigma)$.*

This is a paraphrasing of Theorem 10.2 here. (In fact the bound depends only on $\xi(\Sigma)$ and the quasi-isometric constants of ϕ .) Note that the existence of a quasi-isometric embedding of $\text{Map}(\Sigma)$ into $\text{Map}(\Sigma')$

implies that $\xi(\Sigma) \leq \xi(\Sigma')$ (from the above discussion of rank). It is not clear when such embeddings exist in general.

Given that the pants graph is also a coarse median space, one can apply similar technology there (though the details are different). From this we obtain a quasi-isometric classification for “most” cases as follows:

Theorem 1.2. *Suppose that Σ, Σ' are compact orientable surfaces whose pants graphs are quasi-isometric. Then $\xi(\Sigma) = \xi(\Sigma')$. Moreover, if $\xi(\Sigma) = \xi(\Sigma') \geq 6$ then $\Sigma = \Sigma'$.*

This gives rise to a complete quasi-isometric classification except in the cases of cases of complexity 4 and 5, which we leave unresolved here. Theorem 1.2 is a paraphrasing of Theorem 14.7 here.

We reduce the question of quasi-isometric rigidity to the rigidity of certain subgraphs of the curve graph of a surface. For surfaces with at most one boundary component, this is the same as the separating curve graph, shown to be rigid in [BreM, Ki]. Also, [Bo6] deals with all surfaces with sufficiently high complexity (in fact $\xi(\Sigma) \geq 58$). From this, we we obtain:

Theorem 1.3. *If Σ is a closed surface of genus at least 3, or a surface with one boundary component of genus at least 2, or of complexity at least 58, then the associated pants graph is quasi-isometrically rigid.*

This is a consequence of Theorem 14.8 here. We expect that, in fact, most pants graphs are rigid, as we aim to explore elsewhere.

We remark that the corresponding statements for isometries are known. In [MasW] is shown that if $\xi(\Sigma) \geq 2$, then any isometry of the Weil-Petersson metric is induced by any element of $\text{Map}(\Sigma)$. Similarly, in [Mar] it is shown that any automorphism of the curve graph is induced by an element of $\text{Map}(\Sigma)$ (with a few qualifications of the low-complexity cases). In [BroM] this was used to give another proof of the result of [MasW].

As noted above, we base our account around the notion of a coarse median space, which might be thought of as a coarse version of a median metric space, which is also central to discussion. The definition of a median metric space is quite simple, and is given in Section 2. For further discussion, see [Ve, ChaDH, Bo1]. This has also been studied from a combinatorial viewpoint, see for example, [Che]. 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 [Is, BaH, Ro].) One can associate a “rank” to such a space as the maximal dimension of an embedded cube. One can show that a complete connected median metric space of finite rank, say n , is canonically bilipschitz equivalent to a CAT(0) metric, [Bo4].

The asymptotic cone (see [VaW, G]) of a coarse median space is a topological median algebra. If the space has finite rank, n , then the asymptotic cone is bilipschitz equivalent to a median metric space of rank at most n , (see [Bo2] and Theorem 6.5 here). Also, the dimension of any compact subset thereof has dimension at most n (see [Bo1]). 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 [BehDS]. Much of this is elaborated upon in [Bo1, Bo2]. Here we obtain more information about the flats in such spaces, which we use for the rigidity result Theorem 1.1. Similar statements can be found in [BehKMM], though more specifically for the mapping class group.

One can also apply this to the pants graph. Much of the general theory is the same, though the details are different. Quasiflats are more complicated to describe, so instead we use quasi-isometrically embedded direct products of hyperbolic spaces. There arise from certain classes of multicurves. They give rise to products of trees in the asymptotic cone. which we show can be recognised just from the topology. From this one constructs a certain combinatorial complex, and derive Theorems 1.2 and 1.3 from this. It would be interesting to generalise this to quasi-isometric embeddings but there are complications in applying the same arguments as with the marking graph.

We remark that, in [RaS], 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 [Ha, BehKMM]) where an inverse quasi-isometry is needed. It would also be interesting to know whether the (hyperbolic) pants graph in complexity-2 is rigid.

The outline of this paper is as follows. In Sections 2 to 4 we discuss median metric spaces quite generally. In Section 5 we review properties of asymptotic cones. In Section 6 we discuss general coarse median spaces. This is then applied to the mapping class group (or marking graph) in Sections 7 to 10, though some of the discussion is more general, and will be used again for the pants graph. In Section 11 we return to general coarse median spaces to describe quotient maps. This is used in Section 12 to show that the pants graph is coarse median of rank $\xi_0(\Sigma)$. This is then applied to rigidity of the pants graph in

Sections 13 and 14. Finally, we in Section 15, we give a more or less self-contained account of the existence of medians.

Much of the work of 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.

2. MEDIAN METRIC SPACES

We begin with some general discussion of median metric spaces. For elaboration relevant to this paper, see for example, [Ve, ChaDH, Bo4].

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 that 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 [Sho] that (M, μ) is a median algebra (see [Ve, ChaDH, Bo4]). Moreover, $[a, b]$ is exactly the median interval between a and b , i.e. $[a, b] = \{x \in M \mid \mu(a, b, x) = x\}$. Also, the map $\mu : M^3 \rightarrow M$ is continuous.

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: $\{-1, 1\}^n$. 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 [Che]).

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. We will generally write ν to denote $\text{rank}(M)$.

It was shown in [Bo2] 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.) Moreover, M is locally convex, that is, every point has a base of convex neighbourhoods. It was also shown in [Bo1] that any compact subset of M has topological dimension at most ν . (For more discussion of dimension, see Section 4.)

The following was shown in [Bo4].

Theorem 2.1. *If (M, ρ) is a proper median metric space, then there is a canonically associated bilipschitz equivalent metric, σ_ρ , on M for which (M, σ_ρ) is $CAT(0)$.*

In fact, we can arrange that $\rho/\sqrt{\nu} \leq \sigma_\rho \leq \rho$.

A simple example is \mathbb{R}^n in 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 parallelopipeds. 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 [Bo4].

Lemma 2.2. *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 Υ 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.)

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.

We say that two pairs $(a, b), (c, d)$ in M^2 are *parallel* if $a = b$ and $c = d$ and/or if $a = c$ and $b = d$ or if a, b, d, c is a square. Parallelism is an equivalence relation on M^2 . In this case, $\phi_{a,b}|_{[c, d]}$ is an isometry (hence a median isomorphism) from $[c, d]$ to $[a, b]$. Its inverse is $\phi_{c,d}|_{[a, b]}$.

More generally, 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 $c \in C$. 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.

A *wall* in M is a partition of M into two non-empty 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 $\phi \times \psi : M \rightarrow \{-1, 1\}^2$ is surjective. The rank of M can be equivalently defined as the maximal cardinality of a set of pairwise crossing walls.

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.

3. BLOCKS

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

Let M be a proper median metric space. Set $\nu = \text{rank}(M)$.

Definition. An n -block in M is a convex subset isometric to an n -dimensional l^1 cube.

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. *The following are equivalent for a subset $P \subseteq M$:*

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

Proof. The fact that (2) implies (1) was proven in [Bo4]. 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 previous fact, 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 (2) 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. *Suppose that P, P' 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 non-empty. 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, $Q(P \cup P') \subseteq \langle Q(P) \cup Q(P') \rangle$.

We have the following procedure for subdividing blocks. Suppose that $P \equiv \prod_{i=1}^n I_i$. 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 see that we can realise $\bigcup \mathcal{P}$ as an l^1 cube complex in M all of whose cells are blocks. We will refer to such a subset as a *straight cube complex* in M .

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 $L \cap P$ lies 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, 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.

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 non-singular:

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 H_1 is 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 $I \subseteq \mathbb{R}^\nu$ 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 ν -block in Φ_S , the $f^{-1}|_P$ is a median isomorphism to a block $f^{-1}(P)$ in \mathbb{R}^ν . 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 , with y as a vertex. The above shows that the leaf structure of Υ is topologically standard, and so the link of y in Υ is a cross polytope, and so y is regular.

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

The following will also be useful later (see Proposition 8.7).

Lemma 3.9. *Suppose that P is a ν -block, and that C is a convex subset of M disjoint from P . Let $\phi : M \rightarrow P$ be the projection (gate map) to P . Then $\phi(C)$ is contained in a proper coordinate plane of P .*

Proof. Suppose not. Then $\phi(C)$ contains a ν -cube, Q . There is a collection of ν walls of M which respectively separate the pairs of opposite $(\nu - 1)$ -faces of Q . There is another wall which separates C from P . Together these give a collection of $\nu + 1$ pairwise crossing walls in M , contradicting the fact that $\nu = \text{rank}(M)$. \square

Corollary 3.10. *Suppose that P is a ν -block in M with projection $\phi : M \rightarrow P$. If $K \subseteq M \setminus P$ is compact, then $\phi(K)$ is nowhere dense in P .*

Proof. By local convexity of M , P is contained in a finite union of convex subsets of M disjoint from P . By Lemma 3.9, we see that $\phi(K)$ is in fact contained in a finite union of proper coordinate planes. \square

In particular, if P, P' are ν -blocks, and the projection from P' to P is injective, then $P' \subseteq P$.

For reference elsewhere, we note that there is a variation on Lemma 3.8, where \mathbb{R}^ν is replaced by a real cube, $[-1, 1]^\nu$, and $\Phi = f([-1, 1]^\nu)$. In fact, it is enough to assume that the relative interior, $f((-1, 1)^\nu) \subseteq \Phi$, is cubulated (in the sense that any compact subset of $f((-1, 1)^\nu)$ 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 [KIL, KaKL, BehKMM]

etc. Our argument is analogous to those the be found there. Here, we interpret this in terms of cubulations. We will only use dimension of (locally) compact sets, so all the standard definitons 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 [KIL], where, from context, this is defined in terms of singular homology.)

Suppose that M is a complete median metric space. If M is homeomorphic to \mathbb{R}^ν . Then $\nu = \text{rank}(M)$. To see this, note that it was shown in [Bo1] that (any compact subset) of M has topological dimension at most $\text{rank}(M)$, and so $\nu \leq \text{rank}(M)$. For the other direction, note that by Lemma 2.2, any n -cube in M is the vertex set of an embedded l^1 cube in M , and so $n \leq \nu$. We see that $\text{rank}(M) \leq \nu$.

It follows that M is a proper median metric space. In fact:

Lemma 4.1. *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 4.1. 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 2.2, 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 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 between f and the identity map in 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, ρ , on M , as given by Theorem 2.1. 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 $f : B_0 \rightarrow \Upsilon$. Combining this with the homotopy constructed above, we get a continuous map $g : B_0 \rightarrow M$ which restricts to inclusion on ∂B_0 . Therefore $B_0 \subseteq g(B_0)$. Since B_1 misses the image of the homotopy, we see 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, by a simple dimension argument, B_1 must lie in the union of these ν -cells. 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.

Lemma 4.2. *Suppose that M is a proper 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.

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

Lemma 4.3. *Suppose that X is a hausdorff topological space and that $B, P \subseteq X$ as 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 to the “separation dimension” as defined in [Bo1].

Lemma 4.4. *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 [KIL]. A related, but slightly different statement can be found in [BehKMM]. 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 Chapter IX of [EiS]. (Note that in [KIL], it is implicit from context that singular homology is begin 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 $C = f_0(M)$, $D = f_1(B)$, $S = f_0(\partial B) = f_1(\partial B)$ 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 Lemma 4.2. We recall that M is contractible [Bo4], and has compact dimension at most ν [Bo1].

Proof of Lemma 4.2. This is an extention of the argument for Lemma 4.1. 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 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 4.4 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 a simple dimension argument, 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 4.3 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 non-empty. 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 Lemmas 4.2 and 3.8, we get:

Lemma 4.5. *Suppose 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 .

This is all we will need for the discussion of the marking graph. The following is relevant to the pants graph (see Section 13).

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

Proposition 4.6. *Suppose that M is a 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 Lemma 4.5, and so $f|_{\Phi}$ is a median homomorphism. Now any two points, a, b lies 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 4.6 applies if M is also a product of ν \mathbb{R} -trees, and it follows that f splits as a direct product of embeddings, up to permutation of the factors. Some further discussion of this can be found in [Bo5].

Definition. A *tree product*, T , in M is a convex subset median isomorphic to a direct product of ν non-trivial 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 4.6).

While the above is sufficient for the applications in this paper, we note for applications elsewhere that we can relax the “furriness” condition somewhat. For example, we say that 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. It is easy to see that the above discussion, in particular Theorem 4.6, generalises to products of almost furry trees. We can apply the arguments to the maximal subset which is a product of furry trees, and take the closure. We have already observed that the key statements, in particular Lemma 3.8 and 4.2 have local versions which can be applied to this case. We use the above generalisation in [Bo7].

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 [VaW] (see also [G]). 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 [G], rather than in the 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 non-principal 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 ζ , we say that P holds *almost always* (a.a.) 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 $\mathbf{X} = (X_\zeta)_\zeta$ for such a sequence. If these are all sets, we write $\prod \mathbf{X} = \prod_\zeta X_\zeta$ for their product. Given $\mathbf{x}, \mathbf{y} \in \prod \mathbf{X}$, we write $\mathbf{x} \approx \mathbf{y}$ to mean that $x_\zeta = y_\zeta$ almost always. Thus, \approx is an equivalence relation on $\prod \mathbf{X}$, and we write $\mathcal{U}\mathbf{X} = \prod \mathbf{X} / \approx$ for the quotient — the *ultraproduct* of \mathbf{X} . Note that we only need to have x_ζ defined almost always to determine an element of $\mathcal{U}\mathbf{X}$. We write $x = [\mathbf{x}]$ for this element.

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

Given two \mathcal{Z} -sequences of sets, \mathbf{X} and \mathbf{Y} , we can form the direct product $\mathbf{X} \times \mathbf{Y}$ as $(X_\zeta \times Y_\zeta)_\zeta$, and we see that $\mathcal{U}(\mathbf{X} \times \mathbf{Y})$ is naturally identified with $\mathcal{U}\mathbf{X} \times \mathcal{U}\mathbf{Y}$. A \mathcal{Z} -sequence of relations on $X_\zeta \times Y_\zeta$ give rise to a relation on $\mathcal{U}\mathbf{X} \times \mathcal{U}\mathbf{Y}$ via the map from $\mathcal{U}\mathcal{P}(\mathbf{X} \times \mathbf{Y})$ to $\mathcal{P}(\mathcal{U}\mathbf{X} \times \mathcal{U}\mathbf{Y})$. In other words, x is related to y if x_ζ is almost always related to y_ζ . If the relation on $X_\zeta \times Y_\zeta$ is (almost always) the graph of a function, then the relation induced on $\mathcal{U}\mathbf{X} \times \mathcal{U}\mathbf{Y}$ is also the graph of a function. In other words, a \mathcal{Z} -sequence of functions $f_\zeta : X_\zeta \rightarrow Y_\zeta$ determines a function $\mathcal{U}f : \mathcal{U}\mathbf{X} \rightarrow \mathcal{U}\mathbf{Y}$, where $y = \mathcal{U}f(x)$ is equivalent to the statement that $y_\zeta = f_\zeta(x_\zeta)$ almost always. Note that the above discussion also applies to finite products of sets, and so to

n -ary relations and n -ary operations for any finite n . For example, if Γ is a sequence of groups, the $\mathcal{U}\Gamma$ has the structure of a group. If each Γ_ζ acts on a set X_ζ , then $\mathcal{U}\Gamma$ acts on $\mathcal{U}\mathbf{X}$.

Suppose that $X_\zeta = X$ is constant. In this case, we write $\mathcal{U}X = \mathcal{U}\mathbf{X}$, and refer to $\mathcal{U}X$ as the *ultrapower* of 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 in the ultrapower, $\mathcal{U}\mathbb{R}$, of the real numbers is an ordered field. We say that $x \in \mathcal{U}\mathbb{R}$ is *limited* if $|x| \leq y$ for some $y \in \mathbb{R}$. Otherwise it is *unlimited*. We say that x is *infinitesimal* if $|x| \leq y$ for all positive standard y . Note that 0 is the only standard infinitesimal, and that non-zero infinitesimals are exactly the reciprocals of unlimited numbers.

There is a well defined map $\text{st} : \mathcal{U}\mathbb{R} \rightarrow \mathbb{R} \cup \{\infty\}$ such that $\text{st}(x) = \infty$ if x is unlimited, and $x - \text{st}(x)$ is infinitesimal if x is limited. We refer to $\text{st}(x)$ as the *standard part* of x . We will usually restrict attention to non-positive 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 *ultrametric* 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 an ultrametric, the composition $\hat{\sigma} = \text{st} \circ \sigma : M^2 \rightarrow [0, \infty]$ is an extended pseudometric on M . Here, we use the term *extended* 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 extended metric on \hat{M} .

Given an extended metric space, \hat{M} , we say that two points $x, y \in \hat{M}$ *lie in the same component* if the distance between them is not ∞ . Note that components are both open and closed in the topology induced on \hat{M} . They are also convex in the median structure.

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

As with real numbers, 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.

Lemma 5.1. *$(\hat{X}, \hat{\sigma})$ 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 \mathbf{x}^i be some representative of x^i in $\mathbf{X} = (X_\zeta)_\zeta$. 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 $\mathbf{y} = (y_\zeta)_\zeta$ 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, let $i : \mathcal{Z}(\infty) \rightarrow \mathbb{N}$ be any bijection. We set $y_\zeta = x_\zeta^{i(\zeta)}$, and argue as before. \square

Suppose that $A_\zeta \subseteq X_\zeta$ (almost always). As discussed earlier, this gives rise to a subset of $\mathcal{U}\mathbf{X}$ which can be identified with $\mathcal{U}\mathbf{A}$. 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 5.1, 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*.

Suppose that $f_\zeta : X_\zeta \rightarrow Y_\zeta$ is a \mathbb{Z} sequence of maps between the metric spaces (X_ζ, σ_ζ) and (Y_ζ, σ'_ζ) . We get a map, $\mathcal{U}f : \mathcal{U}\mathbf{X} \rightarrow \mathcal{U}\mathbf{Y}$ as before. Suppose there is a constant, $k \in [0, \infty)$, and an infinitesimal constant $(h_\zeta)_\zeta$ in $\mathcal{U}[0, \infty)$, such that for almost all ζ and all $x, y \in X_\zeta$ we have $\sigma'_\zeta(f_\zeta(x), f_\zeta(y)) \leq k\sigma_\zeta(x, y) + h_\zeta$. (In other words, the f_ζ are uniformly coarsely lipschitz.) 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. If the maps f_ζ are all quasi-isometric embeddings, then \hat{f} is bilipschitz onto its range.

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, $t \in \mathcal{UR}$, and set $\sigma_\zeta = t_\zeta \rho$ 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 $(\hat{X}, t\hat{\rho})$ and passing to its Hausdorffification. We refer to (X^*, ρ^*) as the *extended asymptotic cone* of X with respect to t .

Note that X^* has a preferred basepoint, namely that given by any constant sequence in \mathbf{X} . This, in turn, determines a preferred component, X^∞ , of X^* , namely that containing this basepoint. We refer to X^∞ as the *asymptotic cone* of X with respect to t . By Lemma 5.1, the asymptotic cone is always complete. If X is a geodesic space, so is X^∞ . Note also that quasi-isometric spaces give rise to bilipschitz equivalent asymptotic cones.

One can generalise the above to a sequence of metric spaces, (X_ζ, ρ_ζ) , rescaled by an infinitesimal 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 base point and base component of X^* . As before, a sequence of uniformly coarsely Lipschitz maps $f_\zeta : X_\zeta \rightarrow Y_\zeta$ between such spaces gives rise to a Lipschitz map, $f^* : X^* \rightarrow Y^*$. If the f_ζ are quasi-isometric embeddings, then f^* is bilipschitz onto its image, which is necessarily closed. If they are all quasi-isometries, then f^* is a bilipschitz equivalence.

An example of the above construction is given by a sequence, $\mathbf{G} = (G_\zeta)_\zeta$ 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{UV} , so as to give it the structure as the vertex set, $V(\mathcal{UG})$ of a graph \mathcal{UG} . If each G_ζ is connected, the combinatorial distance functions on V_ζ give us a limiting ultrametric and hence an extended metric on \mathcal{UV} , with values

in $\mathbb{N} \cup \{\infty\}$. This is the same as the combinatorial extended metric given by $\mathcal{UV} = V(\mathcal{UG})$. 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 $\mathbf{\Gamma} = (\Gamma_\zeta)_\zeta$ is a \mathcal{Z} -sequence of groups. Then $\mathcal{U}\mathbf{\Gamma}$ is also a group. If each Γ_ζ acts on a set X_ζ , then $\mathcal{U}\mathbf{\Gamma}$ acts on $\mathcal{U}\mathbf{X}$. If Γ_ζ acts by isometry in some metric space, then so does $\mathcal{U}\mathbf{\Gamma}$. If Γ and X are fixed, then any two points of $X \subseteq \mathcal{U}\mathbf{X}$ in the same $\mathcal{U}\mathbf{\Gamma}$ -orbit also lie in the same Γ -orbit (since if $y = gx$ for some $g \in \mathcal{U}\mathbf{\Gamma}$, 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 $\mathcal{U}\mathbf{\Gamma}$ on the extended asymptotic cone, X^* (with respect to any infinitesimal t). Note that we can identify Γ as a normal subgroup of $\mathcal{U}\mathbf{\Gamma}$. In fact, we have normal subgroups, $\Gamma \triangleleft \mathcal{U}^1\mathbf{\Gamma} \triangleleft \mathcal{U}^0\mathbf{\Gamma} \triangleleft \Gamma_\infty$ of Γ_∞ , where $\mathcal{U}^0\mathbf{\Gamma}$ is the stabiliser of the basepoint of X^* , and $\mathcal{U}^0\mathbf{\Gamma}$ is the setwise stabiliser of the asymptotic cone, X^∞ . Note that $\mathcal{U}^1\mathbf{\Gamma}$ and $\mathcal{U}^0\mathbf{\Gamma}$ may depend to 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 $\mathcal{U}\mathbf{\Gamma}$ on X^* and of $\mathcal{U}^0\mathbf{\Gamma}$ 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.

6. COARSE MEDIAN SPACES

Coarse median spaces were defined in [Bo1]. 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 [BehM2], this is applicable to the mapping class group, as shown in [Bo1]. In Section 12 here, we show that it also applies to the pants graph (or equivalently, the Weil-Petersson metric).

Let (Λ, ρ) be a geodesic metric space. We say that a ternary operation, $\mu : \Lambda^3 \rightarrow \Lambda$, is a ‘‘coarse median’’ if it satisfies the following:

(C1): There are constants, $k, h(0)$, such that for all $a, b, c, a', b', c' \in \Lambda$ 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 $A \subseteq \Lambda$ with $1 \leq |A| \leq p < \infty$, then there is a finite median algebra, (Π, μ_Π) and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ 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 Λ has *rank at most n* if we can always take Π to have rank at most n (as a median algebra). We refer to (Λ, ρ, μ) as a *coarse median space*. We refer to k, h as the *parameters* of Λ .

From (C2) we can deduce that, if $a, b, c \in \Lambda$, 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 .

We will refer to the constants in the definition (i.e. k and the function h) as the *parameters* of Λ .

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

An example of a coarse median space is a Gromov hyperbolic space, where the median of three points is the centre of any geodesic triangle with vertices at these points. This has rank 1. In fact, any rank-1 coarse median space arises in this way.

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 *h -quasimorphism* if $\rho(\phi\mu_X(x, y, z), \mu_Y(\phi x, \phi y, \phi z)) \leq h$ 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 a $h(p)$ -quasimorphism.)

Lemma 6.1. *Suppose that Π is a median algebra generated by a finite subset, $B \subseteq \Pi$. Suppose that $\lambda, \lambda' : \Pi \rightarrow \Lambda$ are h -quasimorphisms with $\rho(\lambda b, \lambda' b) \leq h$ for all $b \in B$. Then, for all $x \in \Pi$, $\rho(\lambda x, \lambda' x)$ is bounded above by some linear function of h , depending only on the parameters of Λ 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 h . 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 .

Lemma 6.2. *Given $n \in \mathbb{N}$, there are constants k_0, h_0 and h_1 depending only on n and the parameters of Λ such that the following holds. Suppose that $Q = \{-1, 1\}^n$ and that $\psi : Q \rightarrow \Lambda$ is an h -quasimorphism for some $h \geq 0$. Then there is an h_0 -quasimorphism, $\phi : Q \rightarrow \Lambda$, with $\rho(\phi x, \psi x) \leq k_0 h + h_1$ 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 . We take $\omega(Q)$ to be the vertex set of such an n -cell. Note that ω is not canonically determined.)

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

Let $\lambda' = \lambda \circ \theta : \Pi \rightarrow \Lambda$. Thus λ' is a h -quasimorphism, and $\lambda'|_Q = \psi = \lambda|_Q$. By Lemma 6.1, we have $\rho(\lambda x, \lambda' x) \leq k_0 h + h_1$ for all $x \in \Pi$, where k_0, h_1 depend only on the parameters of Λ . But $\lambda' \circ \omega|_Q = \lambda \circ \theta \circ \omega|_Q = \lambda|_Q = \psi$, and so we see that $\rho(\phi x, \psi x) \leq k_0 h + h_1$ for all $x \in Q$ as required. \square

Suppose that $((\Lambda_\zeta, \rho_\zeta, \mu_\zeta))_\zeta$ is a \mathcal{Z} -sequence of coarse median spaces, with uniform parameters, and let $t \in \mathcal{UR}$ be a positive infinitesimal. We get a limiting space, $(\Lambda^*, \rho^*, \mu^*)$, where (Λ^*, ρ^*) is the extended asymptotic cone, and where (Λ^*, μ^*) is a topological median algebra (that is, the map $\mu^* : (\Lambda^*)^3 \rightarrow \Lambda^*$ is continuous). If each Λ_ζ has rank at most n (as a coarse median space) then Λ^* has rank at most n (as a median algebra). Note that (Λ^*, μ^*) need not be a median metric space, though it satisfies a weaker metric condition described in [Bo1, Bo2]. (In those papers, we restricted to the asymptotic cone, Λ^∞ , but that does not affect the above observations.)

Lemma 6.3. *Suppose that $Q \subseteq \Lambda^*$ is an n -cube. Then we can find a sequence of h_0 -quasimorphisms, $\phi_\zeta : Q \rightarrow \Lambda_\zeta$ such that for all $x \in Q$, $\phi_\zeta x \rightarrow x$, where h_0 depends only on n and the parameters of Λ .*

Proof. To begin, take any sequence of maps $\psi_\zeta : Q \rightarrow \Lambda_\zeta$, 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^* , 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 Λ^*). Let $\phi_\zeta : Q \rightarrow \Lambda_\zeta$ be the h_0 -monomorphism given by Lemma 6.2. 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

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

From property (C1) we see that any quasiconvex set is quasi-isometrically embedded in Λ .

Similarly as with a median algebra, we say that a map $\phi : \Lambda \rightarrow C$ is a *h-quasiprojection* (or *coarse gate map*) if for all $x \in \Lambda$ and $c \in C$, $\rho(x, \mu(x, \phi x, c)) \leq h$. Such a is well defined up to bounded distance. Note (again using Property (C1)) that $\rho(x, \phi(x))$ is bounded above by a linear function of $\rho(x, C)$.

Note that, if we have a sequence of uniformly quasiconvex sets, $C_\zeta \subseteq \Lambda_\zeta$, we have a limiting bilipschitz embedded closed convex subset, $C^* \subseteq \Lambda^*$ in the extended asymptotic cone Λ^* . If $\phi_\zeta : \Lambda_\zeta \rightarrow C_\zeta$ are a sequence of uniform quasiprojections, the limiting map $\hat{\phi} : \Lambda^* \rightarrow C^*$ is a gate map.

As in [Bo1], we say that a median algebra, Π , is *n-colourable* if there is an *n-colouring* of the walls that no two walls of the same colour cross. We say that a coarse median space Λ is *n-colourable* if in (C2) we can always choose Π to be *n-colourable* as a median algebra. Clearly this implies that Λ has rank at most n . The following was shown in [Bo2].

Theorem 6.4. *Suppose that $((\Lambda_\zeta, \rho_\zeta, \mu_\zeta))_\zeta$ is a sequence of *n-colourable* coarse median spaces, for some fixed n . Then Λ^* admits a metric, ρ' , bilipschitz equivalent to ρ^* , such that (Λ^*, ρ') is an (extended) median metric space with median μ^* . Moreover, Λ^* is *n-colourable* as a median algebra.*

Here, μ^* is the limiting median constructed earlier. In fact, the bilipschitz constants depend only on n . 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 space is bilipschitz equivalent to a proper median metric space, and hence in turn to a CAT(0) space (by Lemma 2.1). In fact, the same holds for a sequence finite rank coarse median spaces. This relies on the following variation of Theorem 6.4.

Theorem 6.5. *Suppose that $((\Lambda_\zeta, \rho_\zeta, \mu_\zeta))_\zeta$ is a sequence of coarse median spaces of rank n , for some fixed n . Then Λ^* admits a metric, ρ' , bilipschitz equivalent to ρ^* , such that (Λ^*, ρ') is an (extended) median metric space of rank n , with median μ^* .*

Proof. This calls for a slight variation on the argument of Section 6 of [Bo2]. Given a finite median algebra, $\Pi \subseteq \Lambda^*$, define a metric λ^Π , on Π as in Section 5 of [Bo2]. Now λ^Π is a median metric on Π . (This is easily seen from the construction, see also the discussion in [Bo4].) A compactness argument as in Section 6 of [Bo2] now gives us a metric, ρ' , on Λ^* (as a convergent subsequence of metrics on larger and larger finite subalgebras). Lemmas 5.2 and 6.2 of [Bo2] now tell us that ρ' is bilipschitz equivalent to ρ . \square

A particular example arises when $(\Lambda_\zeta)_\zeta$ is a sequence of uniformly hyperbolic spaces. (These are the same as coarse median of rank 1.) In this case, (Λ^*, μ^*) is a rank-1 median algebra (variously known in the literature as a “tree algebra”, “median pretree” etc.). As already observed, (Λ^*, ρ^*) is an \mathbb{R}^* -tree, and $(\Lambda^\infty, \rho^\infty)$ is an \mathbb{R} -tree.

It is shown in [Bo2] that if Λ is a coarse median space of rank at most n , then there is no quasi-isometric embedding of \mathbb{R}^{n+1} into Λ (since this would give rise to an injective map of \mathbb{R}^{n+1} into Λ^∞ , contradicting the fact that Λ^∞ 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 6.6. *Let Λ be a coarse median space of rank at most n . Given parameters of Λ 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 Λ 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, $\phi_i : B_i \rightarrow \Lambda$. Now pass to the asymptotic cone with scaling factors, i . We end up with a continuous injective map, $\phi^\infty : B_1 \rightarrow \Lambda^\infty$, contradicting the fact that Λ^∞ has locally compact dimension at most n .

To see that only the parameters of Λ are relevant to the value of r , we should allow Λ also to vary among coarse median spaces with these parameters when taking the asymptotic cone. \square

Note that it follows, for example, that Λ admits no quasi-isometrically embedded euclidean half-space.

7. SURFACES

In this section, we discuss the coarse geometry of the mapping class group, $\text{Map}(\Sigma)$, of a compact orientable surface Σ , from the point view of medians. This uses technology introduced in [MasM2], and elaborated on by many authors since. Medians or “centroids” in the $\text{Map}(\Sigma)$ were introduced in [BehM2], see also Section 15 of this paper. We will express this in terms of the “marking graph”, $\mathcal{M}(\Sigma)$, which is quasi-isometric to $\text{Map}(\Sigma)$.

Let Σ be a compact orientable surface. Let $\xi(\Sigma)$ be the complexity of Σ , 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 $S_{g,p}$ for the “model” surface of genus g and p boundary components.

For $\xi(\Sigma) \geq 1$, let $\mathcal{G} = \mathcal{G}(\Sigma)$ be the curve graph of Σ . Its vertex set, $V(\mathcal{G}) = \mathcal{G}^0$, is the set of free homotopy classes of essential non-peripheral simple closed curves in Σ . Two curves, $\alpha, \beta \in \mathcal{G}^0$ are adjacent if $\iota(\alpha, \beta)$ is equal to 2 if Σ is an $S_{0,4}$, 1 if Σ is an $S_{1,1}$, or 0 if $\xi(\Sigma) \geq 2$ (the non-exceptional case). Here $\iota(\alpha, \beta)$ denotes the geometric intersection number.

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

Theorem 7.1. *There is a universal constant, k , such that for any compact surface, Σ , $\mathcal{G}(\Sigma)$ is k -hyperbolic.*

The existence of such a k , depending on $\xi(\Sigma)$ was proven by Masur and Minsky [MasM1]. The fact that it is uniform (independent of $\xi(\Sigma)$) was proven independently in [Ao, Bo3, CIRS, HePW]. (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. One can therefore simply assert dependence of constants on $\xi(\Sigma)$ at the relevant points.)

Given non-empty $a, b \subseteq \mathcal{G}^0$, let $\iota(a, b) = \max\{\iota(\alpha, \beta) \mid \alpha \in a, \beta \in b\}$. We write $\iota(a) = \iota(a, a)$. If $\iota(a) = 0$, we refer to a as a *multicurve*. We say that $a \subseteq \mathcal{G}^0$ *fills* Σ if $\iota(a, \gamma) \neq 0$ for all $\gamma \in \mathcal{G}^0$.

Given $p, q \in \mathbb{N}$, define a graph $\mathcal{M} = \mathcal{M}(\Sigma, p, q)$ by taking the vertex set, $V(\mathcal{M}) = \mathcal{M}^0$ to be the set of $a \subseteq \mathcal{G}^0$ such that a fills Σ and where $\iota(a) \leq p$, and by deeming $a, b \in \mathcal{M}^0$ 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 non-empty and connected. For definiteness, we can set $\mathcal{M}(\Sigma) = \mathcal{M}(\Sigma, 2, 4)$, though the actual choice will not matter. (The inclusion of $\mathcal{M}(\Sigma, 2, 4)$

into larger $\mathcal{M}(\Sigma, p, q)$ is a quasi-isometry.) We refer to $\mathcal{M}(\Sigma)$ as the *marking graph* of Σ (cf. the marking complex defined in [MasM2]).

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

Definition. By a *subsurface realised in Σ* 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 non-peripheral 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 non-peripheral annuli.

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

We will sometimes abuse notation and use the same symbol for a subsurface and some realisation of it in Σ .

We write $\mathcal{X} = \mathcal{X}(\Sigma)$ for the set of subsurfaces of Σ . We write $\mathcal{X} = \mathcal{X}_A \sqcup \mathcal{X}_N$ where \mathcal{X}_A and \mathcal{X}_N are respectively the sets of annular and non-annular subsurfaces. Given $\gamma \in \mathcal{G}^0$, we will write $X(\gamma)$ for the annular neighbourhood of γ in Σ . This gives a natural bijective correspondence between \mathcal{X}_A and \mathcal{G}^0 .

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

- (1) $X = Y$: X and Y are homotopic.
- (2) $X \prec Y$: $X \neq Y$, and X can be homotoped into Y but not into ∂Y .
- (3) $Y \prec X$: $Y \neq X$, 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 $X \prec Y$ or $X = Y$.

If $X \in \mathcal{X}_N$, we can define $\mathcal{G}(X)$ and $\mathcal{M}(X)$ intrinsically as above. If $X \in \mathcal{X}_A$, one needs to define \mathcal{G} as an arc complex in the annular cover of Σ corresponding to X (see [MasM2]). This is quasi-isometric to the real line. In this case, we set $\mathcal{M}(X) = \mathcal{G}(X)$. (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 $\mathcal{G}(\gamma) = \mathcal{G}(X)$ and $\mathcal{M}(\gamma) = \mathcal{M}(X)$, when $\gamma \in \mathcal{G}^0$, where X is the annular neighbourhood.

We will write σ_X and ρ_X respectively for the combinatorial metrics on $\mathcal{G}(X)$ and $\mathcal{M}(X)$.

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 $x \sim y$. The bound can be explicitly determined by following through the steps of the argument. Eventually these bounds will normally depend only on $\xi(\Sigma)$. We will make explicit any other dependence. 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 normally only interested in maps defined up to bounded distance. For a graph it is therefore 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 are only interested in the metric defined up to bilipschitz equivalence.

The maps defined in this section will involve choices, though they are all canonical up to bounded distance.

We have various natural maps between curve graphs and marking graphs, which are generally termed ‘‘subsurface projections’’. These were originally defined for curve graphs in [MasM2]. All the maps discussed below are uniformly coarsely lipschitz.

Given $X \in \mathcal{X}$ we have a $\chi_X : \mathcal{M}(X) \rightarrow \mathcal{G}(X)$. If $X \in \mathcal{X}_A$, this is the identity. If $X \in \mathcal{X}_N$, it just chooses some curve from the marking.

Given $X, Y \in \mathcal{X}$ with $X \preceq Y$, we can define a subsurface projection, $\psi_{XY} : \mathcal{M}(Y) \rightarrow \mathcal{M}(X)$. If $X \preceq Y \preceq Z$, then $\psi_{XY} \circ \psi_{YZ} \sim \psi_{XZ}$. Write $\theta_X = \chi_X \circ \psi_{XY} : \mathcal{M}(Y) \rightarrow \mathcal{G}(X)$. When the domain is clear from context, we will abbreviate $\psi_X = \psi_{XY}$ and $\theta_X = \theta_{XY}$.

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

In [BehM2], the authors defined a ‘‘centroid’’, $\mu_\Sigma(a, b, c) \in \mathcal{M}^0(\Sigma)$ of three markings, $a, b, c \in \mathcal{M}^0(\Sigma)$. In [Bo1] it was shown that the map μ_Σ is a coarse median on $\mathcal{M}(\Sigma)$. Moreover, the maps χ_X and ψ_{XY} , and hence θ_{XY} are all uniform quasimorphisms. Indeed the fact

that the maps $\theta_X : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(X)$ are all uniform quasimorphisms characterises the median map μ_Σ . (Note that the curve graphs $\mathcal{G}(X)$ are hyperbolic.) A direct proof of the existence of medians (broadly following [BehM2, BehKMM]) is given in Section 15.

Given $a, b \in \mathcal{M}^0(X)$ and $X \in \mathcal{X}$, we will abbreviate $\sigma_X(a, b) = \sigma_X(\theta_X a, \theta_X b)$ and $\rho_X(a, b) = \rho_X(\psi_X a, \psi_X b)$. Given $r \geq 0$, we write $\mathcal{A}(a, b; r) = \{X \in \mathcal{X} \mid \sigma_X(a, b) > r\}$.

Proposition 7.2. (1) *There is some universal $r_0 \geq 0$, such that if $a, b \in \mathcal{M}^0(\Sigma)$, then $\mathcal{A}(a, b; r_0)$ is finite.*

(2) *Given $r \geq 0$, there is some $r' \geq 0$, depending only on r and $\mathcal{A}(a, b; r) = \emptyset$, then $\rho(a, b) \leq r'$.*

At least if one admits dependence on $\xi(\Sigma)$ in (1), then both the above statements are immediate consequences of the following distance estimate given in [MasM2]. Given $a, b \in \mathcal{M}^0(\Sigma)$ and $r \geq 0$, let $D(a, b; r) = \sum_{X \in \mathcal{A}(a, b; r)} \sigma_X(a, b)$. The following is due to Masur and Minsky [MasM2]:

Theorem 7.3. [MasM2]. *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 $a, b \in \mathcal{M}^0(\Sigma)$, then $k_1 \rho_\Sigma(a, b) - h_1 \leq D(a, b; r) \leq k_2 \rho_\Sigma(a, b) + h_2$.*

In fact, it can be seen from the argument in [MasM2] that the constant r_0 depends only on the constants of hyperbolicity of the curve graphs, and the constant of the Bounded Geodesic Image theorem, and can therefore be made independent of $\xi(\Sigma)$ (see the discussion of Theorem 7.1 and Lemma 7.5 here). This therefore gives the uniform constants in Proposition 7.2. (Though this is not central to our discussion.)

We will sometimes abbreviate this to $D(a, b; r) \asymp \rho_\Sigma(a, b)$. In other words, these two quantities agree up to certain linear bounds which are implicit.

The following is an immediate corollary:

Corollary 7.4. *There is some $r_0 \geq 0$ depending only on $\mathcal{X}_I(\Sigma)$ 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 $\mathcal{X}_I(\Sigma)$ such that the following holds. Suppose $a, b \in \mathcal{M}^0(\Sigma)$ and $Y \in X$ and that $X \preceq Y$ whenever $X \in \mathcal{A}(a, b; r)$. Then $k_1 \rho_\Sigma(a, b) - h_1 \leq \rho_Y(a, b) \leq k_2 \rho_\Sigma(a, b) + h_2$.*

More generally, suppose $Y_1, \dots, Y_n \in \mathcal{X}$ and that $X \in \mathcal{A}(a, b; r)$ implies $X \preceq Y_i$ for some i . Then it follows that $\sum_{i=1}^n \rho_{Y_i}(a, b) \asymp \rho_\Sigma(a, b)$.

The following is a simple consequence of the Bounded Geodesic Image Theorem of Masur and Minsky. We will express it in terms of the Gromov product: given $\alpha, \beta, \gamma \in \mathcal{G}(X)$, write $\langle \alpha, \beta : \gamma \rangle_X = \frac{1}{2}(\sigma_X(\alpha, \gamma) + \sigma_X(\alpha, \gamma) - \sigma_X(\beta, \gamma))$.

Lemma 7.5. *There is some universal $r_0 \geq 0$ such that if $X, Y, Z \in \mathcal{X}$, and $\langle \theta_\Sigma X, \theta_\Sigma Y : \theta_\Sigma Z \rangle > r_0$ (so that $X \pitchfork Z$ and $Y \pitchfork Z$), then $\sigma_Z(\theta_Z X, \theta_Z Y) \leq r_0$.*

This follows from the fact 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 . The Bounded Geodesic Image Theorem (with dependence on $\xi(\Sigma)$) was proven in [MasM2]. A simpler proof, with uniform constants, is given in [We].

The following result of Behrstock is also central to the theory:

Lemma 7.6. *There is some universal r_0 such that if $X, Y \in \mathcal{X}$ and $\gamma \in \mathcal{G}^0(\Sigma)$ with $X \pitchfork Y$, $\gamma \pitchfork X$ and $\gamma \pitchfork Y$, then $\min\{\sigma_X(\theta_X(\gamma), \theta_X(Y)), \sigma_Y(\theta_Y(\gamma), \theta_Y(X))\} \leq r_0$.*

This is proven in [Beh] (with r_0 depending on $\xi(\Sigma)$). A simpler proof, which gives explicit universal constants can be found in [Man].

Let τ be a multicurve 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 complementary components which are not $S_{0,3}$'s. Let $\mathcal{X}(\tau) = \mathcal{X}_A(\tau) \cup \mathcal{X}_N(\tau)$. Let $\mathcal{T}(\tau)$ be the direct product of the graphs $\mathcal{M}(X)$ for $X \in \mathcal{X}(\tau)$. Note that $\mathcal{T}(\tau)$ is a coarse median space, with the median defined coordinatewise. We can also combine the maps $\psi_X : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(X)$ for $X \in \mathcal{X}(\tau)$ to give a uniformly Lipschitz quasimorphism, $\psi_\tau : \mathcal{M}(\Sigma) \rightarrow \mathcal{T}(\tau)$.

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

Lemma 7.7. *There are constants, $r_1, r_2 \geq 0$, depending only on $\xi(\Sigma)$, such that for any $p \in \mathcal{T}(\tau)$, there is some $a \in T(\tau; r_1)$ with $\psi_\tau(a) \sim_{r_2} p$.*

Proof. Take the given marking on each component of $\Sigma \setminus \tau$, together with a marking curve for each element of τ which projects to the appropriate element of the corresponding arc graph. (See the construction described in Section 15 for more details.) \square

We will fix some such r_1 , and abbreviate $T(\tau) = T(\tau; r_1)$.

Given $Y \in \mathcal{X}$, we write $Y \pitchfork \tau$ to mean that $\gamma \pitchfork Y$ or $\gamma \prec Y$ for some $\gamma \in \tau$. Let $\mathcal{X}_T(\tau) = \{Y \in \mathcal{X} \mid Y \pitchfork \tau\}$. Note that if $Y \in \mathcal{X}_T(\tau)$, then we can define $\theta_Y(\tau) = \theta_Y(\gamma)$ for any such γ . This is well defined up to bounded distance. In fact, we see that $\theta_Y(a) \sim \theta_Y(\tau)$ for all $a \in T(\tau)$.

Let $\mathcal{X}_I(\tau) = \mathcal{X} \setminus \mathcal{X}_T(\tau)$. Thus, $Y \in \mathcal{X}_I(\tau)$ 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 .) Note that $\theta_Y \circ \psi_X \sim \theta_Y$. We see that the maps θ_Y for $Y \in \mathcal{X}_I(\tau)$ are determined by the map ψ_τ defined above.

In other words, τ together with $\psi_\tau(a)$ determine all the subsurface projections $\theta_Y(a)$ for all $Y \in \mathcal{X}$ for any given $a \in \mathcal{M}^0(\Sigma)$. By Proposition 7.2, these projections determine a up to bounded distance in $\mathcal{M}(\Sigma)$. Therefore, given $p \in \mathcal{M}(\Sigma)$, the element $a \in \mathcal{M}^0(\Sigma)$ given by Lemma 7.7 is unique up to bounded distance. We therefore get a map, $v_\tau : \mathcal{T}(\tau) \rightarrow T(\tau) \subseteq \mathcal{M}(\Sigma)$ such that $\psi_\tau \circ v_\tau : \mathcal{T}(\tau) \rightarrow \mathcal{T}(\tau)$ is the identity up to bounded distance. Note that v_τ is a uniformly lipschitz quasimorphism.

This in turn gives rise to a quasimorphism, $\omega_\tau = v_\tau \circ \psi_\tau : \mathcal{M}(\Sigma) \rightarrow T(\tau)$. 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 $Y \in \mathcal{X}_I(\tau)$. As a consequence, we get:

Lemma 7.8. *Given $r \geq 0$, there is some r_2 depending only on r and $\xi(\Sigma)$, such that if $a \in \mathcal{M}^0(\Sigma)$, with $\sigma_Y(\theta_Y(a), \theta_Y(\tau)) \leq r$ for all $Y \in \mathcal{X}_T(\tau)$, then $\rho(a, T(\tau)) \leq r_2$.*

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 $\mathcal{X} = \mathcal{X}_I(\tau) \cup \mathcal{X}_T(\tau)$, part (2) of Proposition 7.2 tells us that $a \sim b$. \square

In other words, up to bounded distance, $T(\tau)$ can be described as the set of $a \in \mathcal{M}^0(\Sigma)$ such that $\theta_Y a \sim \theta_Y \tau$ for all $Y \in \mathcal{X}_T(\tau)$. Note also that the conclusion of Lemma 7.8 is the same as asserting that $a \in T(\tau; h)$ for some h depending only on r and $\xi(\Sigma)$.

Lemma 7.9. *$T(\tau)$ is uniformly quasiconvex in $\mathcal{M}(\Sigma)$.*

Proof. Suppose $a, b \in T(\tau)$ and $c \in \mathcal{M}^0(\Sigma)$. 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 7.8, $\mu_\Sigma(a, b, c)$ is a bounded distance from $T(\tau)$. \square

Lemma 7.10. *The map $\omega_\tau : \mathcal{M}(\Sigma) \rightarrow T(\tau)$ is a quasiprojection map.*

Proof. Let $x \in \mathcal{M}^0(X)$ and $c \in T(\tau)$. If $X \in \mathcal{X}_T(\tau)$, then $\omega_\tau x \in T(\tau)$, and $\theta_X \mu_\Sigma(x, \omega_\tau x, c) \sim \mu_X(\theta_X x, \theta_X \omega_\tau x, \theta_X c) \sim \mu_X(\theta_X x, \theta_X \tau, \theta_X c) \sim \theta_X \tau \sim \theta_X \omega_\tau x$. If $X \in \mathcal{X}_I(\tau)$, then $\theta_X \mu_\Sigma(x, \omega_\tau x, c) \sim \mu_X(\theta_X x, \theta_X \omega_\tau x, \theta_X c) \sim \mu_X(\theta_X x, \theta_X x, \theta_X c) \sim \theta_X x \sim \theta_X \omega_\tau x$. Thus, by part (2) of Proposition 7.2, we have $\mu_\Sigma(x, \omega_\tau x, c) \sim \omega_\tau x$, as required. \square

Note that a similar discussion applies to $T(\tau; r)$ for any $r \geq r_0$, except that the constants arising will depend on r as well as on $\xi(\Sigma)$. In

summary, we have shown that $T(\tau)$ is a quasi-isometrically embedded copy of $\mathcal{T}(\tau)$.

To continue, we need the following simple topological observation.

Lemma 7.11. *Suppose that $\mathcal{Y} \subset \mathcal{X}$ is any non-empty subset of \mathcal{X} . Then there is a unique subset $\tau \subseteq \mathcal{G}^0(X)$, which is either empty or a multicurve such that $\mathcal{Y} \subseteq \mathcal{X}_I(\tau)$ and if $X \in \mathcal{X}_T(\tau)$ there is some $Y \in \mathcal{Y}$ with $Y \pitchfork X$ or $Y \prec X$. (Here, we adopt the convention that $\mathcal{X}_I(\emptyset) = \emptyset$ and $\mathcal{X}_T(\emptyset) = \mathcal{X}$.)*

Proof. There is a unique smallest, possibly disconnected, subsurface of Σ which contains (up to homotopy) every element of \mathcal{Y} . We let τ be the collection of curves in Σ which are relative boundary components of this subsurface. It is easily verified that τ has the properties stated, and is unique with these properties. \square

We will write $\tau(\mathcal{Y})$ for the multicurve thus constructed.

Given $a, b \in \mathcal{M}^0(X)$, let $\tau(a, b; r) = \tau(\mathcal{A}(a, b; r))$.

Lemma 7.12. *There some r_3 depending only on $\xi(\Sigma)$, such that for all $r \geq r_3$ there is some r'_3 such that if $a, b \in \mathcal{M}^0(\Sigma)$ then $\rho(a, T(\tau)) \leq r'_3$ and $\rho(b, T(\tau)) \leq r'_3$, where $\tau = \tau(a, b; r)$. (Here, we adopt the convention that $\tau(\emptyset) = \mathcal{M}^0(\Sigma)$.)*

Proof. We verify the hypotheses of Lemma 7.8. Let $X \in \mathcal{X}_T(\tau)$. By the definition of τ , there is some $Y \in \mathcal{A}(a, b; r)$ such that either $Y \prec X$ or $Y \pitchfork X$. Thus, $\sigma_Y(\theta_Y a, \theta_Y b) \geq r \geq r_3$, and since $X \notin \mathcal{A}(a, b; r)$ we have $\sigma_X(\theta_X a, \theta_X b) \leq r$. Suppose that $Y \prec X$. If $r \geq r_0$, the constant of Lemma 7.5, then it follows that $\langle \theta_X a, \theta_X b; \theta_X Y \rangle \leq r_0$, and so, by the definition of Gromov product, $\sigma_X(\theta_X a, \theta_X Y) + \sigma_X(\theta_X b, \theta_X Y) \leq 2(r + r_0)$. (Though a more elementary argument shows that $\theta_X Y$ must be a bounded distance from $\theta_X a$ and $\theta_X b$ in this case.) Suppose that $Y \pitchfork X$. If r_3 is bigger than twice the constant r_0 of Lemma 7.6, then, without loss of generality, we must have $\sigma_Y(\theta_Y a, \theta_Y X) > r_0$, so by Lemma 7.6, we must have $\sigma_X(\theta_X a, \theta_X Y) \leq r_0$, and it follows also that $\sigma_X(\theta_X b, \theta_X Y) \leq r + r_0$. Thus, in all cases, we see that $\sigma_X(\theta_X a, \theta_X Y)$ and $\sigma_X(\theta_X b, \theta_X Y)$ are bounded.

Also, since τ is disjoint from Y , we have that $\sigma_X(\theta_X \tau, \theta_X Y)$ is bounded. In summary, this shows that, for all $X \in \mathcal{X}_T(\tau)$, $\sigma_X(\theta_X a, \theta_X \tau)$ and $\sigma_X(\theta_X b, \theta_X \tau)$ are both bounded in terms of r and $\xi(\Sigma)$. The statement now follows from Lemma 7.8 \square

In fact, the same statement holds if we set $\tau = \tau(\mathcal{Y})$ where \mathcal{Y} is any set with $\mathcal{A}(a, b; r) \subseteq \mathcal{Y} \subseteq \mathcal{A}(a, b; r)$.

We now move on to consider “quasicubes” in $\mathcal{M}(\Sigma)$.

Let $Q = \{-1, 1\}^n$, thought of as a median algebra. Given $i \in \{1, \dots, n\}$, let Q_i^\pm be the set of elements of Q whose i th coordinate is equal to ± 1 . Thus, Q_i^\pm is a codimension-1 face of Q , and $\{Q_i^-, Q_i^+\}$ is precisely the set of (unoriented) walls of Q .

Suppose that Λ is a hyperbolic space with the usual median. (That is, $\mu(a, b, c)$ is a bounded distance from each of the sides of a geodesic triangle with vertices a, b, c .) If $\phi : Q \rightarrow \Lambda$ is a quasimorphism, then we can find some $i \in \{1, \dots, n\}$ such that $\phi(x) \sim \phi(y)$ for all $x, y \in Q_i^-$ and for all $x, y \in Q_i^+$. Here, the constants depend only on the hyperbolicity constant and the quasimorphism constant. (For the purposes of our discussion, we could also allow it to depend on n .) This is an exercise in hyperbolic spaces, or can be deduced from the fact that hyperbolic spaces are coarse median of rank 1.

Applying this to the curve graphs, $\mathcal{G}(X)$, for $X \in \mathcal{X}$, we get:

Lemma 7.13. *Given $h \geq 0$, there is some r_4 depending only on h and $\xi(\Sigma)$ such that if $\phi : Q \rightarrow \mathcal{M}(\Sigma)$, is a h -quasimorphism and $X \in \mathcal{X}$, then there is some $i \in \{1, \dots, n\}$, such that $\sigma_X(\phi x, \phi y) \leq r_4$ for all $x, y \in Q_i^-$ and all $x, y \in Q_i^+$.*

Given $s \geq 0$ and $i \in \{1, \dots, n\}$, let $\mathcal{A}(\phi, i; s) = \{X \in \mathcal{X} \mid (\forall x \in Q_i^-, y \in Q_i^+) \sigma_X(\phi x, \phi y) > s\}$, and let $\mathcal{A}(\phi; s) = \bigcup_{i=1}^n \mathcal{A}(\phi, i; s)$. Note that if $s \geq r_4$, then by Lemma 7.13, we have $\sigma_X(\phi x, \phi y) \leq r_4$ for all $x, y \in Q_i^-$ and for all $x, y \in Q_i^+$, since this must be the unique i given by that lemma. In particular, we see that $\mathcal{A}(\phi, i; s) \cap \mathcal{A}(\phi, j; s) = \emptyset$ if $s \geq r_4$. Moreover, $X \in \mathcal{X} \setminus \mathcal{A}(\phi; s)$, then $\sigma_X(\phi x, \phi y) \leq s + 2r_4$ for all $x, y \in Q$. It follows that if $\mathcal{A}(\phi; s) = \emptyset$, then $\text{diam}(\phi(Q))$ is bounded above in terms of s and $\xi(\Sigma)$.

Lemma 7.14. *There is some $s_0 = s_0(h)$, depending only on h and $\xi(\Sigma)$, such that if $X \in \mathcal{A}(\phi, i; s_0)$ and $Y \in \mathcal{A}(\phi, j; s_0)$ with $i \neq j$, then $X \wedge Y$.*

Proof. Given $X \in \mathcal{X}$, and $x, y, z, w \in \mathcal{M}^0(X)$, write $(x, y : z, w)_X = \frac{1}{2} \max(\sigma_X(x, z) + \sigma_X(y, w), \sigma_X(x, w) + \sigma_X(y, z)) - (\sigma_X(x, y) + \sigma_X(z, w))$. Now Lemma 11.7 of [Bo1] states that there is some l_0 depending only on $\xi(\Sigma)$ such that if $a, b, c, d \in \mathcal{M}^0(\Sigma)$, and $X, Y \in \mathcal{X}$ with $(a, b; c, d)_X \geq l_0$ and $(a, c; b, d)_Y \geq l_0$, then $X \wedge Y$.

Set $s_0 \geq l_0 + r_4$. Suppose that $X \in \mathcal{A}(\phi, i; s_0)$ and $Y \in \mathcal{A}(\phi, j; s_0)$. Let $a \in Q_i^- \cap Q_j^-$, $b \in Q_i^+ \cap Q_j^-$, $c \in Q_i^- \cap Q_j^+$ and $d \in Q_i^+ \cap Q_j^+$. $(a, b; c, d)_X \geq l_0$ and $(a, c; b, d)_Y \geq l_0$ so $X \wedge Y$. \square

Now fix s sufficiently large depending on $\xi(\Sigma)$, and abbreviate $\mathcal{A}(\phi, i) = \mathcal{A}(\phi, i; s)$ and $\mathcal{A}(\phi) = \mathcal{A}(\phi; s) = \bigcup_{i=1}^n \mathcal{A}(\phi, i)$. In particular, we assume

that $s \geq \max\{s_0, r_3\}$ where s_0 and r_3 are the constants of Lemmas 7.12 and 7.14 respectively.

Let $\tau = \tau(\phi) = \tau(\mathcal{A}(\phi))$, as given by Lemma 7.11. By Lemma 7.12 we see that $\phi(Q) \subseteq N(T(\tau), r)$, where r depends only on h .

By construction, every element of $\mathcal{X}(\tau)$ is contained in a unique element of $\mathcal{X}(\tau) = \mathcal{X}_N(\tau) \cup \mathcal{X}_I(\tau)$. Moreover, by Lemma 7.14, we see that if $X \in \mathcal{A}(\phi, i)$ and $Y \in \mathcal{A}(\phi, j)$ are contained in the same element of $\mathcal{X}(\tau)$, then $i = j$. In this way we can partition $\mathcal{X}(\tau)$ as $\mathcal{X}(\tau) = \bigsqcup_{i=0}^n \mathcal{B}(\phi, i)$ where each element of $\mathcal{B}(\phi, 0)$ contains no element of $\mathcal{A}(\phi)$, and where each element of $\mathcal{B}(\phi, i)$ contains some element of $\mathcal{A}(\phi, i)$, but no element of $\mathcal{A}(\phi, j)$ for $j \neq i$.

We remark that a simple consequence of Theorem 7.3 is that if a, b are opposite corners of Q , then $\rho(\phi a, \phi b) \asymp \sum_{Y \in \mathcal{B}(\phi)} \rho_Y(\phi a, \phi b)$. (That is, $\rho(\phi a, \phi b)$ agrees with $\sum_{Y \in \mathcal{B}(\phi)} \rho_Y(\phi a, \phi b)$ up to linear bounds depending on n, h and $\xi(\Sigma)$.)

Similarly, if c, d is an edge of Q crossing the i th wall, we have $\rho(\phi c, \phi d) \asymp \sum_{Y \in \mathcal{A}(\phi, i)} \rho_Y(\phi c, \phi d)$.

We say that ϕ is *non-degenerate* if $\mathcal{A}(\phi, i) \neq \emptyset$ for all $i \in \{1, \dots, n\}$. In this case, $n \leq |\mathcal{B}(\phi)| \leq \xi(\Sigma)$. Of particular interest is the case where $n = \xi(\Sigma)$. In this case, $\mathcal{B}(\phi, 0) = \emptyset$, and for all $i \geq 0$, $\mathcal{B}(\phi, i)$ consists of a single element $Y(\phi, i)$ which is an annulus, $S_{0,4}$ or $S_{1,1}$.

Definition. We say that a multicurve, τ , is *big* if each component of the complement is an $S_{0,3}$, $S_{0,4}$ or $S_{1,1}$.

We have shown:

Lemma 7.15. *Let Q be a ξ -cube, and let $\phi : Q \rightarrow \mathcal{M}(\Sigma)$ be a non-degenerate h -quasimorphism. Then there is a canonically associated big multicurve, $\tau = \tau(\phi)$, with the following properties. We have that $\phi(Q)$ lies in a bounded neighbourhood of $T(\tau)$. To each $i \in \{1, \dots, \xi\}$ we can associate either a component, γ_i , of τ or an element, Y_i , of $\mathcal{X}_N(\tau)$ (that is an $S_{1,1}$ or $S_{0,4}$ component of the complement of τ). If $X \in \mathcal{A}(\phi, i)$, then X is either an annular neighbourhood of γ_i , or a subsurface of Y_i , accordingly. If c, d is an edge of Q crossing the i th wall, then $\rho(\phi c, \phi d) \asymp \sigma_{\gamma_i}(\phi c, \phi d)$ or $\rho(\phi c, \phi d) \asymp \rho_{Y_i}(\phi c, \phi d)$. Here all constants depend only on h and the constants of the definitions.*

Note that, by the pigeon-hole principle, each component of τ has the form γ_i for some unique i , and each element of $\mathcal{U}\mathcal{X}_N(\tau)$ has the form Y_i for some unique i .

Suppose that τ is a complete multicurve. Note that, in this case, $\mathcal{X} = \mathcal{X}_A(\tau) \sqcup \mathcal{X}_T(\tau)$. We see that $\mathcal{T}(\tau)$ is quasi-isometric to \mathbb{R}^ξ , and

that $T(\tau)$ is (up to bounded distance) a uniformly quasi-isometrically embedded copy of \mathbb{R}^ξ .

Let $G(\tau) \leq \text{Map}(\Sigma)$ be the group generated by Dehn twists about the elements of τ . Thus $G(\tau) \cong \mathbb{Z}^\xi$. Clearly, $G(\tau)$ preserves $T(\tau)$ setwise. Note that, for all $X \in \mathcal{X}_T(\tau)$, $\sigma_X(a, b)$ is uniformly bounded (since $\theta_X a \sim \theta_X \tau \sim \theta_X b$). Moreover, we can find $g \in G(\tau)$ such that $\sigma_\gamma(b, ga)$ is uniformly bounded above. This shows:

Lemma 7.16. *There is some uniform $k \geq 0$ such that if τ is a complete multicurve in Σ and $a \in T(\tau)$, then $T(\tau) \subseteq N(G(\tau)a; k)$.*

We refer to a set of the form $T(\tau)$ for a complete multicurve, τ , as a “coarse Dehn twist flat”.

Note that, in fact, since there are only finitely many possibilities of γ up to the action $\text{Map}(\Sigma)$, we see that the maps $[g \mapsto ga] : G(\tau); \rightarrow T(\tau) \subseteq \mathcal{M}(\Sigma)$ are uniform quasi-isometric embeddings (which constants depending only on ξ). Here we take the word metric on $G(\tau)$ with respect to the Dehn twist generators. (This was shown in [FLM]).

Note that an element of $T(\tau)$ is determined up to bounded distance by its projections to $\mathcal{G}(\gamma)$ for $\gamma \in \tau$.

Suppose now that τ' is another complete multicurve, and that $\tau \cap \tau' = \emptyset$ (i.e. they have no curve in common). Note that there is some $c \in T(\tau)$, unique up to bounded distance, with $\theta_\gamma c \sim \theta_\gamma \tau'$ for all $\gamma \in \tau$. Thus, if $b \in T(\tau')$, we have $\theta_\gamma c \sim \theta_\gamma b$. It follows from the characterisation of ω_τ described above that $c \sim \omega_\tau b$ (since such γ account for all of $\mathcal{X}_I(\tau)$ in this case). This shows that the diameter of $\omega_\tau(T(\tau'))$ is uniformly bounded above. Since ω_τ is a quasiprojection, we see that $\mu(a, b, c) \sim c$ for all $a \in T(\tau)$ and all $b \in T(\tau')$. In fact, if $d \in \omega_{\tau'}(T(\tau))$, we also have $\mu(a, b, d) \sim d$, $\mu(a, d, c) \sim c$ and $\mu(b, c, d) \sim d$. It follows that for any $X \in \mathcal{X}$, we have $\sigma_X(a, b) \sim \sigma_X(a, c) + \sigma_X(c, d) + \sigma_X(d, b)$. Using the distance formula of [MasM2] (stated as Theorem 9.3 here), we see that $\rho(a, b)$ agrees with $\rho(a, c) + \rho(c, d) + \rho(d, b)$ up to uniform linear bounds.

Here we just need to extract the following information from this, which will be used in Section 10.

Lemma 7.17. *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 $\rho(x, T(\tau')) \geq r$ and $\rho(x, y) \leq kr + t$.*

Proof. If $\tau \cap \tau' = \emptyset$, this follows easily by considering a uniform quasi-geodesic ray in $T(\tau')$ with basepoint in $\omega_{\tau'}(T(\tau))$ and containing x .

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) an copies of $\mathcal{M}(X)$ as X ranges over the elements of $\mathcal{X}_N(\tau_0)$. Applying the above to the restrictions of τ and τ' to any such X we deduce the general case. \square

Of course one can say a lot more about the way in which coarse Dehn twist flats intersect. (See [BehKMM] and Section 11 here.)

8. ASYMPTOTIC CONES OF THE MARKING GRAPH

Let \mathcal{Z} be a countable set with a non-principal ultrafilter, as in Section 5. Let $\mathcal{UG} = \mathcal{UG}(\Sigma)$ and $\mathcal{UM} = \mathcal{UM}(\Sigma)$ be the ultrapowers of $\mathcal{G}(\Sigma)$ and $\mathcal{M}(\Sigma)$. These are graphs with vertex sets \mathcal{UG}^0 and \mathcal{UM}^0 , respectively. Note that the intersection number, ι , extends to a map $\mathcal{U}\iota : (\mathcal{UG}^0)^2 \rightarrow \mathcal{UN}$, and $\mathcal{U}\iota : (\mathcal{UM}^0)^2 \rightarrow \mathcal{UN}$. We also have an ultrapower, $\mathcal{UX} = \mathcal{UX}_A \sqcup \mathcal{UX}_N$. There is a natural bijection between \mathcal{UX}_A and \mathcal{UG}^0 .

We can extend the notation introduced in Section 5. For example, if $X, Y \in \mathcal{UX}$, 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 each of \mathcal{UG} , \mathcal{UM} and \mathcal{UX} with finite quotient. To simplify terminology, we refer to an element of \mathcal{UG}^0 as a *curve* and an element of $\mathcal{G}^0 \subseteq \mathcal{UG}^0$ as *standard curve*. We similarly refer to “markings”, “standard markings”, “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 8.1. *Suppose that $a \subseteq \mathcal{UG}^0$ and $\mathcal{U}\iota(a) \in \mathbb{N}$, then there is some $g \in \mathcal{U}\text{Map}(\Sigma)$ with $ga \subseteq \mathcal{G}^0$.*

Proof. By hypothesis, $\iota(a_\zeta)$ is almost always constant. Therefore, we can find $g_\zeta \in \text{Map}(\Sigma)$ such that $g_\zeta a_\zeta \subseteq \mathcal{G}^0(\Sigma)$ lies in one of only finitely many possible subsets of $\mathcal{G}^0(\Sigma)$. 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 a $S_{1,1}$ or $S_{0,4}$ (up to the action of $\mathcal{U}\text{Map}(\Sigma)$). We can also refer to boundary curves of a surface, or 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 7.

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.

Suppose that $t \in \mathcal{UR}$ is a positive infinitesimal. Rescaling, as in Section 7, 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 write ρ^* , σ^* , respectively, for the limiting metrics. Thus, (\mathcal{M}^*, ρ^*) and $(\mathcal{G}^*, \sigma^*)$ 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. In fact, $\mathcal{G}^\infty(\Sigma)$ is the *universal \mathbb{R} -tree*, that is the unique complete \mathbb{R} -tree with every vertex of valence 2^{\aleph_0} . The maps $\chi_\zeta : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$ are uniformly coarsely lipschitz, and so give rise to a lipschitz map $\chi^* : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$.

Moreover, the coarse median μ on $\mathcal{M}^*(\Sigma)$ gives rise to a ternary operation, $\mu^* : \mathcal{M}^* \rightarrow \mathcal{G}^*$. In fact, $(\mathcal{M}^*, \rho^*, \mu^*)$ is a topological median algebra of rank equal to $\xi = \xi(\Sigma)$. The same applies restricting to the asymptotic cone, $(\mu^\infty, \rho^\infty, \mu^\infty)$. It was shown in [BehDS] that ρ^∞ is bilipschitz equivalent to a median metric, ρ_M^∞ , inducing the median ρ^∞ (see also [Bo4]). This construction is not canonical. In what follows, ρ_M^∞ could be interpreted as any median metric bilipschitz equivalent to ρ^∞ , and inducing μ^∞ . Note that, by Theorem 2.1, ρ_M^∞ is in turn, bilipschitz equivalent to a CAT(0) metric. The latter construction is canonical, given ρ_M^* .

Note that in the case when Σ is an $S_{1,1}$ or $S_{0,4}$, $\mathcal{M}(\Sigma)$ is a Farey graph, and so $\mathcal{M}^*(\Sigma)$ is an \mathbb{R}^* -tree. In this case, the metric, ρ^∞ , on \mathcal{M}^∞ is already a median metric, so we can take $\rho_M^\infty(\Sigma) = \rho^\infty$. This again gives a universal \mathbb{R} -tree. The preimage of every point under the map $\chi^\infty : \mathcal{M}^\infty(\Sigma) \rightarrow \mathcal{G}^\infty(\Sigma)$ is also a universal \mathbb{R} -tree.

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 also have limiting lipschitz homomorphism $\psi_X^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{M}^*(X)$. In fact, as observed above, up to the action of $\mathcal{U}\text{Map}(\Sigma)$, we can 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, $\psi_{XY}^* : \mathcal{M}^*(X) \rightarrow \mathcal{M}^*(Y)$, with $\psi_{XY}^* \circ \psi_Y^* = \psi_X^*$. 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 $Y \preceq X$ 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 $\mathcal{M}^*(\gamma) = \mathcal{G}^*(\gamma) = \mathcal{M}^*(X)$ where X is an annular neighbourhood of γ . This is isometric to \mathbb{R}^* .

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 $\mathcal{M}^*(\Sigma)$, and from Lemma 7.9, we see that it is convex. (Note that it is also the limit of 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 8.1, 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 7, we defined maps $\psi_{\tau_\zeta} : \mathcal{M}(\Sigma) \rightarrow \mathcal{T}(\tau_\zeta)$, $v_{\tau_\zeta} : \mathcal{T}(\tau_\zeta) \rightarrow T(\tau_\zeta)$ and $\omega_{\tau_\zeta} = v_{\tau_\zeta} \circ \psi_{\tau_\zeta} : \mathcal{M}(\Sigma) \rightarrow T(\tau_\zeta)$. These are all uniformly coarsely lipschitz quasimorphisms, and so give rise to maps, $\psi_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{T}^*(\tau)$, $v_\tau^* : \mathcal{T}^*(\tau) \rightarrow T^*(\tau)$ and $\omega_\tau^* = v_\tau^* \circ \psi_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow T^*(\tau)$. In fact, from Lemma 7.10, we see that $\omega_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow T^*(\tau)$ is the gate map of $\mathcal{M}^*(\Sigma)$ to $T^*(\tau)$. Note also that if $\gamma \in \mathcal{UG}^0(\Sigma)$ with $\tau \pitchfork \gamma$ and $a \in T^*(\tau)$, then $\theta_\gamma^*(a) = \theta_\tau^*(a)$.

A special case is when τ is a complete multicurve (or pants decomposition). 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 a $S_{0,3}$, $S_{0,4}$, or a $S_{1,1}$), then again $\mathcal{UX}(\tau)$ has ξ elements, and $T^*(\tau)$ is a direct product of ξ \mathbb{R}^* -trees.

We restrict the above construction to the asymptotic cone, $\mathcal{M}^\infty(\Sigma) \subseteq \mathcal{M}^*(\Sigma)$. If $X \in \mathcal{UX}$, then $\mathcal{M}^*(X)$ and $\mathcal{G}^*(X)$ have preferred base-points. This is defined as follows. Fix any standard $a \in \mathcal{M}(\Sigma)$ and let $e_X \in \mathcal{M}^*(X)$ be the limit of the points $\psi_{X_\zeta}(a) \in \mathcal{M}^*(X_\zeta)$. This limit is independent of a . We similarly define $f_X \in \mathcal{G}^*(\Sigma)$ as the limit of $\theta_{X_\zeta}(a)$ (or equivalently, as $f_X = \chi^* e_X$). Let $\mathcal{M}^\infty(X)$ and $\mathcal{G}^\infty(X)$ be the components containing e_X and f_X respectively. By Lemma 8.1,

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 $\theta_X^*(\mathcal{M}^\infty(\Sigma)) \subseteq \mathcal{G}^\infty(\Sigma)$. We will denote the restriction of θ_X^* to \mathcal{M}^∞ by θ_X^∞ .

If τ is a complete multicurve, we write $T^\infty(\tau) = T^*(\tau) \cap \mathcal{M}^\infty(\Sigma)$. 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 $\mathcal{M}^\infty(X)$ for $X \in \mathcal{UX}_T$. In particular it is a flat in the sense defined in Section 2.

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

By Lemma 8.1, up to the action of \mathcal{UG}^0 , we can take τ to be standard. One way to construct $T^\infty(\tau)$ in this case is as follows. Recall that $G(\tau) \cong \mathbb{Z}^\xi$ be the subgroup of $\text{Map}(\Sigma)$ generated by Dehn twists about the components of τ . Let a be any element of $\mathcal{M}(\Sigma)$. 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 $\mathcal{M}^\infty(\Sigma)$. 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 $T^\infty(\tau) = T^*(\tau) \cap \mathcal{M}^\infty(\Sigma)$ 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 flat in $\mathcal{M}^\infty(\Sigma)$ 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, $\mathcal{G}(\Sigma)$ is a Farey graph, and (up to quasi-isometry) $\mathcal{M}(\Sigma)$ 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 $\mathcal{G}(\Sigma)$). To each $\gamma \in \mathcal{G}^0(\Sigma)$ we can associate a bi-infinite geodesic, or *axis*, in $\mathcal{M}(\Sigma)$. Up to bounded Hausdorff distance, we can identify this axis with the space $T(\gamma)$ defined as before, which we can, in turn, identify up to quasi-isometry, with $\mathcal{M}(\gamma) = \mathcal{G}(\gamma)$. Any two distinct axes meet in most a single edge of $\mathcal{M}(\Sigma)$.

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

In this case, $\mathcal{M}^\infty(\Sigma)$ and $\mathcal{G}(\Sigma)$ are universal \mathbb{R} -trees. Any closed subset median isomorphic to the real line is necessarily convex, and isometric to \mathbb{R} .

We now return to the general case. The main aim of the following discussion will be to describe maximal dimensional cubes and flats in \mathcal{M}^* or in \mathcal{M}^∞ .

Suppose that $a, b \in \mathcal{M}^*(\Sigma)$. Let $A(a, b) = \{X \in \mathcal{UX} \mid \theta_X^* a \neq \theta_X^* b\}$ and $B(a, b) = \{X \in \mathcal{UX} \mid \psi_X^* a \neq \psi_X^* b\}$. We will write $A^0(a, b) = \{\gamma \in \mathcal{UG}^0 \mid \theta_\gamma^* a \neq \theta_\gamma^* b\}$, so we can think of $A^0(a, b)$ as a subset of $A(a, b)$ or of $B(a, b)$.

Note that $A(a, b) \subseteq B(a, b)$, and if $X \in B(a, b)$ and $Y \in \mathcal{UX}$ with $X \preceq Y$, then $Y \in B(a, b)$. We also note that if a', b' is parallel to a, b , then $A(a', b') = B(a, b)$ and $B(a', b') = B(a, b)$. In particular, if $Q \equiv \{-1, 1\}^n \subseteq \mathcal{M}^*(\Sigma)$ is an n -cube, we can define $A_i(Q) = A(c, d)$ and $B_i(Q) = B(c, d)$, where $\{c, d\}$ is a side of Q crossing the i th wall (that is, c, d differ in precisely the i th coordinate).

Lemma 8.2. *Suppose that $a, b, c, d \in \mathcal{M}^*(\Sigma)$, $X, Y \in \mathcal{UX}$, $c \in [a, d]$, $b \in [a, c]$, $X \in A(a, b) \cap A(c, d)$ 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 Lemma 7.6. 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 Lemma 7.6, $\theta_{Y_\zeta} X_\zeta$ 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 $Y \notin A(b, c)$.

We also need to rule out the possibility that $X \prec Y$. If that were the case, we could derive a similar contradiction using the Bounded Geodesic Image Theorem of [MasM2] (see Lemma 7.5 here). Briefly, if $\sigma_{X_\zeta}(a_\zeta, b_\zeta)$ is large then 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 contradiction. (We omit details, since we do not need this case.) \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 $y \in [x, z]$. Note that any rank-1 subset of $\mathcal{M}^*(\Sigma)$ is monotone.

We write $C(O) = \bigcap \{A(x, y) \mid x, y \in O, x \neq y\}$. In other words, $C(O)$ 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 8.2.

Corollary 8.3. *If $O \subseteq \mathcal{M}^*(\Sigma)$ is monotone, $|O| \geq 4$, and $X, Y \in C(O)$, then either $X = Y$, or $X \wedge Y$.*

(Again, what we really need is that X and Y cannot cross.)

We similarly define $C^0(O) = \bigcap \{A^0(x, y) \mid x, y \in O, x \neq y\} \subseteq \mathcal{UG}^0$. We can identify $C^0(O)$ as a subset of $C(O)$. Corollary 8.3 tells us that, if $|O| \geq 4$, then $C^0(O)$ is a multicurve.

Suppose now, that $Q \equiv \{-1, 1\}^n \subseteq \mathcal{M}^*(\Sigma)$ is an n -cube.

Lemma 8.4. *If $X \in A_i(Q)$ and $Y \in A_j(Q)$ with $i \neq j$, then $X \wedge Y$.*

Proof. Let $\phi_\zeta : Q \rightarrow X_\zeta$ be uniform quasimorphisms with $\phi_\zeta x \rightarrow x$ for all $x \in Q$, as given by Lemma 6.3. If c, d is a side of Q crossing the i th wall, then $\sigma_{X_\zeta}(c, d) \rightarrow \infty$, so $X_\zeta \in \mathcal{A}(\phi_\zeta, i)$ for almost all ζ . Similarly, $Y_\zeta \in \mathcal{A}(\phi_\zeta, j)$ for almost all ζ . Thus, by Lemma 7.14, $X_\zeta \wedge Y_\zeta$ for almost all ζ , so $X \wedge Y$. \square

Of course, the above lemmas apply equally well to the subsets $A^0(a, b)$ and $A_i^0(Q)$ etc. of $\mathcal{UG}^0(\Sigma)$.

By a similar argument, we also get:

Lemma 8.5. *Suppose that $\gamma \in A_i^0(Q)$, that $X \in B_j(Q)$ which is an $S_{1,1}$ or $S_{0,4}$, with $i \neq j$. Then $\gamma \wedge X$.*

Proof. Similarly as in the proof of Lemma 8.4. we get $\rho_{\gamma_\zeta}(c, d) \rightarrow \infty$, so $\gamma_\zeta \in A(\phi_\zeta, i)$ for almost all ζ . If e, f is a side crossing the j th wall, then $\rho_{X_\zeta}(e, f) \rightarrow \infty$. Therefore, by Proposition 7.2, we have $\sigma_{Y_\zeta}(e, f)$ arbitrarily large for some $Y_\zeta \preceq X_\zeta$, and so $Y_\zeta \in A(\phi_\zeta, j)$. Therefore, by Lemma 7.9, $\gamma_\zeta \wedge Y_\zeta$, and so $\gamma_\zeta \wedge X_\zeta$. It follows that $\gamma \wedge X$ as claimed. \square

Now, suppose that $n = \xi$. Let $I^0(Q)$ be the set of $i \in \{1, \dots, \xi\}$ such that $|A_i^0(Q)| = 1$. For such i , write $A_i^0(Q) = \{\gamma_i(Q)\}$ with $\gamma_i = \gamma_i(Q) \in \mathcal{UG}^0(\Sigma)$. By Lemma 8.4, the γ_i are disjoint, so $\tau^0(Q) = \{\gamma_i \mid i \in I^0(Q)\}$ is a multicurve. We claim:

Lemma 8.6. *Let $Q \subseteq \mathcal{M}^*(\Sigma)$ be a ξ -cube. Then $\tau^0 = \tau^0(Q)$ is a big multicurve. Moreover, for each $i \in \{1, \dots, \xi\} \setminus I^0(Q)$, there is some $Y_i \in \mathcal{UX}_N(\tau^0)$ with $Y_i \in B_i(Q)$, and such that $Y_i \neq Y_j$ if $i \neq j$. Moreover, if $\gamma \in A_i^0(Q)$ where $i \notin I^0(Q)$, then $\gamma \prec Y_i$.*

Note that $\mathcal{UX}(\tau^0(Q)) = \mathcal{UX}_A(\tau^0(Q)) \sqcup \mathcal{UX}_N(\tau(Q))$ has exactly ξ elements, and $\tau^0(Q)$ can be identified with $\mathcal{UX}_A(\tau^0(Q))$. It follows that $\mathcal{UX}_N(\tau^0(Q))$ consists precisely of the surfaces Y_i for $i \notin I^0(Q)$.

Proof. Let $\phi_\zeta : Q \rightarrow \mathcal{M}(\Sigma)$ be as given by Lemma 7.9. Note that, by Lemma 7.15, if c, d is a side of Q crossing the i th wall, then $\rho_\zeta(c, d) \asymp \sum_{Y \in \mathcal{B}(\phi_\zeta, i)} \rho_{Y_\zeta}(c, d)$. Since $c \neq d$, $\rho_\zeta(c, d) \rightarrow \infty$, and so the right-hand side is also unbounded. In particular, it is non-zero for almost all ζ . In other words, $\mathcal{B}(\phi_\zeta, i) \neq \emptyset$, and so $\mathcal{A}(\phi_\zeta, i) \neq \emptyset$. Thus, ϕ_ζ is non-degenerate, and so $\mathcal{B}(\phi_\zeta, i)$ consists of a single element, which we denote by $Y_{\zeta, i}$. Now, $Y_{\zeta, i}$ is an annulus, $S_{1,1}$ or $S_{0,4}$. In other words, $\tau(\phi_\zeta)$ is big, and as above, we have $\rho_\zeta(c, d) \asymp \rho_{Y_{\zeta, i}}(c, d)$. Let τ^1 be the limit of the $\tau(\phi_\zeta)$, and let Y_i be the limit of the $Y_{\zeta, i}$. By the above, we see that $\rho_{Y_i}(c, d) \neq 0$, since $\rho(c, d) \neq 0$. In particular, $\psi_{Y_i}c \neq \psi_{Y_i}d$, so $Y_i \in B_i(Q)$. Let I^1 be the set of $i \in \{1, \dots, \xi\}$ for which Y_i is an annulus, with core, γ_i , say. That is, $\tau^1 = \{\gamma_i \mid i \in I^1\}$. Then $\{Y_i \mid i \notin I^1\}$ is precisely the set of $S_{1,1}$'s or $S_{0,4}$'s. components of the complement of τ .

By construction, $\gamma_i \in A_i^0(Q)$. We claim that in fact $A_i^0(Q) = \{\gamma_i\}$. For if $\beta \in A_i^0(Q) \setminus \{\gamma_i\}$, then by Lemma 7.6, β cannot cross any γ_j for $j \neq i$, and cannot lie in any Y_j for $j \notin I^1$. Thus, β must cross γ_i . But then, $\beta_\zeta \in \mathcal{A}(\phi_\zeta, i)$, and $\beta_\zeta \cap \gamma_{i, \zeta}$ for almost all ζ , contradicting the construction of τ^1 , and proving the claim. This shows that $I^1 \subseteq I^0(Q)$. In other words, $\tau^1 \subseteq \tau^0$ and so τ^0 is also big. We have also shown that for each $j \in \{1, \dots, \xi\} \setminus I^1$, there is some $Y_i \in \mathcal{UX}_N(\tau)$ with $Y_i \in A_i(Q)$, and such that $Y_i \neq Y_j$ if $i \neq j$. So certainly this also holds with I^1 replaced by $I^0(Q)$.

Finally, suppose $\gamma \in A_i^0(Q)$ with $i \notin I^0(Q)$. Then $i \notin I^1$, so we must have $\gamma \in Y_j$ for some $j \in \{1, \dots, \xi\}$. We claim that $j = i$. For suppose not. Then, $\gamma_\zeta \in \mathcal{A}(\phi_\zeta, i)$ for almost all ζ . Also, $Y_{j, \zeta} \in \mathcal{B}(\phi_\zeta, j)$, so there must be some $X_\zeta \preceq Y_{j, \zeta}$ with $X_\zeta \in \mathcal{A}(\phi_\zeta, j)$. But then, since X_ζ is an $S_{1,1}$ or $S_{0,4}$, either $\gamma_\zeta \preceq X_\zeta$, or $\gamma_\zeta \cap X_\zeta$, contradicting Lemma 7.14. \square

Note, in particular, that for any i , $A_i^0(Q)$ is empty, a single curve, or a set of curves filling an $S_{1,1}$ or an $S_{0,4}$ in Σ . (In the last case, $A_i^0(Q)$ might be infinite, or even uncountable, since we cannot assume it to be standard.)

We will use a variation on Lemma 8.5 as follows. Let Q be a ξ -cube, as before. Suppose that c, d is a side of Q . Then $[c, d]$ has rank 1. In a particular it is monotone. We can therefore define $C_i(Q) = C([c, d]) \subseteq \mathcal{UX}(\Sigma)$. In other words, it is $C_i(Q) = \{X \in \mathcal{UX} \mid \theta_X|_{[c, d]} \text{ is injective}\}$. Since any two such sides are parallel, this is well defined, independently

of the choice of c, d . We similarly define $C_i^0(Q) \subseteq \mathcal{UG}^0(\Sigma)$, which we can identify as a subset of $C_i^0(Q)$.

Now, clearly $C_i^0(Q) \subseteq A_i^0(Q)$, and so, by the above description of $A_i^0(Q)$, we see that $|C_i^0(Q)| \leq 1$. Let $I(Q) = \{i \mid C_i^0(Q) \neq \emptyset\}$. Thus, if $i \in I(Q)$, then $C_i^0(Q) = \{\gamma_i\}$. By Lemma 8.4, the γ_i are disjoint. Set $\tau = \tau(Q) = \{\gamma_i \mid i \in I(Q)\}$. We claim:

Proposition 8.7. *Let $Q \subseteq \mathcal{M}^*(\Sigma)$ be a ξ -cube. Then $\tau = \tau(Q)$ is a big multicurve, and $Q \subseteq T^*(\tau)$. If $Y \in \mathcal{UX}_N(\tau)$, then $Y \in C_i(Q)$ for some $i \in \{1, \dots, \xi\} \setminus I(Q)$.*

Again, by the pigeon-hole principle, it follows that if $i \in \{1, \dots, \xi\}$ then $Y_i \in C_i(Q)$ for a unique $Y_i \in \mathcal{UX}_N(\tau)$.

Before giving the proof, we make a couple of observations. First we claim, $I(Q) \subseteq I^0(Q)$. For suppose $i \in I(Q) \setminus I^0(Q)$. Then $|A_i^0(Q)| \geq 2$, so there must be some $\beta \in I^0(Q)$ with $\beta \pitchfork \gamma_i$. Let c, d be any side of Q crossing the i th wall. Then $\theta_\beta c \neq \theta_\beta d$ and so, by continuity of θ_β , we can find $c', d' \in [c, d]$ with $c < c' < d' < d$, and with $\theta_\beta c' \neq \theta_\beta d'$. But $\theta_{\gamma_i} c \neq \theta_{\gamma_i} c'$ and $\theta_{\gamma_i} d \neq \theta_{\gamma_i} d'$, so this contradicts Lemma 8.2. Thus $I(Q) \subseteq I^0(Q)$ as claimed, and so $\tau \subseteq \tau^0$.

We say that a ξ -cube, Q' , is *smaller* than Q , or that Q is *bigger* than Q' , if $Q' \subseteq \text{hull}(Q)$. Note that there is a natural bijective correspondence between the walls of Q' and those of Q , so we can index them by $\{1, \dots, \xi\}$ consistently. We clearly have $C_i(Q') \supseteq C_i(Q)$ and $A_i(Q') \subseteq A_i(Q)$ for all i .

Proof of Proposition 8.7. We will first show that τ is big. We will argue by “shrinking” the sides of Q indexed by $I^0(Q) \setminus I(Q)$. More precisely, if $i \in I^0(Q) \setminus I(Q)$, then $A_i^0(Q) = \{\gamma_i\}$ and $C_i^0(Q) = \emptyset$. If c, d is an i th side of Q , then we can find distinct $c', d' \in [c, d]$ with $A^0(c', d') = \emptyset$. We let Q' be the cube obtained by replacing the side c, d by c', d' for each such i , and leaving all other sides unchanged (up to the parallel relation). (Recall that $\text{hull}(Q)$ is the direct product of its 1-dimensional faces.) Thus, $A_i^0(Q') = \emptyset$ for all $i \in I^0(Q) \setminus I(Q)$, and $A_i^0(Q') = A_i^0(Q)$ for all other i . In particular, $I^0(Q') = I(Q)$, and so $\tau^0(Q') = \tau(Q)$. By Lemma 8.6 applied to Q' , it follows that this is a big multicurve, as claimed.

Suppose that $Y \in \mathcal{UX}_N(\tau)$. By Lemma 8.6 applied to Q' again, we have $Y \in B_j(Q')$ for some $j \in \{1, \dots, \xi\} \setminus I(Q)$. We claim that $Y \in C_j(Q)$. For if not, let e, f be a side of Q' crossing the j th wall. (This is also a side of Q up to parallelism.) There are distinct $e', f' \in [e, f]$ with $Y \notin B(e, f)$. We now shrink Q' to a smaller cube, Q'' , replacing the j th side by e', f' , and leaving all other sides alone. Thus, $Y \notin B_j(Q'')$.

Moreover, either $I^0(Q'') = I^0(Q')$ or $I^0(Q'') = I^0(Q') \cup \{j\}$. In the former case, $\tau^0(Q'') = \tau^0(Q')$, so, by Lemma 8.6 applied to Q'' , we have $Y \in B_k(Q'')$ for some $k \notin I(Q)$. But $B_k(Q'') \subseteq B_k(Q')$ and $Y \in B_j(Q')$, and so we must have $k = j$, giving a contradiction. Therefore, $I^0(Q'') = I^0(Q') \cup \{j\}$, and so $C_j(Q'') = \{\gamma_j\}$. Since $\gamma_j \in A_j(Q'') \subseteq A_j(Q')$, by Lemma 8.6 applied to Q' , we must have $\gamma_j \prec Y$, and so $Y \in B_j(Q'')$ again giving a contradiction. This shows that $Y \in C_j(Q)$ as claimed.

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. Therefore the map $\psi_\tau|_{\text{hull}(Q)} : \text{hull}(Q) \rightarrow T^*(\tau)$ is injective. Since this is a gate map, it follows that $\text{hull}(Q) \subseteq T^*(\tau)$ (for example using Corollary 3.10). \square

We have a similar result for flats. For this, we will restrict to $\mathcal{M}^\infty(\Sigma)$, since that is what is used in applications. Recall that, if $X \in \mathcal{UX}(\Sigma)$, then $\theta_X^\infty : \mathcal{M}^\infty(\Sigma) \rightarrow \mathcal{G}^\infty(X)$ and $\psi_X^\infty : \mathcal{M}^\infty(\Sigma) \rightarrow \mathcal{M}^\infty(X)$ are just the restrictions of θ_X^* and ψ_X^* .

Let $\Phi \subseteq \mathcal{M}^\infty(\Sigma)$ 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. (Any two such are parallel.) Let $C_i(\Phi) = C(L_i)$, 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 $C_i^0(\Phi) \subseteq \mathcal{UG}^0(\Sigma)$, which we can identify as a subset of $C_i(\Phi)$.

Note that if Q is any ξ -cube in Φ , then $C_i(Q) \supseteq C_i(\Phi)$. In fact, there is some ξ -cube $Q_0 \subseteq \Phi$, with $C_i(Q_0) = C_i(\Phi)$ for all i , and so $C_i(Q) = C_i(\Phi)$ for any cube in Φ bigger than Q_0 . In particular, $|C_i^0(\Phi)| \leq 1$. Let $I(\Phi) = \{i \mid C_i^0(\Phi) \neq \emptyset\}$. If $i \in I(\Phi)$, write $C_i^0(\Phi) = \{\gamma_i\}$, and let $\tau(Q) = \{\gamma_i \mid i \in I(\Phi)\}$. Thus, $\tau = \tau(\Phi) = \tau(Q_0)$ is a big multicurve. If Q is any bigger cube, then Proposition 8.7 tells us that $Q \subseteq T(\tau)$. Since $\text{hull}(Q)$ is exhausted by such hulls, we conclude:

Proposition 8.8. *If $\Phi \subseteq \mathcal{M}^\infty(\Sigma)$ is a flat, then $\tau(\Phi)$ is a big multicurve, and $\Phi \subseteq T(\tau(\Phi))$. Moreover, if $Y \in \mathcal{UX}_N(\tau(\Phi))$, then $Y \in C_i(\Phi)$, for some $i \in \{1, \dots, \xi\} \setminus I(\Phi)$.*

Note also, as in Proposition 8.7 that each $Y \in \mathcal{UX}(\tau)$ lies in $C_i(\Phi)$ for some unique $i \notin I(\Phi)$.

Note that, applying Lemma 8.6 to a large cube in Φ , we see that if $\gamma \in C_i^0(\Phi)$ and $X \in C_j(\Phi)$ is an $S_{1,1}$ or $S_{0,4}$, and $i \neq j$, then $\gamma \wedge X$.

Recall that $\tau(\Phi)$ is completely determined by Φ . It is precisely the set of curves γ for which some coordinate line L maps injectively (therefore bijectively) to $\mathcal{G}^*(\gamma) \cong \mathbb{R}$ via the map θ_γ^* . If it exists, $\gamma = \gamma(L)$ is completely determined by L , and if L' is parallel to L , then $\gamma(L) = \gamma(L')$.

Let $I(Q)$ be the set of coordinage directions for which such a curve, γ_i , exists, and set $\gamma = \{\gamma_i \mid i \in I(Q)\}$.

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

Lemma 8.9. *Let Φ_0, Φ_1 , be two flats with $\Phi_0 \cap \Phi_1$ a codimension-1 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 $\tau_i = \tau(\Phi_i)$. 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 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).

If $1 \in I_0 \cap I_1$, 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 $Y_0 \in C_i(\Phi_1)$ for some $i \neq 1$ (as observed after Proposition 8.7). But $C_i(\Phi_0) = C_i(\Phi_1)$. In other words, we have $\beta \prec Y_0$, $\beta \in C_1(\Phi_0)$, $Y_0 \in C_i(\Phi_0)$ and Y_0 is an $S_{1,1}$ or $S_{0,4}$, so this gives 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 X is an $S_{1,1}$ or $S_{0,4}$. 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, X 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, 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 $\mathcal{M}^\infty(Y)$.
- (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 $\mathcal{M}^\infty(Y)$, 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 project to the axes

in $\mathcal{M}^\infty(Y)$ corresponding to β_0 and β_1 . These axes intersect in a single point.

We next want to characterise Dehn twist flats.

Lemma 8.10. *Suppose that $\Phi \subseteq \mathcal{M}^\infty(\Sigma)$ is a flat. Suppose that for each i there is another flat, $\Phi_i \subseteq \mathcal{M}^\infty(\Sigma)$, with $\Phi \cap \Phi_i$ a codimension-1 coordinate subspace 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 $\gamma_i \in \tau$ be the corresponding curve. By the above, 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 $\tau \setminus \{\gamma_i\}$ 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 $\tau(\Phi) \setminus \{\gamma_i\}$. Since this holds for all $i \in I(\Phi)$ (that is for all components of $\tau(\Phi)$) it follows that $\tau(\Phi)$ 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 $\mathcal{U}^0G = \mathcal{U}G \cap \mathcal{U}^0\text{Map}(\Sigma)$. Then \mathcal{U}^0G acts transitively on Φ , preserving the coordinate directions.

Lemma 8.11. *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 $Y \in \mathcal{U}\mathcal{X}(\tau \setminus \{\gamma\})$ be the component containing γ . Let $\gamma \in \mathcal{G}^0(\Sigma)$ be any other standard curve in Y . Now the axes of β and γ in $\mathcal{G}^\infty(Y)$ 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 above, this is sufficient to prove the result. \square

Putting the above together with Lemma 4.5, we get:

Proposition 8.12. *Suppose that $\Phi \subseteq \mathcal{M}^\infty(\Sigma)$ 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, $\Psi \subseteq \mathcal{M}^\infty(\Sigma)$, 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 $\mathcal{M}^\infty(\Sigma)$, as shown in [BehKMM]. In fact, we only need an injective map. Moreover, we can take two different surfaces with the same complexity. In summary, we conclude:

Theorem 8.13. *Suppose that Σ and Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 2$. Suppose that we have a continuous injective map, $f : \mathcal{M}^\infty(\Sigma) \rightarrow \mathcal{M}^\infty(\Sigma')$ with closed image. If Φ is a Dehn twist flat in $\mathcal{M}^\infty(\Sigma)$, then $f(\Phi)$ is a Dehn twist flat in $\mathcal{M}^\infty(\Sigma')$.*

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

9. CONTROLLING HAUSDORFF DISTANCE

We begin a general statement, which generalises a construction of [BehKMM].

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 \mathbf{t} be a positive infinitesimal, and let M^* be the extended asymptotic cone determined by \mathbf{t} . Given $e \in M^*$, let M_e^* be the component of M^* containing e . Let $\mathbf{r} = 1/\mathbf{t}$.

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

The following is a simple observation (a similar statement is used in [BehKMM]).

Lemma 9.1. *Suppose that $\mathbf{A}, \mathbf{B} \in \mathcal{UP}(M)$, and $\mathbf{e} \in \mathcal{UA}$ (that is $e_\zeta \in A_\zeta$ for almost all ζ). Let $e \in M^*$ be the image of \mathbf{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 $A^* \cap M_e^* = B^* \cap M_e^*$, 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 9.2. *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 \mathbf{e}_n such that A^*, B^* are $\frac{1}{2n}$ -close on $N(\mathbf{e}_n; 2n)$. Write $\mathbf{e}_n = (e_{n,\zeta})_\zeta$. Let \mathcal{Z}_n be the set of $\zeta \in \mathcal{Z}$ such that A_ζ, B_ζ are $\frac{r_\zeta}{n}$ -close on $N(e_{n,\zeta}; nr_\zeta)$. 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 $n(\zeta) > n$ almost always). Let $e_\zeta = e_{n(\zeta),\zeta}$, and let e be the image of $\mathbf{e} = (e_\zeta)_\zeta$ in A^* . Now, for all n , A_ζ, B_ζ are almost always $\frac{r_\zeta}{n}$ -close on $N(e_\zeta; nr_\zeta)$, so A^*, B^* are $\frac{1}{n}$ -close on $N(e; n)$. Since this holds for all n , we have $A^* \cap M_e^* = B^* \cap M_e^* \neq \emptyset$, as required. \square

Suppose now that \mathcal{S} and \mathcal{T} are collections of subsets of M . We write \mathcal{UT} and \mathcal{US} for the respective ultrapowers. We suppose:

- (S1) S is (coarsely) connected for all $S \in \mathcal{S}$.
- (S2) If $\mathbf{T}, \mathbf{T}' \in \mathcal{UT}$ and there is some component, M^0 , of M^* such that $T^* \cap M^0 = (T')^* \cap M^0 \neq \emptyset$, then $\mathbf{T} = \mathbf{T}'$.
- (S3) For all $\mathbf{S} \in \mathcal{US}$, and for all components, M^0 , of M^* , there is some $\mathbf{T} \in \mathcal{UT}$ such that $S^* \cap M^0 = T^* \cap M^0$.

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

(In (S1), “coarsely connected” can be taken to mean that $N(S; s)$ is connected for some fixed s .)

Lemma 9.3. *If \mathcal{S}, \mathcal{T} satisfy (S1)–(S3) above, then there is some $k > 0$ such that for all $S \in \mathcal{S}$, there is some $T \in \mathcal{T}$, such that $\text{hd}(S, T) \leq k$.*

Proof. Suppose not. Let $\epsilon > 0$. Given any $\zeta \in \mathcal{Z}$, there is some $S_\zeta \in \mathcal{S}$ such that for all $T \in \mathcal{T}$, $\text{hd}(S_\zeta, T) > \epsilon r_\zeta$. Let $\mathbf{S} = (S_\zeta)_\zeta \in \mathcal{US}$. Let e_ζ be any element of S_ζ (so that $e \in S^*$). By (S3), there is some $\mathbf{T} \in \mathcal{UT}$ such that $S^* \cap M_e^* = T^* \cap M_e^*$. In particular, for all $R > 4\epsilon$, we have that S_ζ, T_ζ are almost always $\frac{\epsilon r_\zeta}{2}$ -close on $N(e_\zeta; 2Rr_\zeta)$. But $\text{hd}(S_\zeta, T_\zeta) > \epsilon r_\zeta$, so there is some $e'_\zeta \in S_\zeta$ such that S_ζ, T_ζ are not $\frac{\epsilon r_\zeta}{2}$ -close on $N(e'_\zeta; 2Rr_\zeta)$. By (S1), we can find $q_\zeta, q'_\zeta \in S_\zeta$ with $\rho(q_\zeta, q'_\zeta)$ bounded such that S_ζ, T_ζ are $\frac{\epsilon r_\zeta}{2}$ -close on $N(q_\zeta; 2Rr_\zeta)$ but not on $N(q'_\zeta; 2Rr_\zeta)$. But by (S3) again, there is almost always some $T'_\zeta \in \mathcal{T}$ such that S_ζ, T'_ζ are $\frac{\epsilon r_\zeta}{2}$ -close on $N(q'_\zeta; 2Rr_\zeta)$. Clearly $T'_\zeta \neq T_\zeta$. It follows that T_ζ, T'_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 $\mathbf{T}' = (T'_\zeta)_\zeta$. We see that $T^*, (T')^*$ are ϵ -close on $N(q; R)$. Since $R > 4\epsilon > 0$ were arbitrary, it follows from Lemma 9.2 that there is some component, M^0 , of M^* such that

$T^* \cap M^0 = (T')^* \cap M^0 \neq \emptyset$. By (S2), we have $\mathbf{T} = \mathbf{T}'$. But $T'_\zeta \neq T_\zeta$ 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{T} , of subsets of M *linearly diverges* if given any distinct $A, B \in \mathcal{T}$, B linearly diverges from A , with k, t uniform over \mathcal{T} .

Lemma 9.4. *If a family of subsets linearly diverges, then it satisfies (S2) above.*

Proof. Suppose that $\mathbf{A}, \mathbf{B} \in \mathcal{UT}$ 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 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 $\phi : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(\Sigma')$. This gives rise to a continuous map $\phi^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{M}^*(\Sigma')$ with closed image. In fact, each component, $\mathcal{M}_e^*(\Sigma)$, of $\mathcal{M}^*(\Sigma)$ gets sent into the component $\mathcal{M}_{\phi^*(e)}^*(\Sigma')$, of $\mathcal{M}^*(\Sigma')$. Moreover, distinct components get sent into distinct components.

Let $\mathcal{T}(\Sigma)$ be the set of coarse twist flats, $T(\tau)$, as τ ranges over all complete multicurves, τ . This satisfies (S1), and is linearly divergent, by Lemma 7.17. It therefore satisfies (S2) by Lemma 9.4. Note that a Dehn twist flat in a component, M^0 , of $\mathcal{M}^*(\Sigma)$, is by definition, a non-empty set of the form $T^* \cap M^0$ for some $\mathbf{T} \in \mathcal{UT}(\Sigma)$. The same discussion applies to $\mathcal{T}(\Sigma')$.

Let $\mathcal{S} = \{\phi(T) \mid T \in \mathcal{T}(\Sigma)\}$. We claim that $\mathcal{S}, \mathcal{T}(\Sigma)$ satisfies (S3).

Suppose $\mathbf{S} \in \mathcal{US}$. Then $\mathbf{S} = (\phi W_\zeta)_\zeta$, for $W_\zeta \in \mathcal{T}(\Sigma)$. Thus $S^* = \phi^* W^*$, where $\mathbf{W} = (W_\zeta)_\zeta$. Suppose that M^0 is a component of $\mathcal{M}^*(\Sigma')$ with $S^* \cap M^0 \neq \emptyset$. Choose any $e \in W^*$ with $\phi^* e \in M^0$. Thus, $M^0 = \mathcal{M}_{\phi^*(e)}^*(\Sigma')$. We see that $\phi^*(\mathcal{M}_e^*(\Sigma)) = M^0 \cap \phi^*(\mathcal{M}^*(\Sigma))$. Thus, by Theorem 8.13, there is some $T \in \mathcal{T}(\Sigma')$ with $T^* \cap M^0 = S^* \cap M^0$. This verifies property (S3) for $\mathcal{S}, \mathcal{T}(\Sigma')$.

By Lemma 9.3 we now get:

Lemma 9.5. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 2$, and that $\phi : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(\Sigma')$ 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 $\text{hd}(T(\tau'), \phi T(\tau)) \leq k$.*

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 the ξ and the parameters of ϕ . For this, fix some parameters of quasi-isometry, and now take \mathcal{S} to be the set of all images $\phi(T)$, both as T ranges over the set of coarse Dehn twist flats, $\mathcal{T}(\Sigma)$, and as ϕ ranges over all quasi-isometric embeddings from $\mathcal{M}(\Sigma)$ to $\mathcal{M}(\Sigma')$ with these parameters. To verify (Q3) we take $\mathbf{S} = (\phi_\zeta W_\zeta)_\zeta$ and apply Theorem 8.13, to the limiting map ϕ^* of $(\phi_\zeta)_\zeta$. The same argument now gives us a uniform constant, k , independent of any particular ϕ .

10. 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 generalises the result of [Ha, BehKMM].

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. [BehKMM] for example). 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 *properly discontinuously* 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.)

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 $B(G \cap H) = B(G) \cap B(H)$. 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 geodesics if and only if G is two-ended (virtually \mathbb{Z}) and undistorted in X . (Of course, one can say a lot more, but this is all we need here.)

Now, let Σ be a compact surface. Note that $\text{Map}(\Sigma)$ acts discretely on $\mathcal{M}(\Sigma)$. 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 twist flat, $T(\tau)$.

Now if $\gamma \in \mathcal{G}^0(\Sigma)$, 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 $\gamma, \delta \in \mathcal{G}^0(\Sigma)$, then γ, δ are equal or adjacent in $\mathcal{G}(\Sigma)$ 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 $\phi : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(\Sigma')$ is a quasi-isometric embedding.

Suppose that $\tau \subseteq \Sigma$ is a complete multicurve. Now Lemma 9.5 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 9.5, we see that the bound depends only on the complexity of the surfaces and the parameters of quasi-isometry.

Suppose that $\gamma \in \mathcal{G}^0(\Sigma)$. Choose τ, τ' complete multicurves with $\tau \cap \tau' = \{\gamma\}$. Thus $B(\tau) \wedge B(\tau') = B(\gamma) \in \mathcal{Q}(\mathcal{M}(\Sigma))$, and so $B(\theta\tau) \cap B(\theta\tau') \in \mathcal{Q}(\mathcal{M}(\Sigma'))$. It follows that $\theta\tau \cap \theta\tau'$ consists of a single curve, $\delta \in \mathcal{G}^0(\Sigma')$. 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, $\theta : \mathcal{G}^0(\Sigma) \rightarrow \mathcal{G}^0(\Sigma')$ such that $B(\theta\gamma) = \phi B(\gamma)$ for all $\gamma \in \mathcal{G}^0(\Sigma)$. Since, $\phi : \mathcal{B}(\mathcal{M}(\Sigma)) \rightarrow \mathcal{B}(\mathcal{M}(\Sigma'))$ is injective, it follows that θ is injective.

Moreover, if γ, δ are equal or adjacent in $\mathcal{G}(\Sigma)$, 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 $\mathcal{G}(\Sigma')$. In other words, θ gives an injective embedding of $\mathcal{G}(\Sigma)$ into $\mathcal{G}(\Sigma')$.

We now use the following fact from [Sha].

Theorem 10.1. [Sha] *Suppose that Σ and Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 4$. If $\theta : \mathcal{G}(\Sigma) \rightarrow \mathcal{G}(\Sigma)$ is an injective embedding, then $\Sigma = \Sigma'$ and there is some $g \in \text{Map}(\Sigma)$ such that $\theta\gamma = g\gamma$ for all $\gamma \in \mathcal{G}^0(\Sigma)$. The same conclusion holds if Σ, Σ' are both an $S_{2,0}$, both an $S_{0,6}$, 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 $g \in \mathcal{G}^0(\Sigma)$. 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 9.5 gives us a uniform k such that $hd(T(\tau'), \phi T(\tau)) \leq k$ for some multicurve τ' in Σ . But we now know that $\tau' = \tau$, and so we deduce that $hd(T(\tau), \phi T(\tau)) \leq k$ for all multicurves, τ .

Now if $x \in \mathcal{M}(\Sigma)$, we can always find τ, τ' with $\tau \cap \tau' = \emptyset$, and with $\iota(\tau, \tau')$ and $\rho(x, T(\tau))$ and $\rho(x, T(\tau'))$ uniformly bounded. It follows that ϕx is a bounded distance from both $\phi T(\tau)$ and $\phi T(\tau')$ and so 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 10.2. *Suppose that Σ and Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 4$, and that $\phi : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(\Sigma')$ is a quasi-isometric embedding. Then $\Sigma = \Sigma'$ and there is some $g \in \text{Map}(\Sigma)$ such that for all $a \in \mathcal{M}(\Sigma)$, 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 $\mathcal{M}(\Sigma)$ into $\mathcal{M}(\Sigma')$, 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 $\mathcal{M}(\Sigma)$. It is not clear when a quasi-isometric embedding exists if $\xi(\Sigma) < \xi(\Sigma')$.)

One can also describe the lower complexity cases. Write $S_{g,p}$ for the closed orientable surface of genus g with p holes. 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 [Sha], quoted as Theorem 10.1 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, $\mathcal{M}(S_{2,0})$ and $\mathcal{M}(S_{0,6})$

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 10.1). However, if $\Sigma = \Sigma' = S_{1,2}$, then the conclusion of Theorem 10.1 fails without further hypotheses (see [Sha]). 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 $\mathcal{M}(S_{1,2})$ is quasi-isometric to $\mathcal{M}(S_{0,5})$, 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.

Finally 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 classes of quasi-isometries between them up to bounded distance.

This above generalises the results of [Ha, BehKMM], where it is assumed that ϕ is a quasi-isometry and $\Sigma = \Sigma'$.

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

11. QUOTIENT MAPS

We discuss quotients of coarse median spaces, with a view to showing (in Section 12) that the pants graph is a coarse median space.

Lemma 11.1. *Suppose that (Λ, ρ, μ) is a coarse median space. Suppose that Π is a finite median algebra with $|\Pi| \leq q < \infty$, and that $\lambda : \Pi \rightarrow \Lambda$ is a h -quasimorphism. Given $t \geq 0$, there is a finite median algebra Π' , a map $\lambda' : \Pi' \rightarrow \Lambda$ 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 Λ .*

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

We claim that if $x, y \in \Pi$, with $x \sim y$, then $\rho(\theta\lambda x, \theta\lambda y)$ is bounded above in terms of q, h, t and the parameters of Λ . To see this, note that \sim 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 \sim . From the fact that λ is a quasimorphism, and from property (C1) for Λ , we see that at each stage the maximal distance between the λ -images of related elements of Π can increase by most a linear function which depends only on h and the parameters of Λ . This now proves the claim.

Suppose that $z \in \Pi$. By construction, $\lambda'\theta z = \lambda w$, for some $w \sim z$. 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 h' -quasimorphism, where h' depends only on q, h, t and the parameters of Λ . This enables us to give a refinement of (C2) as follows:

Corollary 11.2. *Suppose that (Λ, 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 $A \subseteq \Lambda$ with $1 \leq |A| \leq p < \infty$. Then there is a finite median algebra, (Π, μ_Π) , and maps $\pi : A \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ 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 Λ (so that $\lambda : \Pi \rightarrow \Lambda$ is an $h(p)$ -quasimorphism, with $h(p)$ independent of t). We can assume that $|\Pi| \leq 2^{2^p}$. We now apply Lemma 11.1 to give Π', λ' and $\theta : \Pi \rightarrow \Pi'$. Now replace Π by Π' , π by $\theta \circ \pi$, and λ by λ' . \square

We can also use this construction to give a criterion for bounding the rank of a coarse median space, possibly at the expense of modifying the parameters.

Suppose that (Λ, ρ, μ) is a coarse median space. Given $n \in \mathbb{N}$, let h_0 be the constant of Lemma 6.2, which depends only on n and the parameters of Λ (which make no a-priori reference to rank). We suppose:

(R) There is some $t_0 \geq 0$ such that if Q is an $(n+1)$ -cube and $\phi : Q \rightarrow \Lambda$ is any h_0 -quasimorphism, then there are distinct $x, y \in Q$ with $\rho(\phi x, \phi y) \leq t_0$.

Note that Lemma 6.2 now implies that if $\psi : Q \rightarrow \Lambda$ is any h -quasimorphism for any $h \geq 0$, then there are distinct $x, y \in Q$ with $\rho(\psi x, \psi y) \leq t_1$, where $t_1 = t_0 + 2(k_0 h + h_1)$, where k_0, h_1 are the constants of Lemma 6.2.

Lemma 11.3. *Suppose that (Λ, ρ, μ) is a coarse median space satisfying (R) above. Then (Λ, ρ, μ) has rank at most Λ .*

As mentioned above, the constant h_0 , of (R) refers to the initial parameters of Λ . The parameters with respect to which it has rank at most n will depend on these initial parameters, together with n and t_0 of the hypotheses.

Proof. Let $h : \mathbb{N} \rightarrow [0, \infty)$ be as in (C2) for Λ . Suppose that $A \subseteq \Lambda$ with $|A| \leq p < \infty$. Let Π, π, λ be as given by (C2) for Λ . Set $t_1 = t_0 + 2(k_0 h(p) + h_1)$. Now let Π', λ', θ be as in the proof of Corollary 11.2 with $t = t_1$. We claim that Π' has rank at most n (as a median algebra).

Suppose, for contradiction that $Q' \subseteq \Pi'$ is an $(n+1)$ -cube, which we can take to be convex in Π' (that is, a cell of the corresponding cube complex). Then there is some $(n+1)$ -cube, $Q \subseteq \Pi$, with $\theta|_Q$ an isomorphism onto Q' . But applying (R) as discussed above to $\lambda|_Q$, we see that there exist distinct $x, y \in Q$ with $\rho(\lambda x, \lambda y) \leq t_1$. But x, y lie in distinct \approx -classes, that is, $\rho(\lambda x, \lambda y) > t_1$, giving a contradiction. This shows that $\text{rank}(\Pi') \leq n$ as claimed.

Now, as in Corollary 11.2, we replace Π by Π' , π by $\theta \circ \pi$ and λ by λ' . This also satisfies (C2), where the relevant constants depend only on the original constant $h(p)$ and on n and t . \square

Suppose now that (Λ, ρ, μ) and (Λ', ρ', μ') are coarse median spaces and that $\eta : \Lambda \rightarrow \Lambda'$ is an h_1 -quasimorphism with $\Lambda' = N(\eta(\Lambda); h_2)$.

Lemma 11.4. *Suppose that Q is an n -cube, and that $\psi : Q \rightarrow \Lambda'$ is an h -quasimorphism. Then there is an h' -quasimorphism $\phi : Q \rightarrow \Lambda$ such that for all $x \in Q$, $\rho(\psi x, \eta \phi x) \leq h''$, where h', h'' depend only on h, h_1, h_2 and the parameters of Λ .*

Proof. We proceed similarly as in the proof of Lemma 6.1. Let Π be the free median algebra on Q . Let $\omega : Q \rightarrow \Pi$ and $\theta : \Pi \rightarrow Q$ be homomorphisms, with $\theta|_Q$ the inclusion of Q in Π and with $\theta \circ \omega$ the identity on Q .

Let $\hat{\psi} : Q \rightarrow \Lambda$ be any map such that $\rho'(\psi x, \eta \hat{\psi} x) \leq h''$ for all $x \in Q$. Now apply (C2) in Λ to $\hat{\psi}(Q) \subseteq \Lambda$ to give us an $h(2^n)$ -quasimorphism $\hat{\psi} : \Pi \rightarrow \Lambda$ with $\rho'(\hat{\lambda} x, \hat{\psi} x) \leq h(2^n)$ for all $x \in Q$. Let $\lambda = \eta \circ \hat{\lambda} : \Pi \rightarrow \Lambda'$ and let $\lambda' = \eta \circ \hat{\lambda} \circ \theta : \Pi \rightarrow \Lambda'$. These are both quasimorphisms and $\lambda|_Q = \eta \circ \hat{\lambda}|_Q = \eta \circ \hat{\lambda} \circ \theta|_Q = \lambda'|_Q$. Therefore, Lemma 6.1 applied to Λ' tells us that $\rho'(\lambda x, \lambda' x)$ is bounded for $x \in \Pi$.

Let $\phi = \hat{\lambda} \circ \omega : Q \rightarrow \Lambda$. This is a quasimorphism.

Now $\eta \circ \phi = \eta \circ \hat{\lambda} \circ \omega = \lambda \circ \omega$ and $\eta \circ \hat{\lambda} = \eta \circ \hat{\lambda} \circ \theta \circ \omega = \lambda' \circ \omega$ are a bounded distance apart (since λ and λ' are). Moreover, by construction, $\eta \circ \hat{\lambda}$ and ψ are a bounded distance apart. Therefore $\eta \circ \phi$ and ψ are a

bounded distance apart, as required. Note that all the bounds depend only on h, h_1, h_2 , and the parameters of Λ and Λ' as required. \square

Combining this with the previous results, we get:

Lemma 11.5. *Suppose that Λ, Λ', η are as above. Then there is some h_3 , depending only on the parameters of the hypotheses with the following property. Suppose that there is some $t \geq 0$ such that for any h_3 -quasimorphism $\phi : Q \rightarrow \Lambda$, of an $(n+1)$ cube, Q into Λ , there are distinct $x, y \in Q$ with $\rho'(\eta\phi x, \eta\phi y) \leq t$. Then $\text{rank}(\Lambda') \leq n$.*

For the conclusion, we may have to modify the parameters of Λ , depending on the parameters of the hypotheses and n and t .

Proof. We just need to verify the hypotheses of Lemma 11.3. Let h_0 be as given by Lemma 11.3 for Λ' , with n replaced by $n+1$. Thus, h_0 depends only on the parameters of Λ' and on n . Suppose that $\psi : Q \rightarrow \Lambda'$ is an h_0 -quasimorphism of an $(n+1)$ -cube, Q , into Λ' . Let $\phi : Q \rightarrow \Lambda$ be as given by Lemma 11.4. Thus ϕ is an h_3 -quasimorphism, and $\rho(\eta\phi x, \phi x) \leq h_4$ for all $x \in Q$, where h_3, h_4 depend only on the parameters on n . Therefore, by hypothesis, there are distinct $x, y \in Q$ such that $\rho'(\eta\phi x, \eta\phi y) \leq t$. Thus $\rho'(\psi x, \psi y) \leq t + 2h_4$. We have therefore verified (R) with $t_0 = t + 2h_4$, for an arbitrary h_0 -quasimorphism $\psi : Q \rightarrow \Lambda'$. It follows from Lemma 11.3 that $\text{rank}(\Lambda') \leq n$. \square

The following gives a construction of quotient coarse median spaces, analogous to the process for median algebras described at the end of Section 2.

Suppose that (Λ, ρ, μ) is a coarse median space, and that (Λ', ρ') is a geodesic metric space. Suppose that $\eta : \Lambda \rightarrow \Lambda'$ is a map (not necessarily continuous) satisfying:

$$(Q1) (\exists k_0, h_0 \geq 0)(\forall a, b \in \Lambda), \rho'(\eta a, \eta b) \leq k\rho(a, b) + h_0.$$

$$(Q2) (\exists h_1 \geq 0) \text{ such that } N(\eta(\Lambda); h_1) = \Lambda'.$$

$$(Q3) (\exists k_2, h_2 \geq 0)(\forall a, b, c, d \in \Lambda), \rho'(\eta\mu(a, b, c), \eta\mu(a, b, d)) \leq k\rho'(\eta c, \eta d) + h_2.$$

Note that since we are dealing with geodesic spaces, (Q1) is equivalent to saying that there exist $h, h' > 0$ such that if $\rho(a, b) \leq h$, then $\rho'(\eta a, \eta b) \leq h'$. Similarly, (Q3) is equivalent to saying that there exist $h, h' > 0$, such that if $\rho'(\eta c, \eta d) \leq h$, then $\rho'(\eta\mu(a, b, c), \eta\mu(a, b, d)) \leq h'$.

We can define a ternary operation, μ' , on Λ' as follows. Given $a, b, c \in \Lambda'$, choose $a_0, b_0, c_0 \in \Lambda$ with $\rho'(a, \eta a_0) \leq h_1$, $\rho'(b, \eta b_0) \leq h_1$ and $\rho'(c, \eta c_0) \leq h_1$, and set $\mu'(a, b, c) = \eta\mu(a_0, b_0, c_0)$. By (Q2), this is well defined up to bounded distance, independently of the choice of a_0, b_0, c_0 . Also, by construction, η is quasimorphism.

Lemma 11.6. *(Λ', ρ', μ') is a coarse median space, whose parameters depend only on those of Λ and of (Q1)–(Q3). Moreover, $\text{rank}(\Lambda') \leq \text{rank}(\Lambda)$. Also, if Λ is n -colourable for some $n \in \mathbb{N}$, then so is Λ' .*

Proof. We need to verify (C1) and (C2) of the definition. Property (C1) is a simple consequence of (Q3).

For (C2), suppose that $A \subseteq \Lambda'$ with $|A| \leq p$. Using (Q2), we can find an injective map, $f : A \rightarrow \Lambda$, with $\rho'(a, \eta fa)$ bounded for all $a \in A$. Now let $\Pi, \pi : f(A) \rightarrow \Pi$ and $\lambda : \Pi \rightarrow \Lambda$ be as given by (C2) for $f(A) \subseteq \Lambda$. Let $\pi' = \pi \circ f : A \rightarrow \Pi$ and $\lambda' = \eta \circ \lambda : \Pi \rightarrow \Lambda'$. Thus, λ' is a quasimorphism, and if $a \in A$, then $\lambda'\pi'a = (\eta\lambda)(\pi f)a = \eta(\lambda\pi)(fa)$. Now $\rho(fa, (\lambda\pi)(fa))$ is bounded (from (C2) in Λ), and so, by (Q1), $\rho'(\eta(fa), \eta(\lambda\pi)(fa))$ is bounded. Thus $\eta(a, \lambda'\pi'a)$ is bounded as required. \square

12. THE PANTS GRAPH

In this section, we apply the results of the previous section to the pants graph of a compact surface. The pants graph is quasi-isometric to Teichmüller space in the Weil-Petersson metric [Bro].

Let Σ be a compact orientable surface, with $\xi(\Sigma) \geq 2$. The *pants graph* is traditionally defined as follows. Let \mathcal{E}^0 be the set of complete multicurves in Σ . Let $\mathcal{E} = \mathcal{E}(\Sigma)$ be the graph with vertex set $V(\mathcal{E}) = \mathcal{E}^0$, and where $x, y \in \mathcal{E}^0$ are deemed to be adjacent if there is some $\gamma \in x$ and $\delta \in y$ such that $x \setminus \gamma = y \setminus \delta = u$, say and with $\iota(\gamma, \delta)$ equal to 1 or 2 depending on whether γ (hence also δ) is contained in an $S_{1,1}$ or a $S_{0,4}$ component of $\Sigma \setminus u$. Note that if Σ is an $S_{1,1}$ or $S_{0,4}$, then we can identify $\mathcal{E}(\Sigma)$ and $\mathcal{G}(\Sigma)$.

The notion is quite robust. For example, if $q \geq 2$, let $\mathcal{E}(\Sigma, q)$ be the graph with vertex set \mathcal{E}^0 , and where x, y are adjacent if $\iota(x, y) \leq q$. Thus, $\mathcal{E} \subseteq \mathcal{E}(\Sigma, q)$, and one sees easily that the inclusion is a quasi-isometry.

There is a map $\eta : \mathcal{M} \rightarrow \mathcal{E}$, well defined up to bounded distance. For example, fix a universal constant, sufficiently large, and given $a \in \mathcal{M}^0$ and given $a \in \mathcal{M}^0$, let $\eta(a)$ be any multicurve with $\iota(a, \eta(a)) \leq r$. This gives rise to a coarsely lipschitz map from \mathcal{M} to \mathcal{E} . Note that the image of η is cobounded, and so η satisfies the conditions (Q1) and (Q2) of Section 11. We will see that it also satisfies (Q3).

There are various ways of describing \mathcal{E} up to quasi-isometry. For example, we can start with \mathcal{M} and cone each of the coarse Dehn twist flats, $T(x)$ for $x \in \mathcal{E}^0$. More formally, let \mathcal{E}' be the graph with vertex set $V(\mathcal{E}') = \mathcal{M}^0 \sqcup \mathcal{E}^0$. We deem $a, b \in \mathcal{M}^0$ to be adjacent if a, b are adjacent in \mathcal{M}^0 , and we deem $a \in \mathcal{M}^0$ to be adjacent to $x \in \mathcal{E}^0$ if $a \in T(x)$. Now \mathcal{E}' is quasi-isometric to \mathcal{E} , and η corresponds to the inclusion of \mathcal{M} into \mathcal{E} . Again this is robust: for example we could equally well use any family $\{T'(x)\}_{x \in \mathcal{E}^0}$ with $\text{hd}(T(x), T'(x))$ uniformly bounded.

Distances in \mathcal{E} can be estimated up to linear bounds by a formula given in Chapter 9 of [MasM2], similar to those for the marking graph (given as Theorem 7.3 here). To describe this, let $\mathcal{X}_N \subseteq \mathcal{X}$ be the set of non-annular elements of \mathcal{X} , i.e. subsurfaces of complexity at least 1. Note that if $X \in \mathcal{X}_N$, then $\theta_X(x) \in \mathcal{G}(X)$ is defined for all $x \in \mathcal{E}^0$, and so $\theta_X(x, y) = \theta_X(\theta_X x, \theta_X y)$ is defined for all $x, y \in \mathcal{E}^0$. Given $r \geq 0$, let $\mathcal{A}_N(x, y; r) = \{X \in \mathcal{X}_N \mid \sigma_X(x, y) > r\}$. The formula in [MasM2] is now essentially the same as that for \mathcal{M} , except that we now restrict the sum of the $\sigma_X(x, y)$ to those $X \in \mathcal{A}_N(x, y)$ for sufficiently large Y . (In other words we ignore the contribution of subsurface projections to annuli.)

In particular, we have the following variation of Proposition 7.2:

Proposition 12.1. (1) *There is some universal $r_0 \geq 0$ such that if $x, y \in \mathcal{E}^0(\Sigma)$ then $\mathcal{A}_N(x, y; r_0)$ is finite.*
 (2) *Given $r \geq 0$, there is some $r' \geq 0$, such that if $\mathcal{A}_N(x, y, r) = \emptyset$, then $\rho(x, y) \leq r'$.*

Note that if $X \in \mathcal{X}$, then $\theta_X(\eta a) \sim \theta_X(a)$ for all $x \in \mathcal{M}^0$, and so $\sigma_X(a, b) \sim \sigma_X(\eta a, \eta b)$ for all $X \in \mathcal{X}_N$.

To verify (Q3) for the map $\eta : \mathcal{M} \rightarrow \mathcal{E}$, it is enough to show that $\rho(\eta\mu(a, b, c), \eta\mu(a, b, d))$ is bounded for all $X \in \mathcal{X}$, and so therefore is $\sigma_X(c, d)$. Thus, $\sigma_X(\eta\mu(a, b, c), \eta\mu(a, b, d)) \sim \sigma_X(\mu(a, b, c), \mu(a, b, d))$ is also bounded as required. Thus, by Lemma 11.6, we deduce:

Lemma 12.2. *\mathcal{E} is a coarse median space.*

In fact, we can characterise the coarse median up to bounded distance by saying that $\theta_X(\mu(x, y, z))$ agrees with $\mu(\theta_X x, \theta_X y, \theta_X z)$ up to bounded distance for all $X \in \mathcal{X}_N$.

Note that Lemma 11.6 also tells us that $\text{rank}(\mathcal{E}) \leq \xi = \xi(\Sigma)$. In fact, we can do better.

Definition. A multicurve, τ , is *good* if every curve of τ separates Σ and there are no $S_{0,3}$ elements of $\mathcal{X}_N(\tau)$.

One can easily check that good multicurves are precisely those which minimise $|\mathcal{X}_N(\tau)|$ subject to the constraint that $\mathcal{X}_N(\tau)$ contains no $S_{0,3}$. Note that, if $\xi = \xi(\Sigma)$ is odd then there are $(\xi + 1)/2$ elements, each of complexity 1. If ξ is even, there are $(\xi/2) - 1$ elements of complexity 1, and one element of complexity 2, making $\xi/2$ in total. In general therefore, the maximal value of $|\mathcal{X}_N(\tau)|$, when there is no $S_{0,3}$ is $\lfloor (\xi + 1)/2 \rfloor$, attained in the above situations. We set $\xi_0(\Sigma) = \lfloor (\xi(\Sigma) + 1) \rfloor$.

We make the following simple topological observation.

Lemma 12.3. *Suppose that $n \in \mathbb{N}$, and that for each $i \in \{1, \dots, n\}$ we have a non-empty subset, $\mathcal{Y}_i \subseteq \mathcal{X}_N$ such that $X \wedge Y$ whenever $X \in \mathcal{Y}_i$ and $Y \in \mathcal{Y}_j$ with $i \neq j$. Then $n \leq \xi_0$.*

Proof. One way to see this is to note that there is a multicurve, τ , such that each element of \mathcal{Y}_i lies in some union, B_i , of components of $\Sigma \setminus \tau$, with the B_i disjoint. Thus $n \leq |\mathcal{X}_N(\tau)| \leq \xi_0$. \square

Suppose now that $Q = \{-1, 1\}^n$ is an n -cube for some n , and that $\phi : Q \rightarrow \mathcal{M}$ is a h -quasimorphism. We assume that $\phi(Q) \subseteq \mathcal{M}^0$. Now, choose s sufficiently large, and let $\mathcal{A}(i) = \mathcal{A}(\phi, i; s)$ as described in Section 7. All we require of this is that if $X \in \mathcal{A}(i)$ and $Y \in \mathcal{A}(j)$ for $i \neq j$, then $X \wedge Y$, and moreover that if $v, w \in Q$ differ only in the i th coordinate, then $\sigma_X(\phi v, \phi w)$ is bounded for all $X \in \mathcal{X} \setminus \mathcal{A}(i)$. Here the relevant constants depend only on h .

Now set $\mathcal{A}_N(i) = \mathcal{A}(i) \cap \mathcal{X}_N$ for $i = 1, \dots, n$. Lemma 12.3 tells us that $\mathcal{A}_N(i)$ is empty for all but at most ξ_0 indices i . Therefore if $n > \xi_0$, there must be some i with $\mathcal{A}_N(i) = \emptyset$. It follows that if $v, w \in Q$ differ only in the i th coordinate, then $\sigma_X(\eta\phi v, \eta\phi w) \sim \sigma_X(\phi v, \phi w)$ is bounded (in terms of h) for all $X \in \mathcal{X}_N$. Therefore, by Proposition 12.1, $\rho(\eta\phi v, \eta\phi w)$ is bounded.

On setting $h = h_3$, we have verified the hypotheses of Lemma 11.5, and so:

Lemma 12.4. $\text{rank}(\mathcal{E}) \leq \xi_0$.

In fact, it's not hard to see that $\text{rank}(\mathcal{E}) = \xi_0$. For example, let τ be a good multicurve, and choose any pseudoanosov, g_X , in each subsurface, $X \in \mathcal{X}_N(\tau)$. Let $G \leq \text{Map}(\Sigma)$ be the subgroup generated by $\{g_X \mid X \in \mathcal{X}_N(\tau)\}$, so that $G \cong \mathbb{Z}^{\xi_0}$. It's not hard to verify that the map $[g \mapsto ga] : G \rightarrow \mathcal{E}$ is a quasi-isometric embedding. (This will be discussed further in Section 13.)

We also remark that it follows from the discussion of Section 11 that \mathcal{E} is finitely colourable.

We summarise what we have shown as follows:

Theorem 12.5. *Given $x, y, z \in \mathcal{E}(\Sigma)$, there is some $m \in \mathcal{E}(\Sigma)$ such that for all $X \in \mathcal{X}_N$, we have that $\theta_X(m, \mu(x, y, z))$ is bounded. If $m' \in \mathcal{E}(\Sigma)$ is another such element, then $\rho(m, m')$ is bounded. Setting $\mu(x, y, z) = m$ for any such m , $(\mathcal{E}(\Sigma), \rho, \mu)$ is a coarse median space of rank $\lfloor (\xi(\Sigma) + 1)/2 \rfloor$ and finitely colourable. The natural map, $\eta : \mathcal{M}(\Sigma) \rightarrow \mathcal{E}(\Sigma)$ is a median quasimorphism. Here all constants depend only on $\xi(\Sigma)$.*

We can now go on to apply the various constructions regarding coarse median spaces.

For example, the extended asymptotic cone, \mathcal{E}^* , is a finitely colourable topological median algebra of rank ξ_0 . The asymptotic cone, \mathcal{E}^∞ , embeds in a finite product of \mathbb{R} -trees, by a median homomorphism which is bilipschitz onto its range. Therefore, \mathcal{E}^∞ is bilipschitz equivalent to a median metric space, inducing the same median structure. (Hence, it is, in turn, bilipschitz equivalent to a CAT(0) space — but in this case, we already knew this, since \mathcal{E} is quasi-isometric to Teichmüller space in the Weil-Petersson metric which is CAT(0) [Wo]) Moreover, \mathcal{E}^∞ has (locally) compact dimension equal to ξ_0 . From this, we can deduce the following, proven by different methods, in [BehM1] (see also [EsMR]):

Theorem 12.6. *If \mathbb{R}^n quasi-isometrically embeds in $\mathcal{E}(\Sigma)$, then $n \leq \lfloor (\xi(\Sigma) + 1)/2 \rfloor$.*

In particular, $\xi_0(\Sigma)$ is determined by the quasi-isometry class of $\mathcal{E}(\Sigma)$ as the maximal dimension of a quasi-isometrically embedded euclidean space.

Note that, in fact, we can strengthen this using Proposition 6.6, to bound the radius of a euclidean $(\xi_0(\Sigma) + 1)$ -ball which we can quasi-isometrically embed in $\mathcal{E}(\Sigma)$. In particular, it follows that we could replace \mathbb{R}^n by the euclidean half-space, $\mathbb{R}^{n-1} \times [0, \infty)$ in the statement of Theorem 1.2 and the subsequent remark.

Note that if $\xi(\Sigma) \leq 2$, then $\xi_0(\Sigma) = 1$, so $\mathcal{E}(\Sigma)$ is a coarse median space of rank 1, hence hyperbolic [Bo1]. We recover the following result of [BroF] (see also [Ar]):

Theorem 12.7. *The Weil-Petersson metric on the Teichmüller space of $S_{1,2}$ or $S_{0,5}$ is hyperbolic.*

13. CUBES IN THE PANTS GRAPH

Let Σ be a compact surface, and let $\xi_0(\Sigma) = \lfloor (\xi(\Sigma) + 1)/2 \rfloor$. We will suppose that $\xi(\Sigma) \geq 3$, so $\xi_0(\Sigma) \geq 2$.

Recall that we have a map $\eta : \mathcal{M}(\Sigma) \rightarrow \mathcal{E}(\Sigma)$, and $\theta_X : \mathcal{E}(\Sigma) \rightarrow \mathcal{G}(X)$, for $X \in \mathcal{X}_N(\Sigma)$. These are all coarsely lipschitz median quasimorphisms. Moreover, $\theta_X \circ \eta : \mathcal{M}(\Sigma) \rightarrow \mathcal{G}(X)$ agrees (up to bounded distance) with the map denoted θ_X defined in Section 7. We also have a quasimorphism $\psi_X : \mathcal{E}(\Sigma) \rightarrow \mathcal{E}(X)$, such that $\psi_X \circ \eta : \mathcal{M}(\Sigma) \rightarrow \mathcal{E}(X)$ agrees up to bounded distance with $\eta \circ \psi_X : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(X) \rightarrow \mathcal{E}(X)$. (The map ψ_X be defined similarly as with subsurface projection of curves. If $x \in \mathcal{E}(\Sigma)$, then $\psi_X(x)$ is a complete multicurve in X with the property that its intersection with any essential connected component of $x \cap X$ is uniformly bounded. The result is well defined up to bounded distance.) Note that if $\xi(X) = 1$ (i.e. X is an $S_{1,1}$ or $S_{0,4}$), then ψ_X can be identified with θ_X .

Given a multicurve $\tau \subseteq \Sigma$, let $T_E(\tau) = \{x \in \mathcal{E}(\Sigma) \mid \tau \subseteq x\}$, i.e. all complete multicurves that contain τ . (Note that if $y \in \mathcal{E}(\Sigma)$, then $\rho(y, T_E(\tau))$ is bounded above in terms of $\iota(y, \tau)$. This gives an alternative way to define $T_E(\tau)$ up to bounded distance as $\{y \in \mathcal{E}(\Sigma) \mid \iota(x, \tau) \leq r\}$ for any fixed $r \geq 0$.)

Let $\mathcal{T}_E(\tau)$ be the direct product $\prod_{X \in \mathcal{X}_N(\tau)} \mathcal{E}(X)$. (Since we are only interested in $\mathcal{T}_E(\tau)$ up to quasi-isometry, we can define the metric in several different, but quasi-isometrically equivalent, ways. Indeed we can just restrict to the path metric in the 1-skeleton of the product cube complex.) Note that $\mathcal{T}_E(\tau)$ is naturally a coarse median space.

Similarly as in Section 7 (with \mathcal{E} now playing the role of \mathcal{M}) we can define a maps $\psi_\tau : \mathcal{E}(\Sigma) \rightarrow \mathcal{T}(\tau)$ and $v_\tau : \mathcal{T}(\tau) \rightarrow \mathcal{E}(\tau)$. Here $\psi_\tau(x)$ a product of $\psi_X(x)$ for $X \in \mathcal{X}_N(\tau)$, and v_τ takes the multicurve on each factor $\mathcal{E}(X)$ for each $X \in \mathcal{X}_N(\tau)$, and assembles together with τ itself to give a complete multicurve $v_\tau(x)$. Note that $v_\tau(\mathcal{T}_E(\tau)) \subseteq T_E(\tau)$, and and that $\omega_\tau = v_\tau \circ \psi_\tau$ is a gate map of $\mathcal{E}(\Sigma)$ to $T_E(\tau)$. In fact:

Lemma 13.1. *ω_τ is a quasiprojection (in the sense of coarse medians).*

Proof. This follows similarly as with \mathcal{M} (cf. Lemma 7.10). Note that if $X \in \mathcal{X}_N$ with $X \wedge \tau$, then $\theta_X(\omega_\tau(x)) \sim \theta_X(x)$. Otherwise $\theta_X(\omega_\tau(x)) \sim \theta_X(\tau) \sim \theta_X(y)$ for all $y \in T_E(\tau)$. From the characterisation of the median, μ , on $\mathcal{E}(\Sigma)$, it follows that $\mu(\omega_\tau(x), x, y) \sim \omega_\tau(x)$. \square

We see that $T_E(\tau)$ is a quasi-isometrically embedded copy of $\mathcal{T}_E(\tau)$. The main interest to us is the case where τ is good. Then, all the factors $\mathcal{E}(X)$ are hyperbolic, and all but at most one (in the case where $\xi(\Sigma)$ is odd) is a quasitree.

Next we describe quasicubes in $\mathcal{E}(\Sigma)$.

Given an n -cube, $Q = \{-1, 1\}^n$, a map $\phi : Q \rightarrow \mathcal{E}$, and some $s \geq 0$, write $\mathcal{A}_N(\phi, i; s)$ for the set of $X \in \mathcal{X}_N$ with $\sigma_X(\phi x, \phi y) > s$ whenever $x, y \in Q$ is a pair which differ (at least) in their i th coordinate.

Lemma 13.2. *Given $\xi = \xi(\Sigma)$, n and h , there is some s such that if $\phi : Q \rightarrow \mathcal{E}$ is an h -quasimorphism from an n -cube Q , and $X \in \mathcal{A}_N(\phi, s)$ and $Y \in \mathcal{A}_N(\phi, s)$ with $i \neq j$, then $X \wedge Y$.*

Proof. This is just an elaboration of the proof of Lemma 12.4. By Lemma 11.4, we have a uniform quasimorphism $\hat{\phi} : Q \rightarrow \mathcal{M}$ such that $\eta \circ \hat{\phi}$ agrees with ϕ up to bounded distance. If $X \in \mathcal{X}_N$, we see that $\theta_X \circ \hat{\phi} \sim \theta_X \sim \eta \circ \phi \sim \theta_X \circ \phi$, so if $x, y \in Q$, then $\xi(\phi x, \phi y) \sim \sigma_X(\hat{\phi} x, \hat{\phi} y)$, and the statement follows by Lemma 7.7. \square

We now set $n = \xi_0$ and fix h_0 as in Lemma 6.2 for this n (so that, in practice we only need to deal with h_0 -quasimorphisms). We now choose s as in Lemma 13.2, and abbreviate $\mathcal{A}_N(\phi, i) = \mathcal{A}_N(\phi, i; s)$. We say that ϕ is *degenerate* if there is some i with $\mathcal{A}_N(\phi, i) \neq \emptyset$, and *non-degenerate* otherwise. Note that, by Proposition 12.1, if ϕ is degenerate, then there are distinct $x, y \in Q$ with $\rho(\phi x, \phi y)$ uniformly bounded in terms of ξ (and our choice of s).

In what follows, the discussion largely splits into two cases, depending on the parity of $\xi(\Sigma)$.

Definition. We say that a compact orientable surface, Σ , is *even* (respectively *odd*) if $\xi(\Sigma)$ is even (respectively odd).

The odd case is somewhat simpler to describe, and we will deal mainly with that case first. For the moment, we will assume that Σ is odd, unless otherwise stated. In statements of lemmas (in particular Lemmas 13.7 to 13.12) we will omit this hypothesis where it is not necessary. However, we will first only give proofs in the odd case, and describe later how to prove them in the even case.

We now proceed with the description of quasicubes. Let $Q = \{-1, 1\}^{\xi_0}$ be a ξ_0 -cube, and let $\phi : Q \rightarrow \mathcal{E}$ be a non-degenerate h_0 -quasimorphism.

Lemma 13.3. *Suppose that Σ is odd. Then there is a good multicurve, τ , such that for each $i \in \{1, \dots, \xi_0\}$, there is some $Y_i \in \mathcal{X}_N(\tau)$ with $\mathcal{A}(\phi, i) = \{Y_i\}$.*

In other words, $\{Y_1, \dots, Y_{\xi_0}\}$ is precisely the set of complementary components of τ which are $S_{1,1}$'s or $S_{0,4}$'s.

Proof. By assumption, each $\mathcal{A}(\phi, i) \neq \emptyset$, so choose any $Y_i \in \mathcal{A}(\phi, i)$. By Lemma 13.2, the Y_i are all disjoint, they must be precisely the

complementary components of a good multicurve τ . If $Y \in \mathcal{A}(\phi, i)$, again, by Lemma 13.2, we have $Y \wedge Y_j$ for all $j \neq i$, and so $Y = Y_i$. \square

Note that τ is uniquely determined, and we write it as $\tau(\phi)$. To describe the even case, if τ is a good multicurve, we write $W(\tau)$ for the complexity-2 element of $\mathcal{X}_N(\tau)$. (We will later use the same notation for non-standard good multicurves.)

We now pass the (extended) asymptotic cones. Choose a positive infinitesimal and basepoint, and let $\mathcal{E}^* = \mathcal{E}^*(\Sigma)$ be the extended asymptotic cone of \mathcal{E} , with extended metric, ρ^* . As before, $\mathcal{U}\text{Map}(\Sigma)$ acts transitively on $\mathcal{E}^*(\Sigma)$. The map $\eta : \mathcal{M} \rightarrow \mathcal{E}$ gives rise to a map $\eta^* : \mathcal{M}^* \rightarrow \mathcal{E}^*$. Similarly, if $X \in \mathcal{U}\mathcal{X}_N(\Sigma)$, we get maps $\psi_X^* : \mathcal{E}^*(\Sigma) \rightarrow \mathcal{G}^*(X)$ and $\psi_X^* : \mathcal{E}^*(\Sigma) \rightarrow \mathcal{E}^*(X)$. Note that $\theta_X^* = \theta_X^* \circ \eta : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{G}^*(X)$, using the same notation, as in Section 8. Also $\eta^* \circ \psi_X^* = \psi_X^* \circ \eta^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{E}^*(X)$. All of the above maps are uniformly lipschitz median homomorphisms.

Suppose that $Q = \{-1, 1\}^{\xi_0}$ is a ξ_0 -cube in \mathcal{E} . Given $i \in \{1, \dots, \xi_0\}$, let a, b be a side of Q crossing the i th wall (i.e. a, b differ precisely in i th coordinate). We write $A_i(Q) = \{X \in \mathcal{U}\mathcal{X}_N \mid \theta_X^* a \neq \theta_X^* b\}$, $B_i(Q) = \{X \in \mathcal{U}\mathcal{X}_N \mid \psi_X^* a \neq \psi_X^* b\}$, $C_i(Q) = \{X \in \mathcal{U}\mathcal{X}_N \mid \theta_X^*|_{[a, b]}$ is injective and $D_i(Q) = \{X \in \mathcal{U}\mathcal{X}_N \mid \psi_X^*|_{[a, b]}$ is injective}. Note that any two such sides are parallel, so the above are well defined independently of the choice of a, b . Clearly $C_i(Q) \subseteq A_i(Q) \subseteq B_i(Q)$, and $C_i(Q) \subseteq D_i(Q) \subseteq B_i(Q)$. Note also that if $Q' \subseteq Q$ is a smaller cube (i.e. $Q' \subseteq \text{hull}(Q)$), then $A_i(Q') \subseteq A_i(Q)$, $B_i(Q') \subseteq B_i(Q)$, $C_i(Q') \supseteq C_i(Q)$ and $D_i(Q') \supseteq D_i(Q)$. (We will only use $D_i(Q)$ in the even case.)

We first describe the situation where $\xi(\Sigma)$ is odd. We revert to the terminology of curves and standard curves, etc., as in Section 8.

Lemma 13.4. *Suppose that Σ is odd, and that $Q \subseteq \mathcal{E}^*(\Sigma)$ is a ξ_0 -cube. Then there is a (non-standard) good multicurve, τ , such that for each $i \in \{1, \dots, \xi_0\}$, there is some $Y_i \in \mathcal{U}\mathcal{X}_N(\tau)$ such that $A_i(Q) = C_i(Q) = \{Y_i\}$.*

Note that τ is determined by Q , and we write $\tau = \tau(Q)$.

Proof. This is a repeat of the argument of Section 8, but now a lot simpler. By Lemma 6.3, we have a \mathcal{Z} -sequence of h_0 -quasimorphisms, $\phi_\zeta : Q \rightarrow \mathcal{E}(\Sigma)$ with $\phi_\zeta x \rightarrow x$ for all $x \in Q$. By Lemma 13.3, we have standard good multicurves, τ_ζ , with $\mathcal{X}_N(\tau) = \{Y_{1, \zeta}, \dots, Y_{\xi_0, \zeta}\}$ and with $\mathcal{A}(\phi_\zeta, i) = \{Y_{i, \zeta}\}$. This gives us a multicurve τ and surfaces Y_i with $\tau_\zeta \rightarrow \tau$, $Y_{i, \zeta} \rightarrow Y_i$ and with $\mathcal{U}\mathcal{X}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$. One checks easily that $A_i(Q) = \{Y_i\}$. To see that $Y_i \in C_i(Q)$, note that if $Q' \subseteq Q$ is a

smaller ξ_0 -cube, then $A_i(Q') = A_i(Q)$, so applying the above to Q' , we must have $A_i(Q') = A_i(Q) = \{Y_i\}$, and the result follows easily. \square

Note that it also follows that if Q, Q' are ξ_0 -cubes with Q' bigger than Q , then $\tau(Q) = \tau(Q')$.

Suppose now that $\tau \subseteq \mathcal{UG}^0$ is a multicurve. We have maps $\psi_{\tau_\zeta} : \mathcal{M}(\Sigma) \rightarrow \mathcal{T}(\tau_\zeta)$, $v_{\tau_\zeta} : \mathcal{T}(\tau_\zeta) \rightarrow T(\tau_\zeta)$ and $\omega_{\tau_\zeta} = v_{\tau_\zeta} \circ \psi_{\tau_\zeta} : \mathcal{M}(\Sigma) \rightarrow T(\tau_\zeta)$. Similarly as in Section 8, these are all uniformly coarsely lipschitz quasimorphisms, and so give rise to maps, $\psi_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow \mathcal{T}^*(\tau)$, $v_\tau^* : \mathcal{T}^*(\tau) \rightarrow T(\tau)$ and $\omega_\tau^* = v_\tau^* \circ \psi_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow T^*(\tau)$. From Lemma 13.1, we see that $\omega_\tau^* : \mathcal{M}^*(\Sigma) \rightarrow T^*(\tau)$ is the gate map of $\mathcal{M}^*(\Sigma)$ to $T^*(\tau)$.

We see that $T^*(\tau)$ is a closed convex subset of \mathcal{E}^* , median isomorphic to a direct product of copies of $\mathcal{E}^*(X)$ for $X \in \mathcal{UX}_N(\tau)$. If τ is good, then each $\mathcal{E}^*(X)$ has rank 1, hence is an \mathbb{R}^* -tree. (We do not need that Σ is even for this.)

Restricting to $\mathcal{E}^\infty(\Sigma)$, let $T^\infty(\tau) = T^*(\tau) \cap \mathcal{E}^\infty(\Sigma)$. Note that (\mathcal{E}^*, ρ^*) admits a bilipschitz equivalent metric ρ' such that (\mathcal{E}^*, ρ') is a median metric space. If this is non-empty, then it is a direct product of \mathbb{R} -trees in this metric. Note that, by Proposition 4.6, its structure as a median algebra, hence its decomposition into factors, is completely determined by its topology (see also [Bo5]). Note that if τ is good, then $T_E^\infty(\tau)$ is a tree product in the terminology introduced at the end of Section 4.

Lemma 13.5. *Suppose that $\xi(\Sigma)$ is odd. Suppose that $\tau \subseteq \mathcal{UG}^0(\Sigma)$ is a good multicurve, and that $Q \subseteq T_E^*(\tau)$ is a ξ_0 -cube. Then $\tau(Q) = \tau$.*

Proof. Note that $\text{hull}(Q) \subseteq T_E^*(\tau)$, and we can identify this set with $\prod_i \mathcal{E}^*(Y_i)$ via v_τ^* , where Y_1, \dots, Y_{ξ_0} are the complementary components of τ . Thus, $\text{hull}(Q) = \prod_i I_i$, where each I_i is a non-trivial interval in $\mathcal{E}^*(Y_i)$. In particular, we see that $\theta_{Y_i}|_{I_i}$ is injective, and so $Y_i \in C_i(Q)$. It follows by Lemma 13.4 that $C_i(Q) = \{Y_i\}$, and so $\tau(Q) = \tau$. \square

Note that it follows for an odd surface that if τ and τ' are good multicurves and $T_E^*(\tau) \cap T_E^*(\tau')$ contains a ξ_0 -cube, then $\tau = \tau'$.

Lemma 13.6. *If Σ is odd, and $Q \subseteq \mathcal{E}^*(\Sigma)$ is a ξ_0 -cube, then $Q \subseteq T^*(\tau(Q))$.*

Proof. Write $\tau = \tau(Q)$. Similarly as with the proof of Lemma 8.7, we see that the map $\psi_\tau|_{\text{hull}(Q)} : \text{hull}(Q) \rightarrow T^*(\tau)$ is injective. (Since for each $X \in \mathcal{UX}_N(\tau)$, $\psi_X = \theta_X$ is injective on each 1-dimensional face of $\text{hull}(Q)$.) Thus, by Corollary 3.10, we get $Q \subseteq \text{hull}(Q) \subseteq T^*(\tau)$. \square

It also follows that if $\Phi \subseteq \mathcal{E}^*$ is a flat, then $\Phi \subseteq T^\infty(\tau(\Phi))$.

The following few statements (13.7 to 13.12) will be valid for both odd and even surfaces, so we restrict proofs for the moment to the odd case.

Lemma 13.7. *Suppose that $T \subseteq \mathcal{E}^\infty(\Sigma)$ is a tree product. Then there is a good multicurve $\tau \subseteq \mathcal{UG}^0(\Sigma)$ such that $T \subseteq T_E^\infty(\tau)$.*

Proof. (For Σ odd.) Let $Q \subseteq T$ be any ξ_0 -cube in T , and set $\tau = \tau(Q)$. If $x \in T$, then there are ξ_0 -cubes P and Q' , with Q, Q' both bigger than P , and with $x \in Q'$. Now $\tau(Q') = \tau(P) = \tau(Q) = \tau$, and so by Lemma 13.6, $x \in Q' \subseteq T_E^\infty(\tau)$. Thus $T \subseteq T_E^\infty(\tau)$. \square

Note that (by Lemma 13.5), if Σ is odd, then τ is unique.

Corollary 13.8. *If τ is a good multicurve, then $T_E^\infty(\tau)$ is a maximal tree product.*

Proof. (For Σ odd.) Let $T \supseteq T_E^\infty(\tau)$ be a tree product. By Lemma 13.7, $T \subseteq T_E^\infty(\tau')$ for some τ' . It follows that $\tau = \tau'$ and that $T = T_E^\infty(\tau)$. \square

Putting the above together, we see that a (closed) subset of $\mathcal{E}^\infty(\Sigma)$ is a maximal tree product if and only if it has the form $T_E^\infty(\tau)$ for some good multicurve τ such that $T_E^\infty(\tau) \neq \emptyset$. In particular, each factor must be a furry tree. Moreover, by Proposition 4.6, a closed subset of $\mathcal{E}^\infty(\Sigma)$, homeomorphic to a product of furry trees is a tree product. Therefore the collection of maximal tree products is determined topologically (as the collection of maximal products of furry trees topologically embedded in $\mathcal{E}^\infty(\Sigma)$). We deduce (in the odd case):

Lemma 13.9. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi(\Sigma) = \xi(\Sigma')$ and that $f : \mathcal{E}^\infty(\Sigma) \rightarrow \mathcal{E}^\infty(\Sigma')$ is a homeomorphism. Then if $\tau \subseteq \mathcal{UG}^0(\Sigma)$ is a good multicurve, there is a unique good multicurve, $\tau' \subseteq \mathcal{UG}^0(\Sigma')$, such that $f(T_E^\infty(\tau)) = T_E^\infty(\tau')$.*

We can now use this to show that subsets of $\mathcal{E}(\Sigma)$ of the form $\mathcal{T}_E(\tau)$ where τ is a standard multicurve are determined by the coarse metric structure of $\mathcal{E}(\Sigma)$. In what follows, all curves, multicurves and subsurface are assumed standard.

Suppose that τ and τ' are good multicurves with $\tau \cap \tau' = \emptyset$. Then, we can find $c \in T_E(\tau)$ (unique up to bounded distance) such that $\theta_X c \sim \theta_X \tau'$ for all $X \in \mathcal{X}_N(\tau)$. (Take subsurface projection of τ' to each such X and assemble these curves together with τ itself to give a complete multicurve, $c \supseteq \tau$.) Note that if $b \in T_E(\tau')$, then $\theta_X b \sim \theta_X \tau'$ for all such X , and so, by the definition of ω_τ , it follows that $\omega_\tau b \sim c$. Thus $\omega_\tau(T_E(\tau'))$ has bounded diameter.

We can now argue similarly as in Section 7 (with \mathcal{E} replacing \mathcal{M} , and \mathcal{X}_N replacing \mathcal{X}). If $a \in T_E(\tau)$, $b \in T_E(\tau')$, $c \in \omega_\tau(T_E(\tau'))$ and $d \in \omega_{\tau'}(T_E(\tau))$, then $\sigma_X(a, b) \sim \sigma_X(a, c) + \sigma_X(c, d) + \sigma_X(d, b)$ for all $X \in \mathcal{X}_N$, and so $\rho(a, b)$ agrees with $\rho(a, c) + \rho(c, d) + \rho(d, b)$ up to linear bounds. We can now deduce, by essentially the same argument as Lemma 7.17, that (if Σ is odd):

Lemma 13.10. *There are uniform constants, $k, t \geq 0$ such that if τ, τ' are good multicurves, with $\tau \neq \tau'$, $x \in T_E(\tau')$, and $r \geq 0$, then there is some $y \in T_E(\tau)$ with $\rho(x, T_E(\tau)) \geq r$ and $\rho(x, y) \leq kr + t$.*

We can now proceed as in Section 9 to deduce:

Lemma 13.11. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi_0(\Sigma) = \xi_0(\Sigma') \geq 2$. Suppose that $\phi : \mathcal{M}(\Sigma) \rightarrow \mathcal{M}(\Sigma')$ is a quasi-isometry. Then, given any good multicurve, τ , in Σ , there is a good multicurve τ' in Σ' such that $\text{hd}(T_E(\tau'), \phi T_E(\tau))$ is bounded above by some constant depending only on ξ and the parameters of ϕ .*

Proof. (If Σ and Σ' are odd.) This proceeds exactly as with Lemma 7.5 using Lemma 7.3. For hypothesis (S2), we use Lemma 13.10 (instead of Lemma 7.17), and for hypothesis (S3), we use Lemma 13.9 (instead of Lemma 8.13). The reason why the bound can be assumed to depend only on the parameters of ϕ is explained in the last paragraph of Section 7. \square

Recall that $T_E(\tau)$ is, up to bounded distance, the image of a quasi-isometric embedding, $v_\tau : T_E(\tau) \rightarrow \mathcal{E}(\Sigma)$, where $T_E(\tau) = \prod_{X \in \mathcal{X}_N(\tau)} \mathcal{E}(X)$. We can elaborate on Lemma 13.11.

Proposition 13.12. *Suppose that Σ, Σ', ϕ are as in Lemma 13.11. Then, given any good multicurve, τ , in Σ , there is a good multicurve τ' in Σ' , a bijection, $\pi : \mathcal{X}_N(\tau) \rightarrow \mathcal{X}_N(\tau')$, and a quasi-isometry, $\phi_X : \mathcal{E}(X) \rightarrow \mathcal{E}(\pi(X))$ for each $X \in \mathcal{X}_N(\tau)$, such that the maps $v_\tau \circ (\prod_X \phi_X)$ and $\phi \circ v_{\tau'} : T_E(\tau) \rightarrow \mathcal{E}(\Sigma')$ agree up to bounded distance. The bound, and the parameters of the maps ϕ_X depend only on ξ and the parameters of ϕ .*

Proof. (If Σ is odd.) By Lemma 13.11, we see that $\phi|_{T_E(\tau)}$ is a bounded distance from quasi-isometry from $T_E(\tau)$ to $T_E(\tau')$. Thus, via the quasi-isometric embedding v_τ and $v_{\tau'}$, we get a quasi-isometry $\hat{\phi} : \mathcal{T}_E(\tau) \rightarrow \mathcal{T}_E(\tau')$. Now each of the factors of $\mathcal{T}_E(\tau)$ and $\mathcal{T}_E(\tau')$ is a bushy hyperbolic space (in this case, a quasitree). Here ‘‘bushy’’ means that every point is a bounded distance from the centre of an ideal uniformly quasigeodesic triangle. It therefore follows from [KaKL] (see

also [Bo5]) that, up to bounded distance and permutation of factors, $\hat{\phi}$ splits as a product of quasi-isometries of the factors. (To apply the result of [KaKL] as stated one needs to observe, in addition, that each of the factors admits a cobounded isometric action. However, bushy is all that is really required. See [Bo5] for further discussion of this.) \square

We move on to consider the even case.

First we need the following variation on Lemma 13.3:

Lemma 13.13. *Suppose that ξ is even. Then there is a good multicurve, τ_0 , with the following property. For each i , either there is some $Y_i \in \mathcal{X}_N(\tau_0)$ which is an $S_{1,1}$ or $S_{0,4}$ such that $\mathcal{A}(\phi, i) = \{Y_i\}$, or else every element of $\mathcal{A}(\phi, i)$ lies in $W(\tau_0)$.*

Note that in the latter condition occurs for precisely one index, say $i_0 \in \{1, \dots, \xi_0\}$.

Lemma 13.13 is proven by a similar argument to Lemma 13.3.

We split this into the following cases:

Case (1a): $\mathcal{A}(\phi, i_0) = \emptyset$.

Case (1b): The elements of $\mathcal{A}(\phi, i_0)$ fill $W(\tau)$.

Case (2): $\mathcal{A}(\phi, i_0)$ consists of a single element which is of complexity-1.

Note that in Case (1), the multicurve, τ , is canonically determined, and we write it as $\tau_1(\phi)$. In case (2) write $\mathcal{A}(\phi, i_0) = \{Y_{i_0}\}$. Now $Y_{i_0} \prec W(\tau)$ has one boundary component, α , which is non-peripheral in $W(\tau)$. Set $\tau_2(\phi) = \tau \cup \{\alpha\}$. Thus, $\mathcal{A}(\phi, i) = \{Y_i\}$ for all $i \in \{1, \dots, \xi_0\}$. In this case, $\tau_2(\phi)$ is almost good and canonically determined.

We now pass again to the asymptotic cone. The following two lemmas (13.14 and 13.15) are valid in general.

Lemma 13.14. *Suppose $a, b, c, d \in \mathcal{E}^*(\Sigma)$, with $c \in [a, d]$ and $b \in [a, c]$, with $\theta_X a \neq \theta_X b$, $\theta_X c \neq \theta_X d$ and $\theta_Y c \neq \theta_Y d$. Then either $X = Y$ or $X \pitchfork Y$.*

Proof. This is by the same argument as Lemma 8.2. \square

Corollary 13.15. *Suppose $Q \subseteq \mathcal{E}^*(\Sigma)$ is a ξ_0 -cube, and that $i \in \{1, \dots, \xi_0\}$. If $X \in C_i(Q)$, $Y \in A_i(Q)$, then $X = Y$ or $X \pitchfork Y$.*

Proof. Let a, d be an i th side of Q . Choose $b, c \in [a, d] \setminus \{a, d\}$, and apply Lemma 13.14. \square

We say that a multicurve, τ , in an even surface, Σ , is *almost good* if its complement has exactly ξ_0 complexity-1 components and one $S_{0,3}$ component. (This implies that every curve in τ separates Σ .) We denote the $S_{0,3}$ component by $Z(\tau)$.

Note that such a muticurve can be obtained by taking a good multicurve τ_0 , and adding a curve to split $W(\tau_0)$ into a complexity-1 surface and an $S_{0,3}$. As with good multicurves, the dual graph is a tree.

Lemma 13.16. *Suppose that Σ is even and that $Q \subseteq \mathcal{E}^*(\Sigma)$ is a ξ_0 -cube in $\mathcal{E}^*(\Sigma)$. Then it is one of the following two types.*

Type (1): There is a good multicurve τ_1 in Σ and some $i_0 \in \{1, \dots, \xi_0\}$ such that for all $i \neq i_0$, there is some $Y_i \in \mathcal{UX}_N(\tau_1) \setminus \{W(\tau_1)\}$ such that $C_i(Q) = \{Y_i\}$ and such that $C_{i_0}(Q) \subseteq \{W(\tau_1)\}$ and $D_{i_0}(Q) \cap \mathcal{UX}_N(\tau_1) = \{W(\tau_1)\}$.

Type (2): There is an almost good multicurve, τ_2 , in Σ such that for all $i \in \{1, \dots, \xi_0\}$, there is some $Y_i \in \mathcal{UX}_N(\tau_2)$ with $C_i(Q) = \{Y_i\}$.

Note that on both cases, the Y_i are all distinct, and account for all of the complexity-1 components of the complement of τ_1 or of τ_2 . Also, by Corollary 13.15, we have $A_i(Q) = \{Y_i\}$, provided $i \neq i_0$ in case (1). Note that the two cases are mutually exclusive, and that τ_1 or τ_2 are uniquely determined. We write them as $\tau_1(Q)$ and $\tau_2(Q)$. (As usual in this context, we are referring to non-standard multicurves and subsurfaces here.)

Proof. The argument proceeds as with Lemma 13.4. We can assume that the type of $\phi_\zeta(Q)$ is constant. Type (2) gives us Type (2) here. If they are all of Type (1), then we get a good multicurve τ_1 in Σ . If there is some $Y \preceq W(\tau)$ with $Y \in C_j(Q)$ for some j , then we add a boundary curve of Y to τ_1 to give us an almost good multicurve τ_2 , and set $Y_j = Y$. Note that $Y_j \in \mathcal{UX}_N(\tau_2)$ and that, by Corollary 13.15, $C_j(Q) = \{Y_j\}$. \square

Lemma 13.17. *Suppose that Σ is even, and that $Q \subseteq \mathcal{E}^\infty(\Sigma)$ is a ξ_0 -cube.*

(1) If Q is of Type (1), then $Q \subseteq T_E^\infty(\tau_1(Q))$.

(2) If Q is of Type (2), and $\tau \subseteq \tau_2(Q)$ is a good multicurve, then $Q \subseteq T_E^\infty(\tau)$.

Note that in case (2), τ can be obtained from τ_2 by removing any one of the boundary curves of $Z(\tau)$ (the $S_{0,3}$ component of $\Sigma \setminus \tau_2$), provided this curve is non-peripheral in Σ .

Proof. (1) For each i , $Y_i \in C_i(Q) \subseteq D_i(Q)$, so ψ_{Y_i} is injective on the i th face of $\text{hull}(Q)$, and the argument proceeds as with Lemma 13.6.

(2) Again, this follows by the same argument, using $W(\tau)$ instead of the element Y_j of $\mathcal{UX}_N(\tau_2)$ contained in $W(\tau)$. Since $Y_i \in C_j(Q)$, we have $W(\tau) \in D_j(Q)$. \square

Suppose that $T \subseteq \mathcal{E}^\infty(\Sigma)$ is a tree product. Write $T = \prod_{i=1}^{\xi_0} \Delta_i$, where Δ_i are the \mathbb{R} -tree factors of T , as a median algebra. If $X \in \mathcal{UX}_N$, then we have well defined maps $\theta_X^\infty|_{\Delta_i} : \Delta_i \rightarrow \mathcal{G}^\infty(X)$ and $\psi_X^\infty|_{\Delta_i} : \Delta_i \rightarrow \mathcal{E}^\infty(X)$.

Lemma 13.18. *Suppose that Σ is even, and that $T \subseteq \mathcal{E}^\infty(\Sigma)$ is a tree product. Then there is a good multicurve, τ , and some $i_0 \in \{1, \dots, \xi_0\}$ such that for all $i \neq i_0$, there is a (unique) $Y_i \in \mathcal{UX}_N(\tau) \setminus \{W(\tau)\}$ such that $\theta_{Y_i}|_{\Delta_i}$ is injective, and such that $\psi_{W(\tau)}|_{\Delta_{i_0}}$ is injective.*

(Note that it is possible that there is some (unique) complexity-1 surface $Y \prec W(\tau)$ with $\theta_Y|_{\Delta_j}$ injective for some j . In this case, it would be natural to add a boundary curve of Y to τ to obtain an almost good multicurve τ_2 , and set $Y_j = Y \in \mathcal{UX}_N(\tau_2)$. We will not need this here however.)

Proof. Suppose first that all of the ξ_0 -cubes in T are of Type (2). If Q_0, Q are ξ_0 -cubes in T with Q bigger than Q_0 , then we see that $\tau_2(Q) = \tau_2(Q_0)$. It follows that $\tau_2(Q) = \tau_2$ is constant for all ξ_0 -cubes, $Q \subseteq T$. (For if Q, Q' are such, then we can find ξ_0 -cubes, $Q_0, Q'_0, Q'' \subseteq T$ with Q, Q'' bigger than Q_0 and with Q', Q'' bigger than Q'_0 .) Moreover, we can index the elements of $\mathcal{UX}_N(\tau_2)$ consistently as $Y_1 \dots Y_{i_0}$, so that Y_i the i th side of any ξ_0 -cube $Q \subseteq T$ is (parallel to) an interval in Δ_i . Now, if $a, b \in \Delta_i$ are distinct, let Q be any ξ_0 -cube with i th side $\{a, b\}$. Since $Y_i \in C_i(Q)$, we get $\theta_{Y_i}a \neq \theta_{Y_i}b$. In other words, $\theta_{Y_i}|_{\Delta_i}$ is injective. We can now (arbitrarily and somewhat artificially) remove one of the non-peripheral boundary curves of $Z(\tau_2)$ to give τ . If Y_{i_0} is the element lying in $W(\tau)$, then since $\theta_{Y_{i_0}}|_{\Delta_{i_0}}$ is injective, so is $\psi_{W(\tau)}|_{\Delta_{i_0}}$.

We can therefore assume that there is a Type (1) cube, $Q \subseteq T$. Let $\tau = \tau_1(Q)$. We claim that if $i \neq i_0$, then $\theta_{Y_i}|_{\Delta_i}$ is injective. For suppose $a, b \in \Delta_i$ are distinct. Let Q' be the ξ_0 -cube with i th side parallel to $\{a, b\}$, and all other sides parallel to the corresponding sides of Q . Thus, for all $j \neq i$, we have $C_j(Q') = C_j(Q)$, so $C_j(Q) = \{Y_j\}$ for $j \neq i, i_0$ and $C_{i_0}(Q) \subseteq \{W(\tau)\}$. It therefore follows that Q' must also be of Type (1) with $\tau_1(Q') = \tau$, and so $C_i(Q') = \{Y_i\}$. In particular, $\theta_{Y_i}a \neq \theta_{Y_i}b$ as claimed. Finally, we claim that $\psi_{W(\tau)}|_{\Delta_{i_0}}$ is injective. For suppose that $c, d \in \Delta_{i_0}$ are distinct. Let Q'' be any ξ_0 -cube with i_0 th side parallel to $\{c, d\}$. If $i \neq i_0$, then (since $\theta_{Y_i}|_{\Delta_i}$ is injective) we have $Y_i \in C_i(Q'')$. It follows that Q'' is either of Type (1) with $\tau_1(Q'') = \tau$, or Type (2) with $\tau_2(Q'') \supseteq \tau$. Either way we get $\psi_{W(\tau)}c \neq \psi_{W(\tau)}d$ as required. \square

Note that it follows that any ξ_0 -cube in T is either of Type (1) with $\tau_1(Q) = \tau$, or of Type (2) with $\tau \subseteq \tau_2(Q)$. Thus, either way, by Lemma

13.17 we get $Q \subseteq T_E^\infty(\tau)$. Since every element of T lies in such a cube, we get that $T \subseteq T_E(\tau)$. Therefore Lemma 13.7 in the even case is an immediate consequence.

Note also that if τ, τ' are distinct good multicurves, it is easily seen that we cannot have $T_E^\infty(\tau) \subseteq T_E^\infty(\tau')$, and so Corollary 13.8 follows in the even case.

We now proceed with Lemmas 13.9, 13.10 and 13.11, and deduce Proposition 13.12 in the even case. In fact, this also applies if say, Σ is even and Σ' is odd. This therefore completes the proof of Proposition 13.12 in general.

Note that retrospectively, this can be used to distinguish odd and even surfaces from the quasi-isometry type of $\mathcal{E}(\Sigma)$. In the odd case, all the factors of any $T_E(\tau)$ are quasitrees, whereas in the even case, the factor corresponding to the complexity-2 surface is not a quasitree. We have already seen that $\xi_0(\Sigma)$ is determined by $\mathcal{E}(\Sigma)$ (by Theorem 12.6). Together with the parity of ξ , this determines ξ . In other words, we have:

Proposition 13.19. *Suppose that Σ, Σ' are compact surfaces with $\mathcal{E}(\Sigma)$ quasi-isometric to $\mathcal{E}(\Sigma')$, then $\xi(\Sigma) = \xi(\Sigma')$.*

Proof. This follows by the above, if $\xi(\Sigma), \xi(\Sigma') \geq 3$. Note that such spaces are never hyperbolic. If $\xi(\Sigma) = 2$, then $\mathcal{E}(\Sigma)$ is hyperbolic but not a quasitree. If $\xi(\Sigma) = 1$, then $\mathcal{E}(\Sigma)$ is a quasitree. This deals with all cases. \square

We will give a much stronger statement later (see Theorem 14.7).

We now apply this to identify certain subsurfaces of Σ in terms of the coarse geometry of $\mathcal{E}(\Sigma)$. The discussion applies to both the odd and even cases, though with a few differences.

Given a multicurve, τ , write $G(\tau) \leq \text{Map}(\Sigma)$ which preserves each component of τ . If $X \in \mathcal{X}$, write $G(X) \leq \text{Map}(\Sigma)$ for the subgroup supported on X .

Recall the notation from Section 10. We write $\mathcal{B} = \mathcal{B}(\mathcal{E}(\Sigma))$ for the set of subsets of $\mathcal{E}(\Sigma)$ defined up to finite Hausdorff distance. Any subgroup, $G \leq \text{Map}(\Sigma)$, determines an element $B(G) \in \mathcal{B}$, namely the class of any orbit of G in $\mathcal{E}(\Sigma)$. We will abbreviate $B(\tau) = B(G(\tau))$ and $B(X) = B(G(X))$. Note that $G(\tau)$ acts coboundedly on $T_E(\tau)$, and so $B(\tau)$ is just the class of $T_E(\tau)$. Also, if X is an annulus, the $B(X)$ is just the class of bounded subsets.

Clearly if $X \preceq Y$, then $G(X) \leq G(Y)$, so $B(X) \leq B(Y)$. If $X \in \mathcal{X}_N(\tau)$, then $G(X) \leq G(\tau)$ so $B(X) \leq B(\tau)$.

Lemma 13.20. *If $X, Y \in \mathcal{X}_N$ with $B(X) \leq B(Y)$, then $X \preceq Y$.*

Proof. Let $a \in \mathcal{E}(\Sigma)$ be any pants decomposition containing the relative boundary of Y in Σ . If X is not contained in Y , then there is some component, α , of a , disjoint from Y (or peripheral in Y) which crosses or is contained in X . If $h \in G(Y)$, then $h\alpha = \alpha$. Write $\beta = \theta_X \alpha \in \mathcal{G}^0(X)$. Let $g \in G(X)$ be any pseudoanov in X . Then $\sigma_X(\beta, g^n \beta) \rightarrow \infty$ as $n \rightarrow \infty$. Thus, $\sigma_X(a, g^n a) \rightarrow \infty$. On the other hand, since α is a component of both a and ha , we have $\sigma_X(a, ha)$ bounded for all $h \in G(X)$. Thus, $\sigma_X(g^n, G(Y)a) \rightarrow \infty$. But $\rho(g^n, G(Y)a)$ is linearly bounded below by projection to X , and so $\rho(g^n a, G(Y)a) \rightarrow \infty$, contradicting the assumption that $B(X) \leq B(Y)$. \square

In particular, it follows that if $B(X) = B(Y)$, then $X = Y$.

Definition. We say that a non-empty subset, $\mathcal{Y} \subseteq \mathcal{X}_N$, is *compatible* if there is a good multicurve, τ , such that $\mathcal{Y} \subseteq \mathcal{X}_N(\tau)$. A subsurface, $X \in \mathcal{X}_N$ is *admissible* if $\{X\}$ is compatible.

If Σ is odd, then X is admissible if and only if either it is an $S_{1,1}$ or it is an $S_{0,4}$ and each component of the complement is even and meets X exactly one curve. If Σ is even, then X is admissible if and only if it is one of the following: an $S_{1,1}$, or an $S_{0,4}$ with all but one of the complementary components even and all meeting X is a single curve; or thirdly it is an $S_{1,2}$ or $S_{0,5}$ with all complementary components even and meeting X in a single curve. One can give a similar description of compatibility.

Note that maximal compatible sets are in bijective correspondence to good multicurves.

We claim that we can recognise compatibility in terms of the coarse geometry of $\mathcal{E}(\Sigma)$.

Lemma 13.21. *Suppose that Σ, Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 3$. Suppose that $\phi : \mathcal{E}(\Sigma) \rightarrow \mathcal{E}(\Sigma')$ is a quasi-isometry. If X is an admissible subsurface of Σ , then there is a unique admissible subsurface, πX , in Σ' such that $B(\pi X) = \phi B(X)$. Moreover, π is a bijection between the set of admissible subsurfaces of Σ and the set of admissible surfaces of Σ' . It preserves the complexity of the surface. Moreover, a set, \mathcal{Y} , of admissible surfaces of Σ is compatible, if and only if their images, $\pi \mathcal{Y}$, are compatible in Σ' .*

Proof. This is just putting together Proposition 13.12 and Lemma 13.20. \square

Note that, if τ is a good multicurve, then by construction, the map π of Lemma 13.21, when restricted to $\mathcal{X}_N(\tau)$, agrees with the map π given by Proposition 13.12.

Definition. A *terminal subsurface* is a subsurface X , of Σ which is either an $S_{1,1}$ or else an $S_{0,4}$ with all but one boundary components peripheral in Σ .

In other words, it is a complexity-1 surface cut off by a single curve in Σ . We will refer to such a curve as *1-separating* (see the terminology in Section 14).

Note that any terminal surface is admissible, and that two terminal surfaces (or equivalently the corresponding 1-separating curves) are disjoint, if and only if they are compatible.

Lemma 13.22. *With the hypotheses of Lemma 13.21, X is terminal in Σ if and only if πX is terminal in Σ' .*

Proof. We can recognise terminal surfaces among admissible surfaces from the properties already verified for π .

Suppose that X is admissible. We have already noted that we can distinguish complexity, so we can assume X to be complexity-1. If X is terminal, then we can find another admissible surface Y , compatible with X , with the property that if \mathcal{Y} is any maximal compatible family containing Y , then \mathcal{Y} also contains X . (We take Y to be any admissible surface meeting X in its boundary.) Conversely, if X is admissible and there is such a Y , then X must be terminal. For if not, it must be a $S_{0,4}$ with $\Sigma \setminus X$ disconnected. Suppose Y and \mathcal{Y} are as given. Let Z be the component of $\Sigma \setminus Y$ containing X . Now X must be strictly contained in Z , so after applying some element of $G(Z)$ to \mathcal{Y} if necessary, we can certainly arrange that \mathcal{Y} does not contain X .

Note that the above criterion makes reference only to complexity, admissibility and compatibility of subsurfaces, and is hence preserved by π . \square

We can express this as follows. Let $\mathcal{G}_1(\Sigma)$ be the full subgraph of the curve graph $\mathcal{G}(\Sigma)$ with vertex set the set of all 1-separating curves in Σ , then π gives rise to an isomorphism from $\mathcal{G}_1(\Sigma)$ to $\mathcal{G}_1(\Sigma')$.

We will elaborate on this in Section 14, and explain how, for certain surfaces at least, it can be used to deduce quasi-isometric rigidity of the pants graph.

14. VARIATIONS ON THE CURVE GRAPH

Let Σ be a compact surface with $\xi(\Sigma) \geq 3$. Recall that $\mathcal{G}(\Sigma)$ is the curve graph, and $\text{Map}(\Sigma)$ acts cofinitely on $\mathcal{G}(\Sigma)$. Given a subset, $C \subseteq V(\Gamma)$, we write $\mathcal{G}(\Sigma, C)$ for the full subgraph of $\mathcal{G}(\Sigma)$ with vertex set C . Clearly, if C is $\text{Map}(\Sigma)$ -invariant, then $\text{Map}(\Sigma)$ acts on $\mathcal{G}(\Sigma, C)$.

Definition. We say that $\mathcal{G}(\Sigma, C)$ is *rigid* if $\text{Map}(\Sigma)$ is the full automorphism group of $\mathcal{G}(\Sigma, C)$.

(This is often expressed in terms of the curve complex, rather than the curve graph, but since the curve complex is the flag complex with 1-skeleton $\mathcal{G}(\Sigma)$, these notions are equivalent.)

It is natural to ask for which $\text{Map}(\Sigma)$ -invariant subsets, C , the graph is $\mathcal{G}(\Sigma, C)$ rigid. It was shown in [Iv, Ko, L] that $\mathcal{G}(\Sigma)$ itself is rigid for all but finitely many Σ . In fact, $\mathcal{G}(S_{g,p})$ is rigid if $g \geq 2$ or ($g = 1$ and $p \geq 2$) or ($g = 0$ and $p \geq 6$). This is also an immediate consequence of the result of [Sha] given as Theorem 9.1 here.

Various other cases are known. For example for rigidity for (most of) the non-separating curve graphs was established in [Ir]. Of particular interest here however is when C is the set of all separating curves. In this case, we write $\mathcal{G}_s(\Sigma) = \mathcal{G}(\Sigma, C)$. It follows from [BreM, Ki] that this is also rigid for all but finitely many surfaces (if $g \geq 1$). Clearly, if $g = 0$, then $\mathcal{G}_s(\Sigma) = \mathcal{G}(\Sigma)$, and this case was dealt with in [Ko] and independently in [L]. Therefore, combining these we get:

Theorem 14.1. [Ko, L, BreM, Ki] $\mathcal{G}_s(S_{g,p})$ is rigid if $g \geq 3$ or ($g = 2$ and $p \geq 2$) or ($g = 1$ and $p \geq 3$) or ($g = 0$ and $p \geq 6$).

Note that, under the above conditions, we see that the isomorphism class of the graph $\mathcal{G}(\Sigma, C)$ determines Σ — since it determines $\text{Map}(\Sigma)$ up to isomorphism, hence Σ (see, for example, [RaS] for a direct proof of the last implication).

One can reduce certain other cases to this.

Given a separating curve γ , write $X^-(\gamma), X^+(\gamma)$ for the complementary components (as usual defined up to isotopy). Let $\kappa(\gamma) = \max\{\xi(X^-(\gamma)), \xi(X^+(\gamma))\}$.

Definition. We say that a curve γ is *n-separating* if it is separating and $\kappa(\gamma) = n$. We say that γ is *(n+)-separating* if it is separating and $\kappa(\gamma) \geq n$.

Thus, for example, a curve is 0-separating if cuts off a $S_{0,3}$. Similarly it is 1-separating if it cuts off an $S_{1,1}$ or $S_{0,4}$ (as we defined this term in Section 13). In this case, we write $F(\gamma)$ for the surface cut off by γ . (If $\xi(\Sigma) \geq 4$, this is well defined.)

We write C_n and C_{n+} , respectively, for the sets of n -separating and $(n+)$ -separating curves, and write $\mathcal{G}_n(\Sigma) = \mathcal{G}(\Sigma, C_n)$ and $\mathcal{G}_{n+}(\Sigma) = \mathcal{G}(\Sigma, C_{n+})$. (We could write $\mathcal{G}_s(\Sigma) = \mathcal{G}_{0+}(\Sigma)$ in this notation.)

Recall that $\mathcal{G}_1(\Sigma)$ was the graph introduced at the end of Section 13, where we saw that it is determined by the coarse geometry of $\mathcal{E}(\Sigma)$. We can partition C_1 as $C_{1HT} \sqcup C_{4HS}$, depending on whether the curve

bounds an $S_{1,1}$ or an $S_{0,4}$. Note that this partition is not a-priori deemed part of the structure of $\mathcal{G}_1(\Sigma)$. It can however be recovered, at least in most cases, as we show in Lemma 14.3.

In what follows we write $\mathcal{G}^c(\Sigma)$ for the complementary graph of $\mathcal{G}(\Sigma)$; that is, with the same vertex set and complementary edge set. If C is any set of curves, we write $\mathcal{G}^c(\Sigma, C)$ for the full subgraph of $\mathcal{G}^c(\Sigma)$ with vertex set C . Clearly this is complementary to $\mathcal{G}(\Sigma, C)$. We write $\mathcal{G}_1^c(\Sigma) = \mathcal{G}^c(\Sigma, C_1)$.

Given a subsurface X of Σ , we write $P(X) = \{\gamma \in C_1 \mid \gamma \prec X\}$. We say that X is *big* if $\xi(X) \geq 2$ and $P(X) \neq \emptyset$. (Note that the latter condition is redundant if X has only one relative boundary component in Σ .) Clearly, $P(X)$ is invariant under the subgroup, $G(X)$ of $\text{Map}(\Sigma)$, supported on X . From this one sees easily that if X is big, then $P(X)$ is infinite, and the elements of $P(X)$ fill X . Moreover, $\mathcal{G}^c(\Sigma, P(X))$ is connected (of diameter 2).

By a *division* of Σ , we mean an ordered pair, $\underline{X} = (X^-, X^+)$, where X^- and X^+ are big subsurfaces of Σ which can be realised disjointly so that $\Sigma \setminus (X^- \cup X^+)$ is a disjoint union of (non-peripheral) annuli. In other words, it is equivalent to a transversely oriented multicurve in Σ which separates Σ into two big subsurfaces.

Given a subset $P \subseteq C_1$, write $L(P)$ for the set of elements of $C_1 \setminus P$ which are adjacent to some element of C_1 in $\mathcal{G}_1^c(\Sigma)$ (in other words the “1-sphere” about P in $\mathcal{G}_1^c(\Sigma)$). By a *division* of $\mathcal{G}_1^c(\Sigma)$, we mean an ordered pair, $\underline{P} = (P^-, P^+)$, of disjoint infinite subsets $P^-, P^+ \subseteq C_1$ such that $\mathcal{G}^c(\Sigma, P^-)$ and $\mathcal{G}^c(\Sigma, P^+)$ are connected, and $C_1 \setminus (P^- \cup P^+) = L(P^-) = L(P^+)$. In particular, this implies that every curve in P^- is disjoint from every curve in P^+ , and that every curve of $C_1 \setminus P^\pm$ crosses some curve of P^\mp .

Given a division $\underline{X} = (X^-, X^+)$ of Σ , write $\underline{P} = \underline{P}(\underline{X}) = (P^-, P^+)$, where $P^\pm = P(X^\pm)$.

Lemma 14.2. *The map $[\underline{X} \mapsto \underline{P}(\underline{X})]$ is a bijection between divisions of Σ and divisions of $\mathcal{G}_1^c(\Sigma)$.*

Proof. Let \underline{X} be a division of Σ . We have already observed that P^\pm is infinite and that $\mathcal{G}^c(\Sigma, P^\pm)$ is connected. Every curve of P^- is disjoint from every curve of P^+ . Any curve in of $C_1 \setminus P^\pm$ must cross X^\mp and hence some curve of P^\mp . We see that $\underline{P}(\underline{X})$ is a division of $\mathcal{G}_1^c(\Sigma)$.

Conversely, suppose that \underline{P} is a division of $\mathcal{G}_1^c(\Sigma)$. Let X^\pm be the subsurface of Σ filled by the curves of P^\pm . Since $\mathcal{G}^c(\Sigma, P^\pm)$ is connected, so is X^\pm . By definition, X^\pm is big. Also X^- and X^+ are homotopically disjoint (otherwise some element of P^- would cross some element of P^+). Suppose that Z were some non-annular component of $\Sigma \setminus (X^- \cup$

X^+). We can assume that Z meets X^- . Let Y be the subsurface $X^- \cup Z$. This is big, and is filled by elements of C_1 . In particular, there must be some $\gamma \in C_1$, contained in Y but not contained in X^- . Since $\gamma \notin P^-$, by hypothesis, it must cross some element of P^+ , giving a contradiction. This shows that $\Sigma \setminus (X^- \cup X^+)$ is a disjoint union of annuli. We see that $\underline{X} = (X^-, X^+)$ is a division of Σ . Now clearly, $P^\pm \subseteq P(X^\pm)$. In fact, $P^\pm = P(X^\pm)$, since any curve in $P(X^\pm) \setminus P^\pm$ would again have to cross some element of P^\mp , giving a contradiction. This shows that $\underline{P} = \underline{P}(\underline{X})$. $\underline{X} \leq \underline{Y}$ and $\underline{X} \neq \underline{Y}$.

It is clear from the construction that this gives a bijection as claimed. \square

Given two divisions, $\underline{X}, \underline{Y}$, of Σ , write $\underline{X} \leq \underline{Y}$ to mean that $X^- \preceq Y^-$, or equivalently, that $Y^+ \preceq X^+$. We write $\underline{X} < \underline{Y}$ to mean that $X^- \preceq Y^-$ and $Y^+ \prec X^+$. Clearly this defines a partial order on the set of divisions. Similarly, if $\underline{P}, \underline{Q}$, are divisions of $\mathcal{G}^c(\Sigma)$, we write $\underline{P} \leq \underline{Q}$ to mean that $P^- \subseteq Q^-$, or equivalently that $Q^+ \subseteq P^+$. These notions are equivalent under the bijection defined by Lemma 14.2.

We will want to include complexity-1 surfaces in this. To this end we define a *slice* of Σ to be an ordered pair, $\underline{X} = (X^-, X^+)$, where each of X^- and X^+ is either a complexity-1 subsurface of Σ , or a big subsurface of Σ , and such that these can be realised disjointly so that $\Sigma \setminus (X^- \cup X^+)$ is a disjoint union of (non-peripheral) annuli. In other words, it is the same as a division except that we are allowing X^- or X^+ to be an $S_{1,1}$ or $S_{0,4}$.

Again, this can be recognised from $\mathcal{G}_1(\Sigma)$. We can identify a complexity-1 slice, \underline{X} , with an element, γ , of C_1 (the separating curve), together with a sign, \pm , indicating whether the complexity-1 surface is X^- or X^+ . Note that if \underline{Y} is a division, then $\underline{X} < \underline{Y}$ corresponds to saying that $\gamma \in X^-$ and the sign is $-$, etc.

We can now distinguish elements of C_{1HT} and C_{4HS} , at least if $\xi(\Sigma) \geq 6$. Consider the following statement about an element $\gamma \in C_1$:

(*): Suppose that $\underline{X}, \underline{Y}$ are slices of Σ with $\gamma \prec X^+$ and $\gamma \prec Y^-$. Then there exist slices $\underline{Z}, \underline{W}$ of Σ with $\underline{X} < \underline{Z} < \underline{W} < \underline{Y}$.

Note that this is detectable in terms of $\mathcal{G}_1(\Sigma)$. For example, the statement that $\gamma \prec X^+$ is the same as saying that either X corresponds to \underline{P} with $\gamma \in P^+$, or else X corresponds to $(\beta, -)$, where $\beta \in C_1$ is a curve disjoint from γ .

Lemma 14.3. *Suppose $\xi(\Sigma) \geq 6$, and $\gamma \in C_1$. Then $\gamma \in C_{4HS}$ if and only if it satisfies (*).*

Proof. Suppose first that $\gamma \in C_{4HS}$. Let $\underline{X}, \underline{Y}$ be as given. We can realise these so that X^-, Y^+ and γ are pairwise disjoint. Let U be the (possibly disconnected) subsurface $X^+ \cap Y^-$. This contains γ , and so $F(\gamma) \subseteq U$. Let $\beta_1, \beta_2, \beta_3 \subseteq U \cap \partial\Sigma$ be the other boundary components of $F(\gamma)$. Let α_1, α_2 be disjoint arcs in U respectively connecting the relative boundary of X^- to α_1 and connecting α_1 to α_2 . Let $Z^- \subseteq W^-$ be subsurfaces respectively obtained by taking regular neighbourhoods of $X^- \cup \alpha_1 \cup \beta_1$ and $X^- \cup \alpha_1 \cup \beta_1 \cup \alpha_2 \cup \beta_2$. (We can take these disjoint from Y^+ .) Let Z^+, W^+ be the closures of the complements. This gives slices, $\underline{Z}, \underline{W}$ with $\underline{X} < \underline{Z} < \underline{W} < \underline{Y}$ as required, thereby verifying (*).

Suppose instead that $\gamma \in C_{1HT}$. Since $\xi(\Sigma) \geq 3$, $\Sigma \setminus F(\gamma)$ has complexity at least 3. Therefore we can find an arc, β , in $\Sigma \setminus F(\gamma)$ meeting γ precisely in its endpoints, which cuts $\Sigma \setminus F(\gamma)$ into two surfaces each of complexity at least 1. Let H be a regular neighbourhood of $F(\gamma) \cup \beta$ in Σ . Thus H is an $S_{1,2}$. Let X^-, Y^+ be the components $\Sigma \setminus H$, and let X^+, Y^- be their respective complements in Σ . These give us slices, $\underline{X}, \underline{Y}$, with $\underline{X} < \underline{Y}$ and $\gamma \prec X^+$ and $\gamma \prec Y^-$. Suppose that $\underline{Z}, \underline{W}$ are slices with $\underline{X} < \underline{Z} < \underline{W} < \underline{Y}$, as required by (*). Then H is a union of the subsurfaces $X^+ \cap Z^-, Z^+ \cap W^-$ and $W^+ \cap Y^-$, all containing at least an $S_{0,3}$. But clearly there is no room for this in an $S_{1,2}$, thereby giving a contradiction. \square

Note that this implies that we can detect the genus of Σ as the maximal number of pairwise disjoint elements of C_{1HT} we can find in Σ .

To detect the number of holes, we set $m(\Sigma)$ to be the maximal m such that there is a chain of slices, $\underline{X}_1 < \underline{X}_2 < \dots < \underline{X}_m$, of length m .

Lemma 14.4. *Assuming that $\xi(S_{g,p}) \geq 4$, we have: $m(S_{0,p}) = p - 5$, $m(S_{1,p}) = p - 2$ and $m(S_{g,p}) = 2g + p - 3$ for all $g \geq 2$.*

Proof. To begin, recall that $2g + p - 2$ is the number of pants in any pants decomposition of $\Sigma = S_{g,p}$. It is therefore also the maximal number of essential $S_{0,3}$'s we can embed disjointly in Σ .

Suppose first that $g \geq 2$. In this case, we can find a pants decomposition, $F_1, F_2, \dots, F_{2g+p-2}$, such that if $i < j < k$ then F_j separates F_i from F_k in Σ , and moreover such that F_1 and F_{2g+p-2} each have two of their boundary curves identified, and so give rise to $S_{1,1}$'s in Σ . Now let $X_i^- = \bigcup_{j=1}^i F_j$ and $X_i^+ = \bigcup_{j=i+1}^{2g+p-2} F_j$, for $i = 1, \dots, 2g + p - 3$. This gives a chain $\underline{X}_1 < \underline{X}_2 < \dots < \underline{X}_{2g+p-3}$. Conversely, given any chain $\underline{X}_1 < \underline{X}_2 < \dots < \underline{X}_m$, each of the surfaces X_1^-, X_m^+ and $X_i^+ \cap X_{i+1}^-$ for $1 \leq i \leq m - 1$ must contain an $S_{0,3}$, showing that $m + 1 \leq 2g + p - 2$. Therefore $m(\Sigma) = 2g + p - 3$.

The case when $g = 1$ is essentially the same, except in this case we lose 1, since one of the extreme surfaces (i.e. X_1^1 or X_m^+) must be an $S_{0,4}$, and this must accommodate two $S_{0,3}$'s. Similarly, if $g = 2$, we lose 2, since then both extreme surfaces will be $S_{0,4}$'s. \square

Let us summarise what we have so far detected in terms of $\mathcal{G}_1(\Sigma)$. Note first that if $\xi(\Sigma) \leq 2$ then $\mathcal{G}_1(\Sigma) = \emptyset$, and if $\xi(\Sigma) = 3$, then $\mathcal{G}_1(\Sigma)$ is just an infinite set of vertices. If $\xi(\Sigma) = 4$, then $m(\Sigma) = 2$. If $\xi(\Sigma) = 5$, then $m(\Sigma) = 3$. If $\xi(\Sigma) \geq 6$, then $m(\Sigma) \geq 4$ unless $\Sigma = S_{3,0}$, in which case, $m(S_{3,0}) = 3$. However, we can distinguish $S_{3,0}$ from the complexity-5 surfaces (namely $S_{2,2}, S_{1,5}, S_{0,8}$) by the fact that in $S_{3,0}$ we can find three disjoint curves in C_1 . Moreover, if $\xi(\Sigma) \geq 6$, we can determine g . Since we also know $m(\Sigma)$, we can also determine p .

We have shown:

Lemma 14.5. *Suppose that Σ, Σ' are compact orientable surfaces with $\mathcal{G}_1(\Sigma)$ isomorphic to $\mathcal{G}_1(\Sigma')$. Then either $\xi(\Sigma), \xi(\Sigma') \leq 2$ or $\xi(\Sigma) = \xi(\Sigma') \geq 3$. Moreover, if $\xi(\Sigma) = \xi(\Sigma') \geq 6$, then Σ and Σ' are homeomorphic.*

Note that this leaves open the question of distinguishing the different complexity-4 surfaces and the different complexity-5 surfaces (namely the classes $\{S_{2,1}, S_{1,4}, S_{0,7}\}$ and $\{S_{2,2}, S_{1,5}, S_{0,8}\}$).

We now move on to consider rigidity. We say that a slice, \underline{X} , is *simple* if $\Sigma \setminus (X^- \cup X^+)$ is connected (i.e. a single annulus). In other words, a simple slice is essentially the same thing as a (1+)-separating curve together with a transverse orientation.

We claim that we can detect simple slices. First note that if $\xi(\Sigma) \leq 5$, then all slices are simple, so we suppose $\xi(\Sigma) \geq 6$. Now a slice, \underline{X} is simple if and only if $\text{genus}(X^-) + \text{genus}(X^+) = \text{genus}(\Sigma)$. We can assume that X^- and X^+ are big, otherwise \underline{X} is certainly simple. But now we can detect $\text{genus } X^\pm$ as the maximal number of disjoint C_{1HT} curves contained in X^\pm .

We therefore have a means of describing (1+)-separating curves in Σ in terms of $\mathcal{G}_1(\Sigma)$. Such a curve, γ , corresponds to an unoriented simple slice. That is, either it is already an element of C_1 , or else it corresponds to the unordered pair, $\{X^-(\gamma), X^+(\gamma)\}$, of subsets of C_1 , and we have seen that we can recognise the set of unordered pairs which arise in this way. Moreover, we can also detect the disjointness of (1+)-separating curves from this information. In other words, we can reconstruct the whole of $\mathcal{G}_{1+}(\Sigma)$ from $\mathcal{G}_1(\Sigma)$.

We have shown:

Lemma 14.6. *Suppose that Σ, Σ' are compact orientable surfaces, and that $\xi(\Sigma) = \xi(\Sigma') \geq 4$. Then any isomorphism from $\mathcal{G}_1(\Sigma)$ to $\mathcal{G}_1(\Sigma')$ extends to an isomorphism from $\mathcal{G}_{1+}(\Sigma)$ to $\mathcal{G}_{1+}(\Sigma')$.*

(We have not shown that the extension is unique, but in certain cases at least, it must be, as will follow from the discussion below.)

The above shows that if $\mathcal{G}_{1+}(\Sigma)$ is rigid, then so is $\mathcal{G}_1(\Sigma)$. (Indeed, the converse also holds, since it is not hard to recognise 1-separating curves in $\mathcal{G}_{1+}(\Sigma)$.)

Definition. We say that Σ is of *rigid type* if $\mathcal{G}_{1+}(\Sigma)$ is rigid.

This is taken to imply that $\xi(\Sigma) \geq 4$. Note that if $p \leq 1$, then $\mathcal{G}_{1+}(\Sigma) = \mathcal{G}_s(\Sigma)$, and so applying the result of [BreM, Ki], given as Theorem 14.1 here, we see that $S_{g,0}$ is of rigid type if $g \geq 3$ and $S_{g,1}$ is of rigid type if $p \geq 2$. Also in [Bo6], it is shown that Σ is of rigid type if $\xi(\Sigma) \geq 58$.

We now proceed to applications to the pants graph, $\mathcal{E}(\Sigma)$, of Σ . We immediately get:

Theorem 14.7. *Suppose that Σ, Σ' are compact orientable surfaces with $\mathcal{E}(\Sigma)$ quasi-isometric to $\mathcal{E}(\Sigma')$. Then $\xi(\Sigma) = \xi(\Sigma')$. Moreover, if $\xi(\Sigma) = \xi(\Sigma') \geq 6$, then Σ is homeomorphic to Σ' .*

Proof. Note that $\xi(\Sigma) = 1$ if and only if $\mathcal{E}(\Sigma)$ is a quasitree; $\xi(\Sigma) = 2$ if and only if $\mathcal{E}(\Sigma)$ is hyperbolic and not a quasitree; and $\xi(\Sigma) \geq 3$ if and only if $\mathcal{E}(\Sigma)$ is not hyperbolic. We can therefore assume that $\xi(\Sigma) \geq 3$. As observed at the end of Section 13, we then have that $\mathcal{G}_1(\Sigma)$ and $\mathcal{G}_1(\Sigma')$ are isomorphic, and so the statement then follows by Lemma 14.5. \square

It is well known that $\mathcal{E}(S_{1,2})$ is quasi-isometric to $\mathcal{E}(S_{0,5})$ and that $\mathcal{E}(S_{2,0})$ is quasi-isometric to $\mathcal{E}(S_{0,6})$. It remains unclear whether or not $\mathcal{E}(S_{1,3})$ is quasi-isometric to $\mathcal{E}(S_{0,6})$. Also the classes $\{S_{2,1}, S_{1,4}, S_{0,7}\}$ and $\{S_{2,2}, S_{1,5}, S_{0,8}\}$ remain unresolved by the above.

Regarding quasi-isometric rigidity, we can show:

Theorem 14.8. *Suppose that Σ is a compact orientable surface of rigid type, and that $\phi : \mathcal{E}(\Sigma) \rightarrow \mathcal{E}(\Sigma)$ is a quasi-isometry. Then there is some $h \in \text{Map}(\Sigma)$ such that if $a \in \mathcal{E}(\Sigma)$, then $\rho(\phi a, ha) \leq k$, where k depends only on $\xi(\Sigma)$ and the parameters of the quasi-isometry, ϕ .*

As observed above, this applies, in particular, if $\Sigma = S_{g,0}$ for $g \geq 3$ or if $\Sigma = S_{g,1}$ for $g \geq 2$ or if $\xi(\Sigma) \geq 58$.

Proof. The map π given by Lemma 13.21 determines an automorphism of $\mathcal{G}_1(\Sigma)$. Therefore, by Lemma 14.6, after applying some element of

$\text{Map}(\Sigma)$, we can assume this to be the identity on $\mathcal{G}_1(\Sigma)$. In other words, if X is any terminal subsurface of Σ , we have $\pi X = X$. But now if X is any admissible subsurface of Σ , each component of $\Sigma \setminus X$ is filled by terminal subsurfaces of Σ (possibly it is a terminal subsurface). Now these subsurfaces determine X uniquely, and so it follows that X must be fixed by π . In other words, we have $\phi B(X) = B(X)$ for every admissible subsurface, X , of Σ .

Now suppose that τ is a good multicurve in Σ . Proposition 13.11 gives us a good multicurve τ' in Σ such that $\text{hd}(T_E(\tau'), \phi(T_E(\tau)))$ is finite and bounded above in terms of $\xi(\Sigma)$ and the parameters of ϕ . Also (as observed after Lemma 13.21) the map, π , given by Lemma 13.21, when restricted to $\mathcal{X}_N(\tau)$, agrees with the map π given by Lemma 13.12. Since this is the identity here, it implies that $\mathcal{X}_N(\tau') = \mathcal{X}_N(\tau)$, and so $\tau' = \tau$. In other words, we have shown that $\text{hd}(T_E(\tau), \phi(T_E(\tau)))$ is uniformly bounded above for all good multicurves τ .

The remainder of the proof follows exactly as with Theorem 10.2. Note that if τ, τ' are good multicurves with $\tau \cap \tau' = \emptyset$, then $T_E(\tau)$ and $T_E(\tau')$ diverge, and so any point a bounded distance from both gets moved a bounded distance by ϕ . But this applies to all points, since $\text{Map}(\Sigma)$ acts coboundedly on $\mathcal{E}(\Sigma)$. \square

15. CENTROIDS IN THE MARKING GRAPH

In this section, we give a direct proof of the existence of medians or “centroids” in the mapping class groups, as in [BehM2]. Specifically we will show:

Theorem 15.1. *There is a constant t_0 depending only on $\xi(\Sigma)$, such that if $a, b, c \in \mathcal{M}(\Sigma)$, then there is some $m \in \mathcal{M}(\Sigma)$ such that for all $X \in \mathcal{X}(\Sigma)$, $\sigma_X(\theta_X m, \mu_X(\theta_X a, \theta_X b, \theta_X c)) \leq t_0$. Moreover, if $m' \in \mathcal{M}(\Sigma)$ 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 $\mu : \mathcal{M}^3 \rightarrow \mathcal{M}$ by setting $\mu(a, b, c) = m$. Of course, it is enough to define $\mu(a, b, c)$ for a, b, c in the vertex set, \mathcal{M}^0 , of \mathcal{M} .

The argument follows broadly as in [BehM2] using [BehKMM], though we don't need to appeal directly to the compatibility theorem for projection maps — we effectively reprove this for medians. The constants given in [BehM2] depend on the complexity, $\xi(\Sigma)$. However, we now have the results available to remove this dependency; specifically, the uniform hyperbolicity of the curve graphs and uniform constants for the Bounded Geodesic Image Theorem and Behrstock's lemma (see Theorem 7.1, and Lemmas 7.5 and 7.6. respectively, of this paper). It

therefore seems likely that, suitably formulated, one can choose t_0 in Theorem 15.1 to be universal, though we do not pursue this issue here.

We elaborate on the compatibility of the rank-1 median structures associated to transverse subsurfaces.

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

(Note that the vertex set is a rank-1 median algebra, and that 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 (in other words, πA generates $V(\Delta)$ as a median algebra. 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 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* if $\{\Delta_i\}_{i \in \mathcal{J}}$ 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 $T = T(\{\Delta_i\}_{i \in \mathcal{J}})$. (This is some ambiguity, in that we could swap two trees that are single edges and meet in a single vertex, but we don't need to worry about that here.)

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 non-trivial.

Lemma 15.2. *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 15.3. *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) Without loss of generality, $v_{01} = v_{02}$. 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. In a metric space, (G, σ) , we write $\langle x, y, z \rangle = \frac{1}{2}(\sigma(x, z) + \sigma(y, z) - \sigma(x, y))$ for the Gromov product (as in Section 7).

Lemma 15.4. *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 a maps $\pi : B \rightarrow V(\Delta)$, and $\lambda : V(\Delta) \rightarrow G$ such that for all $x, y, z \in V(\Delta)$, if $\langle x, y, z \rangle \leq t$, then $\pi z \in [x, y]_{V(\Delta)}$. Moreover, λ is a*

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 .

(Recall that “ h -quasimorphism” means that $\sigma(\lambda\mu(x, y, z), \mu(\lambda x, \lambda y, \lambda z)) \leq h$ for all x, y, z .)

Proof. This is proven in [Bo1]. 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 finite tree (with vertex set B). The additive constant depends only 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 11.2 here.) (In [Bo1] 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 curve graphs. Recall that if $X \in \mathcal{X}(\Sigma)$ then $\mathcal{G}(X)$ is k -hyperbolic for some universal k (we could allow k to depend on $\xi(\Sigma)$ here if we want). Given $p \in \mathbb{N}$, we will choose universal $t \geq 0$ sufficiently large as described below. Given $A \subseteq \Sigma$ we apply Lemma 15.4 to $B = \theta_X(A) \subseteq \mathcal{G}(X)$ with t as above, to get at 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 15.10, 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 $\sigma_X(\theta_X, \theta_Y) \leq t$, then $\pi_X a = \pi_X b$. 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 15.10) we also note that λ is an h -quasimorphism, and that for all $a \in A$, $\sigma_X(\theta_X a, \lambda_X \pi_X a) \leq h$, where $h = h(p)$ depends only on p .

Lemma 15.5. *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 $\sigma_X(\theta_X a, \theta_X Y) > r_0$, then $\sigma_Y(\theta_Y a, \theta_Y X) \leq r_0$. 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 $\sigma_X(\theta_X a_{XY}, \theta_X Y) \leq r_0$. 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 $\sigma_X(\theta_X b, \theta_X a_{XY}) \geq t(p)$. Thus, $\sigma_X(\theta_X b, \theta_X Y) \geq t(p) - r_0 > r_0$. Similarly, $\sigma_Y(\theta_Y b, \theta_Y X) > r_0$. This contradicts Lemma 7.6, proving that no such b exists. \square

Note that, by Lemma 15.2, 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.

Note that, by construction, $\sigma_X(\theta_X a_{XY}, \theta_X Y) \leq r_0$. Also, if $Z \in \mathcal{X}$ with $\Delta(Z)$ non-trivial, we have $\sigma_X(\theta_X a_{XZ}, \theta_X Z) \leq r_0$. If $\sigma_X(\theta_X Y, \theta_X Z) < t - 2r_0$, then $\sigma_X(\theta_X a_{XY}, \theta_X a_{XZ}) < s(p)$, so $v_{XY} = \pi_X a_{XY} = \pi_X a_{XZ} = v_{XZ}$. For future reference (Lemma 15.10) we also note that $\sigma_X(\theta_X a_{XY}, \lambda_X v_{XY}) = \sigma_X(\theta_X a_{XY}, \lambda_X \pi_X a_{XY}) \leq h(p)$, so $\sigma_X(\theta_X Y, \lambda_X v_{XY}) \leq r_0 + h(p)$.

We write \mathcal{X}_0 for the set of $X \in \mathcal{X}$ such that $\Delta(X)$ is non-trivial. It follows from Proposition 7.2, 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 15.1 and 15.4. Note that this determines vertices $v_{XY} \in \Delta(X)$ and $v_{YX} \in \Delta(Y)$ which get identified in the common enlargement, $\Delta(X, Y)$.

Lemma 15.6. *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 there (since there must be boundary curves of Y and Z which are disjoint) we must have $\sigma_X(\theta_X Y, \theta_X Z) \leq l$, for some universal constant, l , depending only (or at most) on $\xi(\Sigma)$. Provided we have chosen $t > l + 2r_0$, this implies that $v_{XY} = v_{XZ}$. \square

Lemma 15.7. *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 15.6, 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 Lemma 7.6, $\min\{\sigma_X(\theta_X Y, \theta_X Z), \sigma_Y(\theta_Y X, \theta_Y Z)\} \leq r_0$. Therefore, if $t \geq 3r_0$, we see that either $v_{XY} = v_{XZ}$ or $v_{YZ} = v_{YX}$. Similarly, we have ($v_{YZ} = v_{YX}$ or $v_{ZX} = v_{ZY}$) and ($v_{ZX} = v_{ZY}$ or $v_{XY} = v_{XZ}$), and so the statement follows. \square

We can now start the proof of Theorem 15.1

Suppose $a, b, c \in \mathcal{M}^0$. We want to find a median for a, b, c in \mathcal{M}^0 . First choose any $d \in \mathcal{M}^0$ with $\sigma_\Sigma(\theta_\Sigma d, \mu_\Sigma(\theta_\Sigma a, \theta_\Sigma b, \theta_\Sigma c))$ bounded in the curve graph $(\mathcal{G}(\Sigma), \sigma_\Sigma)$ of Σ . (Choose any centre for $\theta_\Sigma a, \theta_\Sigma b, \theta_\Sigma c$ in $(\mathcal{G}(\Sigma), \sigma_\Sigma)$ and extend arbitrarily to a marking of Σ .)

Now set $A = \{a, b, c, d\}$, and let $\pi_X : A \rightarrow \Delta(X)$ be as described in Lemma 15.4. Let $h = h(4)$. Write $d_X = \pi_X d$ and $e_X = \mu_X(\pi_X a, \pi_X b, \pi_X c)$. Let \mathcal{X}_0 be the (finite) set of $X \in \mathcal{X}$ such that $\Delta(X)$ is non-trivial. 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 $\mu_T d$ 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 can imagine these relations as follows. We imagine any finite set of elements of \mathcal{X} embedded 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 $=, \prec, \wedge$ 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 of hyperbolic 3-manifolds.)

Lemma 15.8. *$X, Y, Z \in \mathcal{X}_1$ 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 15.6 and 15.7, $X \pitchfork Z$, and $\{\Delta(X), \Delta(Y), \Delta(Z)\}$ is coherent. In particular, $e_X = v_{XY} = v_{XZ}$ and $d_Z = v_{ZY} = v_{ZX}$ so $X \ll Z$. \square

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

Lemma 15.9. *Given $X, Y, Z \in \mathcal{X}_1$ with $X \ll Y$ and $Y \prec Z$, then either $X \ll Z$ or $X \prec Z$.*

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 $Z = X$, $Z \prec X$, $Z \ll X$ or $Z \wedge X$. Now $Z = X$ or $Z \prec X$ both give $Y \prec X$ contradicting $X \ll Y$; $Z \ll X$ gives $Z \ll Y$ contradicting $Y \prec Z$, 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 $X \prec Y$. This relation is antisymmetric on \mathcal{X}_1 . It is not transitive, but in view of Lemma 15.9, 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 $U \prec X$. Note that $\Sigma \notin \mathcal{X}_1$, so $U \neq \Sigma$.

From this, we can deduce:

Lemma 15.10. *There is some universal $u_0 > 0$, such that if $a, b, c \in \mathcal{M}^0$, there is some $\alpha \in V(\mathcal{G}(\Sigma))$ such that if $X \in \mathcal{X}$, with $\alpha \pitchfork X$ and $\alpha \prec X$, then $\sigma_X(\theta_X \alpha, \mu_X(\theta_X a, \theta_X b, \theta_X c)) \leq u_0$.*

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 $\alpha \prec X$ or $\alpha \pitchfork X$. Then either $U \prec X$ or $U \pitchfork X$. As in Section 7, we use the notation \sim to mean “up to bounded distance”. In all cases, $\theta_X \alpha$ is defined and $\theta_X \alpha \sim \theta_X U$. Now, $\lambda_X e_X = \lambda_X \mu_{V(\Delta(X))}(\pi_X a, \pi_X b, \pi_X c) \sim \mu_X(\lambda_X \pi_X a, \lambda_X \pi_X b, \lambda_X \pi_X c) \sim \mu_X(\theta_X a, \theta_X b, \theta_X c)$. We therefore want to show that $\theta_X U \sim \lambda_X e_X$. Note that $\lambda_X d_X = \lambda_X \pi_X d \sim \theta_X d$, and

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 centre for $\theta_X a, \theta_X b, \theta_X c$ in $\mathcal{G}(X)$, so if $\theta_X U$ were far enough away (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 $\langle \theta_X a, \theta_X c; \theta_X U \rangle$ are both greater than r_0 (after permuting a, b, c as necessary). By Lemma 7.5, 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 15.5, the trees $\Delta(X)$ and $\Delta(U)$ are coherent. Moreover, since $\Delta(U)$ is non-trivial, we have $\theta_X U \sim \lambda_X v_{XU}$. Suppose that $X \in \mathcal{X}_1$. Then $X \ll U$, so $e_X = v_{XU}$, thus $\theta_X U \sim \lambda_X v_{XU} = \lambda_X e_X$ as required. 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:

Proof of Theorem 15.1. Uniqueness up to bounded distance, depending in $\xi(\Sigma)$, is an immediate consequence of Proposition 7.2, so we prove existence.

Note that the discussion of Section 7 allows us to talk loosely of “markings” as filling sets of curves of bounded pairwise intersection. Such a set is always a bounded distance (in terms of self-intersection) from some element of \mathcal{M} .

Let $a, b, c \in \mathcal{M}$. Given $X \in \mathcal{X}$, let $\delta_X = \mu_X(\theta_X a, \theta_X b, \theta_X c) \in \mathcal{G}(X)$. Let α be a curve as given by Lemma 15.10. We consider only the case when α separates Σ . The non-separating 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 15.1 holds intrinsically to Y and Z . Thus, we can find an intrinsic markings, m_Y , of Y such that if $X = Y$ or $X \prec Y$, then $\sigma_X(\theta_X m, \delta_X)$ is bounded. (This implicitly uses the fact that subsurface projection to Y and then intrinsically in Y to X agrees up to bounded distance with subsurface projection directly to X .) We have a similar marking, m_Z , of Z . Now extend $m_Y \cup m_Z \cup \{\alpha\}$ to a marking, m , of Σ such that $\sigma_\Omega(\theta_\Omega m, \delta_\Omega)$ is bounded, where $\Omega \in \mathcal{X}$ is the annulus with core curve α . (For this, extend to any marking, and then apply a suitable Dehn twist about α .) As observed above, we may as well assume that $m \in \mathcal{M}$.

Suppose that $X \in \mathcal{X}$. If $X = Y$, $X \prec Y$, $X = Z$, $X \prec Z$ or $X = \Omega$, then $\sigma_X(\theta_X m, \delta_X)$ is bounded. If not, then either $\alpha \prec X$ or $\alpha \pitchfork X$. But then, by the choice of α , $\sigma_X(\theta_X \alpha, \delta_X)$ is bounded. But $\sigma_X(\theta_X m, \theta_X \alpha)$ is bounded, so we are done.

If either Y or Z is a $S_{0,3}$, we just take the corresponding m_Y or m_Z to be empty, and proceed in the same way. \square

REFERENCES

- [Ao] T.Aougab, *Uniform hyperbolicity of the graphs of curves* : Geom. Topol. **17** (2013) 3855–2875.
- [Ar] J.Aramayona, *The Weil-Petersson geometry of the five-times punctured sphere* : in “Spaces of Kleinian groups”, London Math. Soc. Lecture Notes Ser. Vol. 329, Cambridge Univ. Press (2006) 219–231.
- [BaH] H.-J.Bandelt, J.Hedlikova, *Median algebras* : Discrete Math. **45** (1983) 1–30.
- [Beh] J.A.Behrstock, *Asymptotic dimension of the mapping class group and Teichmüller space* : Geom. Topol. **10** (2006) 1523–1578.
- [BehDS] J.A.Behrstock, C.Druțu, M.Sapir, *Median structures on asymptotic cones and homomorphisms into mapping class groups* : Proc. London Math. Soc. **3** (2011) 503–554.
- [BehKMM] J.A.Behrstock, B.Kleiner, Y.Minsky, L.Mosher, *Geometry and rigidity of mapping class groups* : Geom. Topol. **16** (2012) 781–888.
- [BehM1] J.A.Behrstock, Y.N.Minsky, *Dimension and rank for mapping class groups* : Ann. of Math. **167** (2008) 1055–1077.
- [BehM2] J.A.Behrstock, Y.N.Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.
- [BesBF] M.Bestvina, K.Bromberg, K.Fujiwara, *Constructing group actions on quasi-trees and applications to mapping class groups* : preprint, 2012, posted at arXiv:1006.1939.
- [Bo1] B.H.Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.
- [Bo2] B.H.Bowditch, *Embedding median algebras in products of trees* : Geom. Dedicata **170** (2014) 157–176.

- [Bo3] B.H.Bowditch, *Uniform hyperbolicity of the curve graphs* : Pacific J. Math. **269** (2014) 269–280.
- [Bo4] B.H.Bowditch, *Some properties of median metric spaces* : preprint, Warwick 2013.
- [Bo5] B.H.Bowditch, *Quasi-isometric maps between direct products of hyperbolic spaces* : preprint, Warwick 2014.
- [Bo6] B.H.Bowditch, *Rigidity of the strongly separating curve graph* : preprint, Warwick 2014.
- [Bo7] B.H.Bowditch, *The coarse geometry of the Teichmüller metric* : preprint, Warwick/Tokyo 2014.
- [BreM] T.E.Brendle, D.Margalit, *Commensurations of the Johnson kernel* : Geom. Topol. **8** (2004) 1361–1384.
- [Bro] J.F.Brock, *The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores* : J. Amer. Math. Soc. **16** (2003) 495–535.
- [BroF] J.Brock, B.Farb, *Curvature and rank of Teichmüller space* : Amer. J. Math. **128** (2006) 1–22.
- [BroM] J.Brock, D.Margalit, *Weil-Petersson isometries via the pants complex* : Proc. Amer. Math. Soc. **135** (2007) 795–803.
- [ChaDH] I.Chatterji, C.Druţu, F.Haglund, *Kazhdan and Haagerup properties from the median viewpoint* : Adv. Math. **225** (2010) 882–921.
- [Che] V.Chepoi, *Graphs of some CAT(0) complexes* : Adv. in Appl. Math. **24** (2000) 125–179.
- [CIRS] M.Clay, K.Rafi, S.Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences* : preprint, 2013, posted at arXiv:1302.5519, to appear in Algebr. Geom. Topol.
- [EiS] S.Eilenberg, N.Steenrod, *Foundations of algebraic topology* : Princeton University Press, Princeton, New Jersey (1952).
- [En] R.Engelking, *Theory of dimensions finite and infinite* : Sigma Series in Pure Math. No. 10, Heldermann Verlag, Lemgo, (1995).
- [EsMR] A.Eskin, H.Masur, K.Rafi, *Large scale rank of Teichmüller space* : preprint, 2013, posted at arXiv:1307.3733.
- [FLM] B.Farb, A.Lubotzky, Y.Minsky, *Rank-1 phenomena for mapping class groups* : Duke Math. J. **106** (2001) 581–597.
- [G] M.Gromov, *Asymptotic invariants of infinite groups* : “Geometric group theory, Vol. 2” London Math. Soc. Lecture Note Ser. No. 182, Cambridge Univ. Press (1993).
- [Ha] U.Hamenstädt, *Geometry of the mapping class groups III: Quasi-isometric rigidity* : preprint, 2007, posted at arXiv:0512429.
- [HePW] S.Hensel, P.Przytycki, R.C.H.Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs* : preprint, 2013, posted at arXiv:1301.5577, to appear in J. Eur. Math. Soc.
- [Ir] E.Irmak, *Complexes of nonseparating curves and mapping class groups* : Michigan Math. J. **54** (2006) 81–110.
- [Is] J.R.Isbell, *Median algebras* : Trans. Amer. Math. Soc. **260** (1980) 319–362.
- [Iv] N.V.Ivanov, *Automorphism of complexes of curves and of Teichmüller spaces* : Internat. Math. Research Notices **14** (1997) 651–666.

- [KaKL] M.Kapovich, B.Kleiner, B.Leeb, *Quasi-isometries and the de Rham decomposition* : Topology **37** (1998) 1193–1211.
- [Ki] Y.Kida, *Automorphisms of the Torelli complex and the complex of separating curves* : J. Math. Soc. Japan **62** (2011) 363–417.
- [KIL] B.Kleiner, B.Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings* : Inst. Hautes Études Sci. Publ. Math. **86** (1997) 115–197.
- [Ko] M.Korkmaz, *Automorphisms of complexes of curves in punctured spheres and on punctured tori* : Topology Appl. **95** (1999) 85–111.
- [L] F.Luo, *Automorphisms of the complex of curves* : Topology **39** (2000) 283–298.
- [Man] J.Mangahas, *Uniform exponential growth of subgroups of the mapping class group* : Geom. Funct. Anal. **19** (2010) 1468–1480.
- [Mar] D.Margalit, *Automorphisms of the pants complex* : Duke Math. J **121** (2004) 457–479.
- [MasM1] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves I: hyperbolicity* : Invent. Math. **138** (1999) 103–149.
- [MasM2] H.A.Masur, Y.N.Minsky, *Geometry of the complex of curves II: hierarchical structure* : Geom. Funct. Anal. **10** (2000) 902–974.
- [MasW] H.Masur, M.Wolf, *The Weil-Petersson isometry group* : Geom. Dedicata **93** (2002) 177–190.
- [RaS] K.Rafi, S.Schleimer, *Curve complexes are rigid* : Duke Math J. **158** (2011) 225–246.
- [Ro] M.A.Roller, *Poc-sets, median algebras and group actions, an extended study of Dunwoody’s construction and Sageev’s theorem* : Habilitationsschrift, Regensburg, 1998.
- [Sha] K.J.Shackleton, *Combinatorial rigidity in curve complexes and mapping class groups* : Pacific J. Math. **230** (2007) 217–232.
- [Sho] M.Sholander, *Medians and betweenness* : Proc. Amer. Math. Soc. **5** (1954) 801–807.
- [VaW] L.Van den Dries, A.J.Wilkie, *On Gromov’s theorem concerning groups of polynomial growth and elementary logic* : J. Algebra **89** (1984) 349–374.
- [Ve] E.R.Verheul, *Multimedians in metric and normed spaces* : CWI Tract, 91, Stichting Mathematisch Centrum, Centrum voor Wiskunde en Informatica, Amsterdam, 1993.
- [We] R.C.H.Webb, *A short proof of the bounded geodesic image theorem* : preprint, 2013, posted at arXiv:1301.6187, to appear in J. Reine Angew. Math.
- [Wo] S.A.Wolpert, *Geometry of the Weil-Petersson completion of Teichmüller space* : in “Surveys in differential geometry, Vol. VIII” Int. Press. (2003) 357–393.

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

CURRENT ADDRESS: DEPARTMENT OF MATHEMATICAL AND COMPUTING SCIENCES, TOKYO INSTITUTE OF TECHNOLOGY, 2-12-1, W8-37 O-OKAYAMA, MEGURO-KU, TOKYO 152-8552, JAPAN.