LARGE-SCALE RANK AND RIGIDITY OF THE WEIL-PETERSSON METRIC

BRIAN H. BOWDITCH

Abstract. We study the large-scale geometry of Weil-Petersson space, that is, Teichmüller space equipped with the Weil-Petersson metric. We show that this admits a natural coarse median structure of a specific rank. Given that this is equal to the maximal dimension of a quasi-isometrically embedded euclidean space, we recover a result of Eskin, Masur and Rafi which gives the coarse rank of the space. We go on to show that, apart from finitely many cases, the Weil-Petersson spaces are quasi-isometrically distinct, and quasi-isometrically rigid. In particular, any quasi-isometry between such spaces is a bounded distance from an isometry. By a theorem of Brock, Weil-Petersson space is equivariantly quasi-isometric to the pants graph, so our results apply equally well to that space.

1. Introduction

In this paper, we investigate the large-scale geometry of Teichmüller space in the Weil-Petersson metric. In particular, we give a formula for the coarse rank of this space (the maximal dimension of a quasi-isometrically embedded copy of the euclidean plane), thereby recovering a result of [EMR1]. We go on to prove quasi-isometric rigidity of the Weil-Petersson metric for all but finitely many surfaces. In other words, any self-quasi-isometry is a bounded distance from an isometry induced by an element of the mapping class group (Theorem 1.4). To obtain these results, we show that the Weil-Petersson space admits a natural ternary operation, well defined up to bounded distance, which gives it the structure of a coarse median space, as defined in [Bo1]. We remark that related results for the Teichmüller metric have been proven elsewhere, see [EMR1, Bo7, EMR2].

The main results only depend on the equivariant quasi-isometry type of the spaces involved. Note that, in [Bro], it was shown that the Weil-Petersson space is equivariantly quasi-isometric to the pants graph. Therefore, all the main results stated here are valid also for that space. Indeed we will work mostly with the pants graph in this paper.

Our proofs will rely heavily on results of [Bo8], where relevant ideas are developed, and applied to the case of mapping class group itself (or equivalently to the "marking graph"). Central to this, is the notion of a "coarse median". (In the case of mapping class group, such a median was constructed in [BeM2].)

Date: First draft, 3rd March 2015. Revised, 13th January 2020.

Let Σ be a compact orientable surface of genus $g = g(\Sigma)$ and with $p = p(\Sigma)$ boundary components. We write $\xi(\Sigma) = 3g + p - 3$ for its *complexity*. We sometimes denote its topological type by $S_{g,p}$. We will also write $\xi_0(\Sigma) = \lfloor (\xi(\Sigma)+1)/2 \rfloor$.

By the *Weil-Petersson space*, $\mathbb{W}(\Sigma)$, of Σ, we mean the Teichmüller space of Σ (that is, the space of marked complete finite-area hyperbolic structures on the interior of Σ) equipped with the Weil-Petersson metric. This is a negatively curved riemannian manifold diffeomorphic to $\mathbb{R}^{2\xi(\Sigma)}$. It is not complete, but is geodesically convex and CAT(0) (see [Wo]). Its completion has a natural stratification, where the lower dimensional strata are geodesically embedded direct products of lower-complexity Weil-Petersson spaces. From this, one can easily see that the completion contains flats (geodesically embedded euclidean spaces) of dimension $\xi_0(\Sigma)$, but of no higher dimension. In other words, it has "euclidean rank" equal to $\xi_0(\Sigma)$. (Indeed, it does not contain any isometrically embedded euclidean ball of dimension $\xi_0(\Sigma) + 1$.) Note also that the mapping class group, $\mathrm{Map}(\Sigma)$, acts by isometry on $\mathbb{W}(\Sigma)$. It was shown in [MasW] that $\mathbb{W}(\Sigma)$ is "rigid", in the sense that every isometry is induced by some element of $\mathrm{Map}(\Sigma)$. Here, we aim to give coarse versions of the above statements.

The following is a central observation for our arguments:

Theorem 1.1. *There is a ternary operation, $\mu : \mathbb{W}(\Sigma)^3 \longrightarrow \mathbb{W}(\Sigma)$, which endows $\mathbb{W}(\Sigma)$ with the structure of a coarse median space of rank $\xi_0(\Sigma)$.*

The notion of a "coarse median space" was defined in [Bo1], and will be discussed further in Section 3 here. Roughly speaking, it says that, when dealing with a subset of the space of bounded finite cardinality, the median operation behaves like the standard median operation on a finite CAT(0) cube complex. Moreover, we can take this cube complex to have dimension at most the (coarse median) rank (which is equal to $\xi_0(\Sigma)$ in the case of $\mathbb{W}(\Sigma)$).

Note that if $\xi_0(\Sigma) = 1$ (i.e. Σ is $S_{0,5}$ or $S_{1,2}$), then $\mathbb{W}(\Sigma)$ is rank 1, and so (by Theorem 2.1 of [Bo1]) we recover the result of [BroF] (see also [Ar]) that $\mathbb{W}(\Sigma)$ is hyperbolic in this case. (In fact, $\mathbb{W}(S_{0,5})$ is quasi-isometric to $\mathbb{W}(S_{1,2})$, since these surfaces have isomorphic pants graphs as we discuss in Section 7.)

We remark that it is shown in [BeHS1] that Weil-Petersson space is "hierarchically hyperbolic" and that this, in turn, implies that it is coarse median. We comment further on this in Section 3.

The following result of [EMR1] is now an immediate consequence of Theorem 1.1 above together with Corollary 2.4 of [Bo1]:

Theorem 1.2. *There is a quasi-isometric embedding of euclidean n-space into $\mathbb{W}(\Sigma)$ if and only if $n \leq \xi_0(\Sigma)$.*

In fact (as in [EMR1]), the same statement holds with "euclidean space" replaced by euclidean half-space. Indeed, one can make a stronger statement: see Proposition 3.2 here.

An immediate consequence of Theorem 1.2 is that if Σ and Σ' are two surfaces with $\mathbb{W}(\Sigma)$ quasi-isometric to $\mathbb{W}(\Sigma')$, then $\xi_0(\Sigma) = \xi_0(\Sigma')$. In fact, we will show the following.

Theorem 1.3. *Suppose that Σ and Σ' are compact orientable surfaces with $\mathbb{W}(\Sigma)$ quasi-isometric to $\mathbb{W}(\Sigma')$, 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 complexity 4 and 5. In fact, one can also distinguish $\mathbb{W}(S_{2,1})$, as we discuss in Section 7. The remaining cases we leave unresolved here. (Though I suspect that these are also quasi-isometrically distinct.)

The main result of this paper refers to quasi-isometric rigidity. We say that a geodesic metric space with an isometric group action (in this case, $\mathbb{W}(\Sigma)$ acted upon by $\mathrm{Map}(\Sigma)$) is *quasi-isometrically rigid* if any self-quasi-isometry is a bounded distance from the isometry induced by some element of the group.

We will show:

Theorem 1.4. *If $g(\Sigma) + p(\Sigma) \geq 7$, or if $g(\Sigma) \geq 3$ and $p(\Sigma) \leq 1$, then $\mathbb{W}(\Sigma)$ is quasi-isometrically rigid.*

Indeed, the distance bound of the conclusion depends only on the quasi-isometry constants and on $\xi(\Sigma)$. Clearly, the result covers all but finitely many topological types of surfaces.

Theorem 1.4 is a consequence of Theorem 7.9 here, which reduces it to a combinatorial rigidity statement for certain curve graphs. The relevant combinatorial rigidity theorems are proven respectively in [BreM], [Ki] and [Bo6] (see Theorem 7.7 here). The result as stated leaves about a dozen cases open, as we discuss further at the end of Section 7.

The proofs of Theorems 1.3 to 1.4 above make use of the asymptotic cone, $\mathbb{W}^\infty(\Sigma)$, of $\mathbb{W}(\Sigma)$ (see [VaW, G, DrK]). This is a complete metric space, in this case CAT(0) (since $\mathbb{W}(\Sigma)$ is). The following is a direct consequence of Theorem 1.1, as shown in [Bo1] (see Corollary 2.4, thereof).

Theorem 1.5. *Any asymptotic cone, $\mathbb{W}^\infty(\Sigma)$, of $\mathbb{W}(\Sigma)$ has locally compact dimension $\xi_0(\Sigma)$.*

The *locally compact dimension* of a topological space is the maximal dimension of a locally compact subset (cf. [BeM1]). As described in [Bo1], Theorem 1.5 is easily seen to imply the "only if" direction of Theorem 1.2. (The "if" direction is relatively straightforward, as noted earlier.)

Theorem 1.5 can also be viewed as a general consequence of the fact that $\mathbb{W}^\infty(\Sigma)$ admits a median metric of rank ξ_0 (see Theorem 3.3 here).

We also have:

Theorem 1.6. *If Σ and Σ' are compact orientable surfaces with $\mathbb{W}^\infty(\Sigma)$ homeomorphic to $\mathbb{W}^\infty(\Sigma')$, then either $\xi(\Sigma) = \xi(\Sigma') \geq 3$, or $\xi(\Sigma), \xi(\Sigma') \leq 2$.*

(In fact, in the above, we could take the asymptotic cones of Σ and Σ' to be with respect to different scaling factors and ultrafilters.)

In fact, much of this paper will be devoted to understanding the structure of the asymptotic cone. To apply this to quasi-isometric rigidity, we make use of some results of [Bo8]. These were inspired by related arguments in [BeKMM].

As noted above, it is shown in [Bro] that $\mathbb{W}(\Sigma)$ is $\mathrm{Map}(\Sigma)$-equivariantly quasi-isometric to the pants graph, denoted here as $\mathbb{P}(\Sigma)$. (The pants graph will be discussed in Section 2 here.) In fact, we will mostly work with $\mathbb{P}(\Sigma)$ instead of $\mathbb{W}(\Sigma)$. All the main theorems stated above apply equally well with $\mathbb{P}(\Sigma)$ substituted for $\mathbb{W}(\Sigma)$ (though the asymptotic cone, $\mathbb{P}^\infty(\Sigma)$ will only be bilipschitz equivalent to a CAT(0) space in this case).

We remark that there are also combinatorial rigidity results for $\mathbb{P}(\Sigma)$. In [Mar] it is shown that any automorphism of $\mathbb{P}(\Sigma)$ is induced by an element of $\mathrm{Map}(\Sigma)$ (with a few qualifications for the low-complexity cases). In [BroMar], this fact was used to give another proof of the result of [MasW] mentioned above.

We note that related results for Teichmüller space in the Teichmüller metric are explored in [Bo7], also using the fact that it admits a coarse median structure (in this case of rank $\xi(\Sigma)$). A different approach to quasi-isometric rigidity of the Teichmüller metric has been given independently in [EMR2]. We also remark that the quasi-isometry types of these and various other metrics on Teichmüller space are studied in [Y].

We remark that some of the results of [BeHS2] apply to Weil-Petersson space: in particular the fact that a top-dimensional quasiflat is a finite Hausdorff distance from a union of quasi-isometrically embedded orthants (see also [Bo9]). This offers some new insights into the geometry of this space, though it does not appear sufficient to prove quasi-isometric rigidity results in this case. (In particular, Assumptions 2 and 3 of their Section 6 fail for $\mathbb{W}(\Sigma)$.)

The proof of rigidity in outline goes as follows. We begin with a general discussion of pants graphs and coarse median spaces in Sections 2 and 3. (The latter includes proofs of Theorems 1.1, 1.2 and 1.5.) In Section 4, we associate to any multicurve, τ, in Σ, a subset, $T(\tau)$, of the pants graph $\mathbb{P}(\Sigma)$. (Under the standard quasi-isometry to $\mathbb{W}(\Sigma)$, this corresponds to a stratum of the completion of Weil-Petersson space.) If τ is a "good" multicurve, then $T(\tau)$ is a quasi-isometrically embedded copy of a direct product of ξ_0 "bushy" hyperbolic spaces (see Proposition 4.4). We refer to such sets loosely as "product regions" in $\mathbb{P}(\Sigma)$. The aim is then to show that these product region, together with their product structure, are coarsely preserved by any quasi-isometry (Lemma 5.10 and Proposition 5.11). To see this, we pass to the asymptotic cone. In the limit, a product region gives rise to a direct product of ξ_0 "furry" $\mathbb{R}$-trees (see Lemma 4.7). A general fact about median metric spaces (Proposition 4.8) allows us to recognise such subsets topologically (Corollary 5.8). A general result about families of subsets of an asymptotic cone then allows us to pass back to the pants graph to deduce Lemma

5.10 and Corollary 5.11. The manner in which product regions are arranged in $\mathbb{P}(\Sigma)$ is encoded in a certain combinatorial graph. It follows (Section 6) that a quasi-isometry of $\mathbb{P}(\Sigma)$ induces an automorphism of this graph. In Section 7 we show that this, in turn, induces an automorphism of what we call the "strongly separating curve graph" of Σ. In the cases where this is known to be rigid, such an automorphism is induced by an element of $\mathrm{Map}(\Sigma)$, which we may as well take to be the identity. This then tells us that each product region in $\mathbb{P}(\Sigma)$ is preserved up to bounded Hausdorff distance. From the general abundance of product regions, it follows easily that the quasi-isometry moves each point of $\mathbb{P}(\Sigma)$ a bounded distance. This proves the quasi-isometric rigidity of $\mathbb{P}(\Sigma)$, or equivalently that of $\mathbb{W}(\Sigma)$, in such cases (Theorem 7.9). This accounts for all but finitely many topological types of Σ (Theorem 7.7). From this Theorem 1.4 follows.

Notation. We will be referring to [Bo1] and [Bo8], which deal with more general spaces, or with analogous spaces in a slightly different context. We comment briefly on how these correspond. In all cases, we have two families of spaces, each indexed by a set $\mathcal{X}$, typically the set of subsurfaces of Σ; together with various "projection maps" between them. In the general set up of [Bo1], we have a family of coarse median spaces, denoted $\Lambda(X)$, as well as a family of hyperbolic spaces, denoted $\Theta(X)$, where $X \in \mathcal{X}$. In [Bo8], these are denoted respectively, by $\mathcal{M}(X)$ and $\mathcal{G}(X)$ (and later specialised, respectively, to the marking graph, $\mathbb{M}(X)$, and curve graph, $\mathbb{G}(X)$). In the present paper, when X is not an annulus, $\Lambda(X)$ (or $\mathcal{M}(X)$) will correspond to the pants graph, $\mathbb{P}(X)$; and $\Theta(X)$ (or $\mathcal{G}(X)$) will again correspond to the curve graph, $\mathbb{G}(X)$. When X is an annulus, $\mathbb{P}(X)$ and $\mathbb{G}(X)$ will both be deemed to be singletons. In fact, annular subsurfaces play no essential role in the present paper. (One could alternatively delete annular subsurfaces from the indexing set, though for consistency with notation elsewhere, we will include them.) The various projection maps are denoted similarly in all cases.

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

2. The pants graph

In this section, we discuss the coarse geometry of the pants graph and define the notion of subsurface projection in this context. The pants graph was originally described for closed surfaces in the Appendix of [HaT]. (There the authors use the term "marking", though this has since commonly come to have a different meaning.)

First, we recall some standard general definitions. Throughout this paper, $N(.;r)$ will denote a closed uniform r-neighbourhood in a metric space. Suppose that $\phi : (\Lambda, \rho) \longrightarrow (\Lambda', \rho')$ is a map between two geodesic metric spaces. We say

that ϕ is *coarsely lipschitz* if

$$(\exists k_1, k_2 \geq 0)(\forall x, y \in \Lambda)(\rho'(\phi x, \phi y) \leq k_1 \rho(x, y) + k_2).$$

We say that ϕ is a *quasi-isometric embedding* if it is coarsely lipschitz and

$$(\exists k_3, k_4 \geq 0)(\forall x, y \in \Lambda)(\rho(x, y) \leq k_3 \rho'(\phi x, \phi y) + k_4).$$

We say that ϕ is a *quasi-isometry* if it is a quasi-isometric embedding and there is some $k_5 \geq 0$ such that $\Lambda' = N(\phi(\Lambda); k_5)$.

By a *curve* in Σ we mean a free homotopy class of non-trivial non-peripheral simple closed curves in Σ. We write $\mathbb{G}^0(\Sigma)$ for the set of curves. Given $\alpha, \beta \in \mathbb{G}^0(\Sigma)$, we write $\iota(\alpha, \beta)$ for their (geometric) intersection number (that is, the minimum number of intersections among representatives of these classes). As usual, we define the *curve graph*, $\mathbb{G}(\Sigma)$, to be the graph with vertex set $\mathbb{G}^0(\Sigma)$, where α, β are deemed adjacent if they have minimal possible intersection for that surface (in other words, disjoint if $\xi(\Sigma) \geq 2$; $\iota(\alpha, \beta) = 1$ if $\Sigma = S_{1,1}$; $\iota(\alpha, \beta) = 2$ if $\Sigma = S_{0,4}$). We write $\sigma = \sigma_\Sigma$ for the combinatorial metric on $\mathbb{G}(\Sigma)$. A key result tells us that $\mathbb{G}(\Sigma)$ is hyperbolic [MasM1]. (In fact the curve graphs are uniformly hyperbolic [Ao, ClRS, HePW, Bo3].) Given finite subsets, $a, b \subseteq \mathbb{G}^0(\Sigma)$, we will write $\iota(a, b) = \max\{\iota(\alpha, \beta) \mid \alpha \in a,\ \beta \in b\}$.

A *multicurve*, a, is a non-empty set of (distinct) pairwise disjoint curves in Σ. We will sometimes abuse terminology by identifying a with a realisation thereof as a submanifold, $a \subseteq \Sigma$. We say that a is a *complete multicurve* (or "pants decomposition") if each component of $\Sigma \setminus a$ is an $S_{0,3}$.

The *pants graph* is traditionally defined as follows. Let $\mathbb{P}^0$ be the set of complete multicurves in Σ. Let $\mathbb{P} = \mathbb{P}(\Sigma)$ be the graph with vertex set $V(\mathbb{P}) = \mathbb{P}^0$, and where $a, b \in \mathbb{P}^0$ are deemed to be adjacent if there is some $\gamma \in a$ and $\delta \in b$ such that $a \setminus \gamma = b \setminus \delta = c$, 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 $\Sigma \setminus c$. Note that if Σ is an $S_{1,1}$ or $S_{0,4}$, then we can identify $\mathbb{P}(\Sigma)$ with $\mathbb{G}(\Sigma)$. In this case, $\mathbb{P}(\Sigma) = \mathbb{G}(\Sigma)$ is a Farey graph.

In all cases, $\mathbb{P}(\Sigma)$ is connected. This was originally proven in [HaT] for closed surfaces, though it is easily seen that the same argument works for compact surfaces. In fact, the distance between two elements, $a, b \in \mathbb{P}(\Sigma)$, is bounded above by some function of their intersection number, $\iota(a, b)$. To see this, note that up to the action of $\mathrm{Map}(\Sigma)$, there are only finitely many possibilities for the pair a, b for any given bound on $\iota(a, b)$; and so by the connectedness of $\mathbb{P}(\Sigma)$, they are connected by a path of bounded length.

We see that, from a coarse geometric perspective, the notion of the pants graph is quite robust. For example, if $q \geq 2$, let $\mathbb{P}(\Sigma, q)$ be the graph with vertex set $\mathbb{P}^0$, and where x, y are adjacent if $\iota(x, y) \leq q$. Thus, $\mathbb{P} \subseteq \mathbb{P}(\Sigma, q)$, and one sees from the observation of the previous paragraph, that the inclusion is a quasi-isometry. (We will implicitly make use of this fact later, since various constructions are well defined up to bounded intersection, hence up to bounded distance.)

Note that there is a map, $\chi_\Sigma : \mathbb{P}(\Sigma) \longrightarrow \mathbb{G}(\Sigma)$, obtained by selecting one curve from the multicurve. It is well defined up to bounded distance (1, in fact). Also, up to bounded distance, it is also characterised by the fact that $\iota(\tau, \chi_\Sigma(\tau))$ is bounded (since distances in the curve graph are also bounded in terms of intersection number).

We write $\mathcal{X}$ for the set of subsurfaces of Σ, where a *subsurface* here is required to have no homotopically trivial boundary components, and not be $S_{0,3}$, or a peripheral annulus. Also, we regard a subsurface as being defined up to homotopy. We will sometimes abuse notion by identifying a subsurface, X, with a realisation thereof in Σ, in which case, we take each component of ∂X to be either a component of $\partial\Sigma$ or a non-peripheral curve disjoint from $\partial\Sigma$. We decompose $\mathcal{X} = \mathcal{X}_A \sqcup \mathcal{X}_N$ into annular and non-annular surfaces. There is a natural identification of $\mathcal{X}_A$ with $\mathbb{G}^0(\Sigma)$, where an annulus gets identified with a core curve thereof.

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

(1): $X = Y$.
(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.

We write $X \preceq Y$ to mean that $X \prec Y$ or $X = Y$.

Given $\gamma \in \mathbb{G}^0(X)$, we write $X(\gamma) \in \mathcal{X}_A$ for the regular neighbourhood of γ. We write $\gamma \prec X$ to mean that $X(\gamma) \prec X$ etc.

Given $X \in \mathcal{X}_N$, we will write $\mathbb{G}(X)$ and $\mathbb{P}(X)$, respectively, for the curve graph and pants graph defined intrinsically to X. Note that we have an intrinsically defined map $\chi_X : \mathbb{P}(X) \longrightarrow \mathbb{G}(X)$.

We will also need to define subsurface projection for complete multicurves. We begin with a more general discussion.

Given $\alpha \in \mathbb{G}^0(\Sigma)$ and $X \in \mathcal{X}_N$, we will generally assume that α and X are realised minimally, that is, in general position so that the number of components of $\alpha \cap X$ is minimal (or equivalently that $|\alpha \cap \partial X|$ is minimal). In this case, we write $\zeta_X(\alpha)$ for the set of components of $\alpha \cap X$ viewed as defined up to homotopy sliding the endpoints of arcs in the respective boundary components of X. (In other words, parallel curves in $\alpha \cap X$ get identified in $\zeta_X(\alpha)$.) If α is homotopic into X, then we take $\zeta_X(\alpha) = \{\alpha\}$. The set $\zeta_X(\alpha)$ is well defined (irrespective of the realisation) and has cardinality bounded in terms of $\xi(\Sigma)$.

Given $\alpha, \beta \in \mathbb{G}^0(\Sigma)$, we can define the *relative intersection number*, $\iota_X(\alpha, \beta)$, to be the maximal intersection number of $\delta \in \zeta_X(\alpha)$ and $\epsilon \in \zeta_X(\beta)$. (Here the intersection number of δ, ϵ is the minimal cardinality of $\delta \cap \epsilon$ among realisations in the respective homotopy classes.)

We can apply the same definitions to multicurves, a, b in Σ. (Here $\zeta_X(a)$ includes any component of a homotopic into X.) Directly form the definitions, we get $\iota_X(a,b) = \max\{\iota_X(\alpha,\beta) \mid \alpha \in a,\ \beta \in b\}$. Note that if $Y \in \mathcal{X}_N$, with $Y \preceq X$, then $\iota_Y(a,b) \leq \iota_X(a,b)$. Also, if a, b both lie in X, then $\iota_X(a,b) = \iota(a,b)$. In practice, at least one of a, b will always lie entirely in the subsurface X.

Now, if a is a multicurve in Σ, write $a_X = a \cap \mathbb{G}^0(X) \subseteq \zeta_X(a)$. We also write $a_{\partial X} = a \cap \partial X$. Suppose that $a \in \mathbb{P}^0(\Sigma)$ is a complete multicurve. Then each component of $X \setminus a$ is a disc; an annulus with one boundary component in a_X; or an $S_{0,3}$, with all boundary components in $a_X \cup a_{\partial X}$.

The main goal of the following discussion will be to construct a map $\psi_{YX} : \mathbb{P}(X) \longrightarrow \mathbb{P}(Y)$ whenever $Y \preceq X$. (This will be an example of a "coarse gate map" as defined in Section 4.) For simplicity, we first describe $\psi_Y = \psi_{Y\Sigma} : \mathbb{P}(\Sigma) \longrightarrow \mathbb{P}(Y)$.

Suppose that $Z \in \mathcal{X}_N$ is a component of $Y \setminus a_Y$ (not an $S_{0,3}$). If $\beta, \gamma \in \mathbb{G}^0(Z)$, then $\iota(\beta,\gamma)$ is bounded above in terms of $\max\{\iota_Z(a,\beta), \iota_Z(a,\gamma)\}$ (since the latter bounds the number of possible intersections of β and γ in any component of $Z \setminus a$). Moreover, we can always find some $b_Z \in \mathbb{P}^0(Z)$ with $\iota_Z(a, b_Z) \leq 2$. (For us, it would be enough for this to be bounded in terms of $\xi(\Sigma)$. This is easy to see, given that there are only finitely many possibilities for the homotopy class of $\zeta_Z(a)$ up to self-homeomorphism of Z.) We choose such a b_Z for each such Z, and set $c_Y = a_Y \cup \bigcup_Z b_Z \in \mathbb{P}^0(Y)$. This is well defined up to bounded intersection in Y. We write $\psi_Y(a) = c_Y$.

In fact, if $X, Y \in \mathcal{X}_N$, with $Y \preceq X$, we can perform this construction intrinsically to X to give us an element $\psi_{YX}(a) \in \mathbb{P}^0(Y)$ for any $a \in \mathbb{P}^0(X)$. We note:

Lemma 2.1. *Suppose that $X, Y \in \mathcal{X}_N$ with $Y \preceq X$ and $a \in \mathbb{P}^0(X)$, $b \in \mathbb{P}^0(Y)$. Then $\rho_Y(b, \psi_{YX}a)$ is bounded above in terms of $\xi(\Sigma)$ and $\iota_Y(a,b)$.*

Proof. From the construction, it is easily seen that after applying Dehn twists about curves in a_Y to the multicurve b, we can arrange that $\iota_Y(b, \psi_{YX}a)$ is bounded in terms of $\iota_Y(a,b)$ and $\xi(\Sigma)$. From this the statement follows. □

Suppose $a, b \in \mathbb{P}^0(X)$. Then it is a simple exercise to find some $c \in \mathbb{P}^0(Y)$, with $\iota_Y(a,c)$ and $\iota_Y(b,c)$ both bounded above in terms of $\iota_Y(a,b) \leq \iota_X(a,b)$ and $\xi(\Sigma)$. By Lemma 2.1, this bounds $\rho_Y(c, \psi_{YX}a)$ and $\rho_Y(c, \psi_{YX}b)$, hence also $\rho_Y(\psi_{YX}a, \psi_{YX}b)$. It now follows that the map $\psi_{YX} : \mathbb{P}^0(X) \longrightarrow \mathbb{P}^0(Y)$ is coarsely lipschitz, and so extends to a coarsely lipschitz map $\psi_{YX} : \mathbb{P}(X) \longrightarrow \mathbb{P}(Y)$. We can think of this map as defining "subsurface projection" for the pants graphs.

Suppose $Z \in \mathcal{X}_N$ with $Z \preceq Y$. Since ψ_{ZY} is also coarsely lipschitz, we have that $\rho_Z(\psi_{ZY}c, \psi_{ZY}\psi_{YX}a)$ is bounded. In fact, we can extend c arbitrarily to a multicurve $d \in \mathbb{P}(X)$, and by construction, we can suppose that $\psi_{ZX}d = \psi_{ZY}c$. Now since $\iota_Z(a,d) = \iota_Z(a,c) \leq \iota_Y(a,c)$ is bounded, by the previous paragraph again,

we see that $\rho_Z(\psi_{ZX}d, \psi_{ZX}a)$ is bounded. It follows that $\rho_Z(\psi_{ZX}a, \psi_{ZY}\psi_{YX}a)$ is bounded. Here all bounds depend only on $\xi(\Sigma)$.

We deduce:

Lemma 2.2. *Given $X, Y \in \mathcal{X}_N$ with $Y \preceq X$, the map $\psi_{YX} : \mathbb{P}(X) \longrightarrow \mathbb{P}(Y)$ is coarsely lipschitz. Moreover, if $Z \in \mathcal{X}_N$, with $Z \preceq Y$, then $\rho_Z(\psi_{ZX}a, \psi_{ZY}\psi_{YX}a)$ is bounded for all $a \in \mathbb{P}(X)$. Here all bounds depend only on $\xi(\Sigma)$.*

Given $X, Y \in \mathcal{X}_N$ with $Y \preceq X$, we write $\theta_{YX} = \chi_Y \circ \psi_{YX} : \mathbb{P}(X) \longrightarrow \mathbb{G}(Y)$. Note that if $a \in \mathbb{P}^0(X)$ and $\gamma \in a$, then $\theta_{YX}a$ has bounded intersection with any component of $\gamma \cap X$. This means that θ_{YX} agrees, up to bounded distance, with the usual subsurface projection to $\mathbb{G}(Y)$ as defined in [MasM2].

We will often abbreviate ψ_{YX} to ψ_Y and θ_{YX} to θ_Y, where there is no confusion regarding the domain. We will abbreviate $\rho_X(a, b) = \rho_X(\psi_X a, \psi_X b)$ and $\sigma_X(a, b) = \sigma_X(\theta_X a, \theta_X b)$ for $a, b \in \mathbb{P}(\Sigma)$ etc.

The following version of the "distance formula" was described in [MasM2].

Given $a, b \in \mathbb{P}(\Sigma)$ and $r \geq 0$, write

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

Theorem 2.3. *There is some r_0 depending only on $\xi(\Sigma)$, such that given any $r \geq r_0$, there exist $k_1 > 0, k_2, k_3, k_4 \geq 0$ such that for all $X \in \mathcal{X}_N$ and all $a, b \in \mathbb{P}(X)$, we have*

$$k_1\rho_X(a, b) - k_2 \leq \sum_{Y \in \mathcal{A}_X(a,b;r)} \sigma_Y(a, b) \leq k_3\rho_X(a, b) + k_4.$$

Some other facts are also immediate consequences of more general statements about subsurface projection. First note that if $X, Y \in \mathcal{X}_N$ with $Y \prec X$ or $Y \pitchfork X$, then we have the standard subsurface projection $\theta_Y X \in \mathbb{G}^0(Y)$, as usual defined up to bounded distance.

We have the following immediate consequence of the Bounded Geodesic Image Theorem of [MasM2]. If $X, Y \in \mathcal{X}_N$, $Y \prec X$, $a, b \in \mathbb{P}(X)$, then $\min\{\sigma_X(a, \partial_X Y) + \sigma_X(b, \partial_X Y) - \sigma_X(a, b),\, \sigma_Y(a, b)\} \leq r_0$. (Here $\partial_X Y$ denotes the relative boundary of Y in X. It has diameter at most 1 in $\mathbb{G}(X)$.) Also the following is a consequence of Behrstock's Lemma [Be] (see also [Man]). If $X, Y, Z \in \mathcal{X}_N$, with $Y \prec X$, $Z \prec X$, $Y \pitchfork Z$ and $x \in \mathbb{P}^0(X) \cup \mathcal{X}_N$, then $\min\{\sigma_Y(x, \partial_Y Z), \sigma_Z(x, \partial_Z Y)\} \leq r_0$. Here, $r_0 \geq 0$ depends only on $\xi(\Sigma)$.

For later reference, we recall a few facts regarding the low complexity cases.

In the cases where $\xi(\Sigma) \leq 2$, we have already noted that $\mathbb{P}(\Sigma)$ (or equivalently $\mathbb{W}(\Sigma)$) is Gromov hyperbolic. When $\xi(\Sigma) = 1$, $\mathbb{P}(\Sigma)$ is the Farey graph, hence a quasitree, and its Gromov boundary, $\partial\mathbb{P}(\Sigma) \cong \partial\mathbb{W}(\Sigma)$, is homeomorphic to the set of irrational numbers. When $\xi(\Sigma) = 2$, however, $\partial\mathbb{W}(\Sigma)$ is connected. In fact, it follows from the description in [BroMas], that it naturally contains the boundary of the corresponding curve graph as a dense subspace. This was, in turn, shown

to be homeomorphic to the Nöbeling curve in [HeP]. In particular, $\mathbb{W}(\Sigma)$ is not a quasitree.

3. Coarse median spaces

In this section, we recall the notion of a coarse median space. We show that $\mathbb{W}(\Sigma)$ is naturally such a space. We give proofs of Theorems 1.1, 1.2 and 1.3, and derive some further results.

Recall that a *median algebra* is a set, M, equipped with a ternary operation $\mu : M^3 \longrightarrow M$, satisfying $\mu(a,b,c) = \mu(b,c,a) = \mu(b,a,c)$, $\mu(a,a,b) = a$ and $\mu(a,b,\mu(c,d,e)) = \mu(\mu(a,b,c),\mu(a,b,d),e)$ for all $a,b,c,d,e \in M$. (See for example, [Is, BaH, Ve, R, Bo1] for further discussion.) A *subalgebra* is a subset closed under μ. An *n-cube* is a (sub)algebra isomorphic to the direct product $\{-1,1\}^n$. Here, $\{-1,1\}$ is equipped with its unique median structure ("majority vote"), and the median on the product $\{-1,1\}^n$ is defined on each factor independently. The *rank* of M is the maximal n such that M contains an n-cube. (This will always be finite for the median algebras we deal with in this paper.) Given $a,b \in M$, write $[a,b] = \{x \in M \mid \mu(a,b,x) = x\}$ for the median *interval* from a to b. A subset $C \subseteq M$ is *convex* if $[a,b] \subseteq C$ for all $a,b \in C$. Intervals are always convex.

A *median metric space* is (equivalent to) a median algebra, M, equipped with a metric, ρ, such that if $a,b,c \in M$, then $c \in [a,b]$ if and only if $\rho(a,b) = \rho(a,c) + \rho(c,b)$. (See, for example, [Ve, ChaDH, Bo4].) The vertex set of a finite CAT(0) cube complex has a natural structure as a median algebra. Indeed it is a median metric space with the metric induced by the combinatorial metric on the 1-skeleton. It turns out that every finite median algebra arises in this way [Che]. In this case, the rank of the median algebra equals the dimension of the cube complex.

If $Q \subseteq M$ is an n-cube in a connected median metric space, M, then its convex hull, $\mathrm{hull}(Q)$, is a median direct product of n real intervals. In fact, it is isometric to such in the l^1-metric. Clearly, Q itself corresponds to the "corners" of $\mathrm{hull}(Q)$.

Suppose now that (Λ,ρ) is a geodesic metric space. As in [Bo1], we define a *coarse median* on Λ to be a ternary operation $\mu : \Lambda^3 \longrightarrow \Lambda$, satisfying:

(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, $h : \mathbb{N} \longrightarrow [0,\infty)$, 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 $\pi : A \longrightarrow \Pi$ and $\lambda : \Pi \longrightarrow \Lambda$ 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.

A map, $\phi : \Lambda \longrightarrow \Lambda'$ between two coarse median spaces is an *h-quasimorphism* if $\rho'(\phi\mu(a,b,c), \mu'(\phi a, \phi b, \phi c)) \leq h$ for all $a, b, c \in \Lambda$. We will use the same terminology when the domain is a median algebra.

Further discussion of coarse medians can be found in [Z1, NWZ].

We recall notion of an asymptotic cone [VaW, G]. Here, we use the notation of [Bo1, Bo8]. Given a countable set, $\mathcal{Z}$, equipped with a non-principal ultrafilter, and a $\mathcal{Z}$-sequence, $(\Lambda_\zeta, \rho_\zeta)_\zeta$, of metric spaces, we obtain a limiting space, $(\Lambda^\infty, \rho^\infty)$, which is a complete metric space. In particular, if $\Lambda_\zeta = \Lambda$ is constant, ρ is a fixed metric on Λ, and $(t_\zeta)_\zeta$ is a $\mathcal{Z}$-sequence of positive numbers (or *scaling factors*) with $t_\zeta \to 0$ (with respect to the ultrafilter) then we can set $\rho_\zeta = t_\zeta \rho$, to get a limiting space, $(\Lambda^\infty, \rho^\infty)$. This is referred to as an *asymptotic cone* of (Λ, ρ) [VaW, G]. If μ is a coarse median on Λ, then we get a limiting lipschitz ternary operation, $\mu^\infty : (\Lambda^\infty)^3 \longrightarrow \Lambda^\infty$, so that $(\Lambda^\infty, \mu^\infty)$ is a median algebra.

If (Λ, ρ, μ) has rank at most n, then $(\Lambda^\infty, \rho^\infty)$ has rank at most n (as a median algebra). It has locally compact dimension at most n (that is, any locally compact subset has dimension at most n). From this, it follows that Λ does not admit any quasi-isometric embedding of $\mathbb{R}^{n+1}$ [Bo1]. Moreover, ρ^∞ is bilipschitz equivalent to a median metric, inducing the same median structure [Bo2, Bo8]. (A more canonical construction is given in [Z2].) In fact, ρ^∞ is also bilipschitz equivalent to a CAT(0) metric [Bo4] (though this information is redundant in the case where $\Lambda = \mathbb{W}(\Sigma)$, since $\mathbb{W}(\Sigma)$ is already CAT(0), and it is easily verified that this property is preserved after taking asymptotic cones). It is an open question in this case as to whether the homeomorphism (or bilipschitz) type of asymptotic cone is independent of the ultrafilter or scaling factors. However, the results we state here hold for any asymptotic cone, and we often refer to "the" asymptotic cone.

It is easily seen that the existence of a coarse median on a geodesic space is invariant under quasi-isometry [Bo1]. We will show that $\mathbb{P}(\Sigma)$ admits such a structure, from which it will follow that $\mathbb{W}(\Sigma)$ does too. This will be a consequence of Theorem 7.2 of [Bo8], which gives hypotheses (A1)–(A10) which imply the existence of medians.

Recall that we have a collection of maps, $\psi_{YX} : \mathbb{P}(X) \longrightarrow \mathbb{P}(Y)$, for $X, Y \in \mathcal{X}_N$ with $Y \preceq X$. To be precise, the hypotheses laid out in Section 7 of [Bo8] require these to be defined for all $X \in \mathcal{X}$. Rather than reformulate the hypotheses to take care of this (which would be easy to do), we will simply set $\mathbb{P}(X) = \mathbb{G}(X)$ to be a singleton for all $X \in \mathcal{X}_A$. In this way, the relevant hypotheses relating to $\mathcal{X}_A$ become trivial, so we simply ignore them.

We now go briefly through the hypotheses. First, (A1) says that $\mathbb{G}(X)$ is uniformly hyperbolic, which is certainly true here, by [MasM1]. (There are only finitely many topological types of X in any given Σ.) Properties (A2) and (A3)

require that the maps χ_X and ψ_{YX} are uniformly coarsely lipschitz. The case of χ_X is trivial, and that of ψ_{YX} is a consequence of Lemma 2.2 here. Property (A2) also requires that χ_X is cobounded, which is immediate here. Property (A4) is the second clause of Lemma 2.2. Property (A5) is a simple and immediate consequence of subsurface projection. As noted in [Bo8], Properties (A6) and (A7) are both immediate from Property (B1) there, which is a formulation of the distance formula given as Theorem 2.3 here. Properties (A8) and (A9) are respectively consequences of the Bounded Geodesic Image Theorem and Behrstock's Lemma, as laid out at the end of Section 2. Finally, (A10) just says that we can combine complete multicurves on disjoint subsurfaces and then extend them to give a complete multicurve on the whole surface.

We have now verified the hypotheses of Theorem 7.2 of [Bo8]. We therefore get:

Theorem 3.1. *There is a coarse median, $\mu : \mathbb{P}(\Sigma)^3 \longrightarrow \mathbb{P}(\Sigma)$, with the property that for all $X \in \mathcal{X}_N$, the maps $\theta_X : \mathbb{P}(\Sigma) \longrightarrow \mathbb{G}(X)$ are uniform quasimorphisms. Here, the constants depend only on $\xi(\Sigma)$.*

Here, the median on $\mathbb{G}(X)$ is the usual "centroid" map on a hyperbolic space (which gives it the structure of a rank-1 coarse median space). The fact that the maps θ_X are uniform quasimorphisms determine the median on $\mathbb{P}(\Sigma)$ up to bounded distance, so in particular, it is necessarily $\mathrm{Map}(\Sigma)$-equivariant up to bounded distance. Note that we can perform this construction intrinsically on subsurfaces, to give a median, μ_X, on each $\mathbb{P}(X)$. Moreover the maps $\psi_{YX} : \mathbb{P}(X) \longrightarrow \mathbb{P}(Y)$ are necessarily all coarse median. (See [Bo8] for details.)

Theorem 7.2 of [Bo8] immediately implies that $\mathbb{P}(\Sigma)$ has rank at most $\xi(\Sigma)$. The proof there used Proposition 10.2 of [Bo1]. In fact, in the present situation, we see by the latter result that $\mathbb{P}(\Sigma)$ has rank at most $\xi_0(\Sigma)$. For this, note that the hypotheses (P1)–(P4) of Proposition 10.2 of [Bo1] are all satisfied. (For the purposes of this discussion, we should interpret the indexing set, $\mathcal{X}$, used in [Bo1], to be $\mathcal{X}_N$ here.) Note that (P3) calls for an upper bound, ν, on the number of disjoint elements of $\mathcal{X}_N$ we can embed disjointly in Σ. A simple topological argument gives $\nu = \xi_0(\Sigma)$ in this case (see Lemma 4.2). Proposition 10.2 of [Bo1] now tells us that $\mathbb{P}(\Sigma)$ has rank at most $\xi_0(\Sigma)$. In fact, as observed in Section 1, $\mathbb{W}(\Sigma)$ (hence $\mathbb{P}(\Sigma)$) has rank at least $\xi_0(\Sigma)$ (since it admits a (quasi)isometric embedding of $\mathbb{R}^{\xi_0(\Sigma)}$). Since these properties are quasi-isometrically invariant, they apply equally well to $\mathbb{P}(\Sigma)$ and $\mathbb{W}(\Sigma)$.

As observed above, from the general results of [Bo1], we can now deduce Theorems 1.2 and 1.5.

In fact, we can strengthen this statement. In view of Lemma 6.10 of [Bo8], we have:

Proposition 3.2. *There is a bound on the radius of a ball in $\mathbb{R}^{\xi_0+1}$ which can be quasi-isometrically embedded into $\mathbb{W}(\Sigma)$, where the bound depends only on the topological type of Σ and the parameters of quasi-isometry.*

In particular, there is no quasi-isometric embedding of the half-space, $\mathbb{R}^{\xi_0} \times [0, \infty)$ into $\mathbb{W}(\Sigma)$.

The above statements can also be viewed as consequences of the following, which will have further applications throughout this paper.

Theorem 3.3. *The asymptotic cone $\mathbb{W}^\infty(\Sigma)$ admits a canonical bilipschitz equivalent median metric (of rank ξ_0) inducing the original median structure on $\mathbb{W}(\Sigma)$ (that is, the ultralimit of the median on $\mathbb{W}(\Sigma)$).*

The existence of such a metric follows from [Bo1] and Theorem 6.9 of [Bo8]. A canonical construction thereof is given in [Z2]. (Its canonical nature is not essential for any argument of the present paper, however.)

Although it is not essential to the present paper, we note the following result about the asymptotic cone:

Theorem 3.4. *$\mathbb{W}^\infty(\Sigma)$ admits a bilipschitz embedding into a finite product of $\mathbb{R}$-trees. Moreover, we can take this embedding to be a median homomorphism with respect to the respective median structures.*

Proof. As in Section 7 of [Bo8], we see that $\mathbb{W}(\Sigma)$, hence $\mathbb{W}^\infty(\Sigma)$, is "finitely colourable", as defined in Section 12 of [Bo1]. The statement then follows by the main result of [Bo2]. □

Note that Theorem 3.3 is an immediate consequence, modulo deleting the word "canonical": the embedding involves making arbitrary choices.

In addition to the asymptotic cone, $\mathbb{W}^\infty(\Sigma)$, we will also need to consider the "extended asymptotic cone" $\mathbb{W}^*(\Sigma)$ (or equivalently $\mathbb{P}^*(\Sigma)$). This is defined in Section 5 of [Bo8]. This definition is the same as for the asymptotic cone, except we do not require distances to be finite. Formally, we can think of $\mathbb{W}^*(\Sigma)$ as an $\mathbb{R}^*$-metric space: that is, taking non-negative values in the extended reals, $\mathbb{R}^*$ (that is, the ultrapower of the reals factored out by infinitesimals). For our purposes here, we can simply think of it as taking values in $[0, \infty]$ (i.e. where we identify all positive unlimited numbers, and denote them all by ∞). Two points of $\mathbb{W}^*(\Sigma)$ lie in the same *component* if the distance between them is limited (i.e. not ∞). Thus, $\mathbb{W}^\infty(\Sigma)$ is a component, and (since $\mathbb{W}(\Sigma)$ is quasi-homogeneous — $\mathrm{Map}(\Sigma)$ acts coboundedly), every component of $\mathbb{W}^*(\Sigma)$ is an isometric copy of $\mathbb{W}^\infty(\Sigma)$. We also note that $\mathbb{W}^*(\Sigma)$ has a natural ternary operation, $\mu^* : (\mathbb{W}^*(\Sigma))^3 \longrightarrow \mathbb{W}^*(\Sigma)$, so that $\mathbb{W}^*(\Sigma)$ is a topological median algebra (that is, the median operation is continuous). As a topological (or metric) space, $\mathbb{W}^*(\Sigma)$ is just an uncountable disjoint union of copies of $\mathbb{W}^\infty(\Sigma)$. The median endows it with additional structure which interrelates these components.

We should finally note that in the cases where $\xi(\Sigma) \leq 2$, the median we have defined on $\mathbb{W}(\Sigma)$ agrees, up to to bounded distance, with the centroid operation on a hyperbolic space. (Indeed any hyperbolic space admits only one rank-1 coarse median structure up to bounded distance.)

Remark. As remarked in the introduction, it is shown in [BeHS1] that $\mathbb{P}(\Sigma)$ satisfies a stronger set of axioms, namely those of a "hierarchically hyperbolic space". The authors also verify that these in turn imply the properties (A1)–(A10) of [Bo8], and therefore recover the fact that $\mathbb{P}(\Sigma)$ is coarse median. In order to obtain the rank bound, one needs to make the additional observation regarding (P3) of [Bo1] as explained above. (As the statement is formulated in [BeHS1], their rank bound is weaker, though it seems that their argument also yields a rank bound of ξ_0.) We should also comment that their indexing set (denoted $\mathfrak{S}$) is larger than our $\mathcal{X}$, in that it includes disconnected subsurfaces (in order for it to satisfy their "orthogonality axiom"). However the curve graphs associated with non-connected subsurfaces all have bounded diameter.

4. Multicurves and product structure

To go further, we need to discuss the local "product structure" in $\mathbb{P}(\Sigma)$. Much of this fits into the more general picture (applicable to the mapping class group and to the Teichmüller metric) as described in [Bo8]. We briefly outline this.

Let $\tau \subseteq \mathbb{G}^0(\Sigma)$ be a multicurve. Let $\mathcal{X}_A(\tau) = \{X(\gamma) \mid \gamma \in \tau\} \subseteq \mathcal{X}_A$ be the corresponding annular neighbourhoods, and let $\mathcal{X}_N(\tau) \subseteq \mathcal{X}_N$ be the set of non-$S_{0,3}$ components of $\Sigma \setminus \tau$. Let $\mathcal{X}(\tau) = \mathcal{X}_A(\tau) \sqcup \mathcal{X}_N(\tau)$. If $Y \in \mathcal{X}_N$, write $\tau \pitchfork Y$ to mean that there is some $\gamma \in \tau$ with either $\gamma \pitchfork Y$ or $\gamma \prec Y$.

We write

$$T(\tau) = \{a \in \mathbb{P}^0(\Sigma) \mid \tau \subseteq a\},$$

and

$$\mathcal{T}(\tau) = \prod_{X \in \mathcal{X}_N(\tau)} \mathbb{P}(X).$$

Here, we give $\mathcal{T}(\tau)$ the l^1 product metric. Note that it also has a product ternary operation, which gives it the structure of a coarse median space. We write $\mathcal{T}^0(\tau) = \prod_{X \in \mathcal{X}_N(\tau)} \mathbb{P}^0(X)$ (which we can think of as the vertex set of $\mathcal{T}(\tau)$ viewed as a cube complex). We can identify $\mathcal{T}(\tau)$ and $T(\tau)$, simply by combining the component multicurves together with τ itself. We write $\psi_\tau : T(\tau) \longrightarrow \mathcal{T}^0(\tau)$ and $\upsilon_\tau : \mathcal{T}^0(\tau) \longrightarrow T(\tau)$ for the inverse bijections. In fact, by combining the subsurface projections ψ_X for $X \in \mathcal{X}_N(\tau)$, we get a map $\psi_\tau : \mathbb{P}(\Sigma) \longrightarrow \mathcal{T}^0(\tau)$. Write $\omega_\tau = \upsilon_\tau \circ \psi_\tau : \mathbb{P}(\Sigma) \longrightarrow T(\tau)$. (We can assume that $\omega_\tau | T(\tau)$ is just inclusion.) All the above maps are uniform quasimorphisms.

The definition of $T(\tau)$ is quite robust. Note that, up to bounded Hausdorff distance, we could have defined $T(\tau)$ to be the set of $a \in \mathbb{P}^0(\Sigma)$ with $\iota(a, \tau)$ bounded by some constant. In fact, to tie this in with the discussion in [Bo8], we

set $T(\tau;r)$ be the set of $a \in \mathbb{P}^0(\Sigma)$ such that $\sigma_Y(a,\tau) \leq r$ for all those $Y \in \mathcal{X}$ satisfying $\tau \pitchfork Y$. (Note that there is a well defined projection $\theta_Y\tau \in \mathbb{G}(Y)$ up to bounded distance.) We claim:

Lemma 4.1. *There is some $r_0 \geq 0$ such that for all $r \geq r_0$, the Hausdorff distance $\mathrm{hd}(T(\tau), T(\tau;r))$ is finite and bounded above in terms of r and $\xi(\Sigma)$.*

Proof. Clearly, if $a \in T(\tau)$, then $\theta_Y(a,\tau)$ is bounded for all $Y \in \mathcal{X}_N$.

Conversely, suppose $a \in T(\tau;r)$. By construction, $\omega_\tau a \in T(\tau)$. We claim that $\sigma_Y(a,\omega_\tau a)$ is bounded for all $Y \in \mathcal{X}_N$. The statement then follows from the distance formula (given as Theorem 2.3 here). If $\tau \pitchfork Y$, this claim is immediate given that $\sigma_Y(a,\tau)$ and $\sigma_Y(\tau,\omega_\tau a)$ are bounded. Otherwise, $Y \preceq X$ for some $X \in \mathcal{X}_N(\tau)$. By definition of ω_τ, we have that $\rho_X(a,\omega_\tau a)$ is bounded, so $\sigma_Y(a,\omega_\tau a)$ is bounded as required. □

Note also that the constructions of υ_τ, ψ_τ and ω_τ agree, up to bounded distance, with those described in Section 9 of [Bo8], so the following are direct consequences of the more general discussion there.

First, $\omega_\tau : \mathbb{P}(\Sigma) \longrightarrow T(\tau)$ is a "coarse gate map". This means that $\rho(\omega_\tau a, \mu(a,\omega_\tau a, c))$ is uniformly bounded for all $a \in \mathbb{P}(\Sigma)$ and $c \in T(\tau)$. It follows that $T(\tau)$ is uniformly (median) quasiconvex in $\mathbb{P}(\Sigma)$; that is, if $a,b \in T(\tau)$ and $c \in \mathbb{P}(\Sigma)$, then $\rho(c,T(\tau))$ is bounded above in terms of $\rho(c,\mu(a,b,c))$. From this, it follows that the map $\upsilon_\tau : \mathcal{T}(\tau) \longrightarrow T(\tau) \subseteq \mathbb{P}(\Sigma)$ is a quasi-isometric embedding of $\mathcal{T}(\tau)$ into $\mathbb{P}(\Sigma)$.

To proceed we make a few topological observations about Σ.

We first note that we can embed at most $\xi_0(\Sigma)$ disjoint surfaces of complexity at least 1 in Σ:

Lemma 4.2. *Suppose that $\mathcal{Y} \subseteq \mathcal{X}_N$ satisfies $X \wedge Y$ for all distinct $X, Y \in \mathcal{Y}$. Then $|\mathcal{Y}| \leq \xi_0(\Sigma)$.*

Proof. We write $e(X)$ for minus the Euler characteristic of X, so that $e(\Sigma) = 2g+p-2$. It is easily seen that $\sum_{Y\in\mathcal{Y}} e(Y) \leq e(\Sigma)$. Write $\mathcal{Y}_1 \subseteq \mathcal{Y}$ for the set of $S_{1,1}$ elements of $\mathcal{Y}$, and write $\mathcal{Y}_2 = \mathcal{Y}\backslash\mathcal{Y}_1$. Thus, $e(Y) = 1$ for all $Y \in \mathcal{Y}_1$, and $e(Y) \geq 2$ for all $Y \in \mathcal{Y}_2$. Write $A = |\mathcal{Y}_1|$ and $B = |\mathcal{Y}_2|$. We have $A \leq g$ and $A + 2B \leq e(\Sigma)$, so $2(A+B) \leq g + e(\Sigma) = \xi(\Sigma) + 1$, and so $|\mathcal{Y}| = A + B \leq \lfloor(\xi(\Sigma)+1)/2\rfloor = \xi_0(\Sigma)$ as required. □

Note that it follows that if τ is any multicurve in Σ then $|\mathcal{X}_N(\tau)| \leq \xi_0(\Sigma)$. (Recall that the notation $\mathcal{X}_N$ excludes all $S_{0,3}$'s.)

Definition. We say that a multicurve τ is *good* if $\Sigma \setminus \tau$ has exactly $\xi_0(\Sigma)$ components, none of which is an $S_{0,3}$, nor an $S_{0,4}$ with two boundary components identified to a single component of τ.

(If we had such an $S_{0,4}$, we should delete this component of τ as to give us an $S_{1,2}$ instead.)

For much of the discussion in Section 5, we will need to split into odd and even cases. (The odd case being somewhat simpler.)

Definition. We say that Σ is *odd* (respectively *even*) if $\xi(\Sigma)$ is odd (respectively even).

In other words, $S_{g,p}$ is odd precisely when $g+p$ is even.

We distinguish the following types of multicurves. In the following, τ is assumed to be a good multicurve.

(T0): Σ is odd and each component of $\Sigma \setminus \tau$ has complexity 1.

(T1): Σ is even and $\Sigma \setminus \tau$ has exactly one component of complexity 2, and all other components have complexity 1.

(T2): Σ is even and each component of $\Sigma \setminus \tau$ has complexity 1.

(Recall that complexity-1 corresponds to $S_{1,1}$ or $S_{0,4}$, and complexity-2 corresponds to $S_{1,2}$ or $S_{0,5}$.)

In case (T1), we will write $W(\tau) \in \mathcal{X}_N(\tau)$ for the complexity-2 component of $\Sigma \setminus \tau$.

Lemma 4.3. *Any good multicurve is one of the types (T0), (T1) or (T2) described above.*

Proof. Let $\mathcal{Y} = \mathcal{X}_N(\tau)$. We use the notation of the proof of Lemma 4.2. Recall that $A \leq g$ and $A + 2B \leq 2g + p - 2$. Let $Z = \Sigma \setminus \bigcup \mathcal{Y}_1$. This is a connected non-empty surface (provided Σ is not $S_{2,0}$).

If Σ is odd, then $2(A+B) \leq 3g+p-2$. We see that $A+B$ is maximised when $A = g$, and then Z is an $S_{0,g+p}$. We can cut Z into $(g+p-2)/2$ $S_{0,4}$'s, and so the maximum is realised with $B = (g+p-2)/2$. This is type (T0).

If Σ is even, then $2(A+B) \leq 3g+p-3$. In this case, the maximum is realised when $A \geq g-1$. If $A = g$, then Z is an $S_{0,g+p}$ which we can cut into $(g+p-5)/2$ $S_{0,4}$'s together with an $S_{0,5}$, giving the maximal $B = (g+p-3)/2$. This is type (T1). If $A = g-1$, then Z is an $S_{1,g+p-1}$ which we can cut either into $(g+p-1)/2$ $S_{0,4}$'s (type (T2)), or else into $(g+p-3)/2$ $S_{0,4}$'s and one $S_{1,2}$ (type (T1) again). Either way, we again get a maximal $B = (g+p-1)/2$. □

If τ is of type (T0), the complement of the union of the $S_{1,1}$ components of $\Sigma \setminus \tau$ is a planar (genus 0) surface (or empty in the case where Σ is an $S_{2,0}$). In case (T2) this complement has genus 1. The set of non-separating curves of τ are cyclically arranged in this subsurface. In case (T1), the complement of the $S_{1,1}$ and $S_{1,2}$ components is a union of planar surfaces. In all cases, we see that no good multicurve can strictly contain another.

Now if τ is a good multicurve, then all the factors of $\mathcal{T}(\tau)$ are hyperbolic spaces. In fact, each is a *bushy* hyperbolic space in the sense that every point is (a bounded distance from) the centroid of three ideal points. In summary, we have shown:

Proposition 4.4. *If τ is a good multicurve, then $T(\tau)$ is a quasi-isometrically embedded copy of a direct product of $\xi_0(\Sigma)$ bushy hyperbolic spaces. Moreover, the embedding is a median quasimorphism (i.e. it respects the coarse median structures up to bounded distance).*

We will need to understand better how product regions, corresponding to good multicurves, are arranged in $\mathbb{P}(\Sigma)$. This will be discussed in Section 5. For the moment, we just give two Lemmas (4.5 and 4.6) describing how product regions diverge. This will be needed for the proof of Lemma 5.10. Analogous statements can be found in Section 9 of [Bo8].

Lemma 4.5. *Suppose that τ, τ' are good multicurves with $\tau \neq \tau'$. There are constants, $k, t \geq 0$, depending only on $\xi(\Sigma)$, such that if $a \in T(\tau')$ and $r \geq 0$, then there is some $b \in T(\tau')$ with $\rho(b, T(\tau)) \geq r$ and $\rho(a, b) \leq kr + t$.*

Proof. In fact, we show that for all $l \geq 0$, we can find $b \in T(\tau')$ with $l = \rho(a, b)$ and with $l \leq k\rho(b, c) + t$ for all $c \in T(\tau)$. The statement then follows setting $l = kr + t$.

In what follows, $\sim$ will denote "up to an additive constant" or "up to bounded distance", where the bound only depends on $\xi(\Sigma)$. Also all "linear bounds" referred to are assumed to depend only on $\xi(\Sigma)$.

First note that there is some complexity-1 subsurface, $X \in \mathcal{X}_N(\tau')$, which intersects τ non-trivially. (Otherwise $\tau \subseteq \tau'$, so $\tau = \tau'$.) In particular, $\theta_X \tau \in \mathbb{G}(X)$ is defined. Moreover, $\theta_X c \sim \theta_X \tau$ for all $c \in T(\tau)$.

Let $\alpha \subseteq a$ be the component of a lying in the interior of X (so $\theta_X a = \alpha$). Now $\mathbb{G}(X)$ is a Farey graph. In particular, we can find a geodesic ray, $\pi \subseteq \mathbb{G}(X)$, emanating from α, such that for all vertices, $\beta \in \pi$, we have $\sigma_X(\theta_X \tau, \beta) \sim \sigma_X(\theta_X \tau, \alpha) + \sigma_X(\alpha, \beta)$.

Given such β, let $b = (\tau' \setminus \alpha) \cup \beta \in T(\tau')$ (so $\theta_X b = \beta$). Note that for all $Y \in \mathcal{X}_N \setminus \{X\}$, $\theta_Y b \sim \theta_Y a \in \mathbb{G}(Y)$. Therefore, by the distance formula (given as Theorem 2.3 here) $\rho(a, b)$ agrees with $\sigma_X(a, b) = \sigma_X(\alpha, \beta)$, up to linear bounds.

Moreover, again by the distance formula, $\sigma_X(b, c)$, is bounded above by a linear function of $\rho(b, c)$. Since $\theta_X c \sim \theta_X \tau$, it follows that $\sigma_X(a, b)$ is also. It follows that $\sigma(a, b) \leq k\rho(b, c) + t$, where k, t depend only on $\xi(\Sigma)$.

Now, by continuity, we can arrange that $\sigma_X(a, b)$ takes whatever positive (integer) value we want (since π also gives us a path in $T(\tau') \subseteq \mathbb{P}(\Sigma)$, and we have seen that $\sigma_X(a, b)$ hence also $\rho(a, b)$ can be made arbitrarily large). Therefore, we can take $\rho(a, b) = l$, so $l \leq k\rho(b, c) + t$ as claimed. □

We will also need:

Lemma 4.6. *Suppose that τ, τ' are good multicurves which together fill Σ. Then $T(\tau)$ and $T(\tau')$ uniformly diverge.*

The hypothesis means that $\tau \cup \tau'$ cuts Σ into discs and peripheral annuli, and the conclusion means that, for any $r \geq 0$, the diameter of $N(T(\tau), r) \cap N(T(\tau'), r)$

is bounded above in terms of r and $\xi(\Sigma)$. Lemma 4.6 follows exactly as in Lemma 9.6 of [Bo8].

We now move on to discuss the extended asymptotic cone, $\mathbb{P}^*(\Sigma)$, and asymptotic cone, $\mathbb{P}^\infty(\Sigma) \subseteq \mathbb{P}^*(\Sigma)$. The asymptotic cone is a well established notion (see [VaW, G]). The "extended" version thereof is used in [Bo8] (see the discussion at the end of Section 3 of the present paper).

As in [Bo8], we write $\mathcal{U}\mathbb{G}^0 \supseteq \mathbb{G}^0$ and $\mathcal{U}\mathcal{X} \supseteq \mathcal{X}$ for the ultrapoducts of $\mathbb{G}^0$ and $\mathcal{X}$, etc. We have $\mathcal{U}\mathcal{X} = \mathcal{U}\mathcal{X}_N \sqcup \mathcal{U}\mathcal{X}_A$. Given $\gamma \in \mathcal{U}\mathbb{G}^0$, we will generally write $(\gamma_\zeta)_\zeta$ for the corresponding $\mathcal{Z}$-sequence in $\mathbb{G}^0$ etc. We extend the notation, $X \preceq Y$, $X \wedge Y$, $X \pitchfork Y$ etc. to $\mathcal{U}\mathcal{X}$. If $X \in \mathcal{U}\mathcal{X}$, we have limiting spaces, $\mathbb{P}^*(X)$, $\mathbb{G}^*(X)$ etc.

In what follows we will refer to elements of $\mathcal{U}\mathbb{G}^0$, $\mathcal{U}\mathcal{X}$ as "curves" and "subsurfaces", and to elements of $\mathbb{G}^0$ and $\mathcal{X}$ as "standard curves" and "standard subsurfaces". We can also speak of "multicurves" and "standard multicurves" etc.

Note that it makes sense to talk of the "topological type" and "complexity" of (non-standard) subsurfaces. Since there are only finitely many topological types possible in Σ, these will always be standard. (A $\mathcal{Z}$-sequence taking finitely many values is almost always constant.) We can also refer to complementary components of multicurves, and so on. These will again be (non-standard) subsurfaces.

Now suppose that $\tau \subseteq \mathcal{U}\mathbb{G}^0$ is a multicurve (in the above sense). We get a limiting closed subset, $T^*(\tau) \subseteq \mathbb{P}^*(\Sigma)$, and uniformly lipschitz maps, $\psi^*_\tau : \mathbb{P}^*(\Sigma) \longrightarrow \mathcal{T}^*(\tau)$, $\upsilon^*_\tau : \mathcal{T}^*(\tau) \longrightarrow T^*(\tau)$, and $\omega^*_\tau = \upsilon^*_\tau \circ \psi^*_\tau : \mathbb{P}^*(\Sigma) \longrightarrow T^*(\tau)$. Recall that the spaces $T(\tau_\zeta) \subseteq \mathbb{P}(\Sigma)$ are uniformly quasiconvex, and that the maps $\omega_{\tau,\zeta} : \mathbb{P}(\Sigma) \longrightarrow T(\tau_\zeta)$ are uniform gate maps. It follows that $T^*(\tau)$ is convex in $\mathbb{P}^*(\Sigma)$, and that ω^*_τ is a gate map (that is, $\omega^*_\tau x \in [x, c]$ for all $x \in \mathbb{P}^*(\Sigma)$ and all $c \in T^*(\tau)$). We also note that $T^*(\tau)$ is naturally median isomorphic to the direct product, $\prod_{X \in \mathcal{U}\mathcal{X}_N(\tau)} \mathbb{P}^*(X)$. Indeed the isomorphism is bilipschitz with respect to the l^1 metric on the product. (Here $\mathcal{U}\mathcal{X}_1(\tau)$ denotes the set of complexity-1 complementary components — in the "non-standard" sense alluded to above.)

Restricting to $\mathbb{P}^\infty(\Sigma)$, we write $T^\infty(\tau) = T^*(\tau) \cap \mathbb{P}^\infty(\Sigma)$. If this is non-empty, then $T^\infty(\tau) = \prod_{X \in \mathcal{U}\mathcal{X}_N(\tau)} \mathbb{P}^\infty(X)$. Now, if τ is good, then each $\mathbb{P}(X_\zeta)$ is (almost always) a bushy hyperbolic space, and it follows that $\mathbb{P}^\infty(X)$ is the complete $2^{\aleph_0}$-regular $\mathbb{R}$-tree. (This is proven for Hadamard manifolds and for hyperbolic spaces with cobounded group actions in [DyP] and [DrK], though only the requirement that the space is bushy is needed for these arguments to go through.) In particular, $\mathbb{P}^\infty(X)$ is *furry* (that is, each point has valence at least 3). In summary, we have shown:

Lemma 4.7. *If $\tau \subseteq \mathcal{U}\mathbb{G}^0$ is a good multicurve, then $T^\infty(\tau)$ is a median direct product of $\xi_0(\Sigma)$ furry $\mathbb{R}$-trees.*

The following fact from [Bo8] will allow us to recognise such products topologically.

Proposition 4.8. *Suppose that D is a direct product of $\xi_0(\Sigma)$ furry $\mathbb{R}$-trees, and that $f : D \longrightarrow \mathbb{P}^\infty(\Sigma)$ is a continuous injective map with closed image. Then f is a median homomorphism, and $f(D) \subseteq \mathbb{P}^\infty(\Sigma)$ is convex.*

Proof. Noting that $\mathbb{P}^\infty(\Sigma)$ is bilipschitz equivalent to a median metric space of rank $\xi_0(\Sigma)$, this is an immediate consequence of Proposition 4.7 of [Bo8]. □

We also remark that, by the results of [KaKL] (see also [Bo5]) the product structure of D, hence that of $f(D)$, are also determined by its topology.

5. Recognising product regions from the coarse geometry

As a brief summary of the last section, we showed that a multicurve, τ, gives rise to a product space, $T(\tau)$, inside $\mathbb{P}(\Sigma)$. This has a maximal number of factors when τ is good, and in this case, each factor is hyperbolic. We will refer loosely to such sets as "maximal product regions". (A formal definition of this notion would require careful quantification, which will not be logically necessary for our proofs.)

In this section, we aim to show that such maximal product regions are recognisable in terms of the coarse geometry of $\mathbb{P}(\Sigma)$. Much of the work is carried out in the asymptotic cone, where these regions correspond to maximal products of trees, and we can bring Proposition 4.8 into play. The description depends on the combinatorial structure of τ, which depends on the parity of $\xi(\Sigma)$. The odd case is simpler, and we deal with that first. To get the general idea, one could just ignore the even case described in Section 5.3, on a first reading.

We begin with a discussion of quasicubes in Section 5.1. We will prove Theorem 1.6 in Section 5.4.

Throughout this section, we assume that $\xi(\Sigma) \geq 3$, unless otherwise stated.

5.1. **Quasicubes.**

Let $Q = \{-1, 1\}^{\xi_0}$ be a ξ_0-cube. (In fact, all "cubes" in this section will be ξ_0-cubes.) By an *ith side* of Q we mean a pair of elements which differ precisely in their ith coordinates. Suppose that $\phi : Q \longrightarrow \mathbb{P}(\Sigma)$ is an h-quasimorphism.

Recall, from Section 2, that if $a, b \in \mathbb{P}(\Sigma)$, then $\mathcal{A}(a, b; r) = \{X \in \mathcal{X}_N \mid \sigma_X(a, b) \geq r\}$.

Lemma 5.1. *There is some k_0, depending only on h and $\xi(\Sigma)$, such that if c, d and c', d' are respectively ith and jth sides of Q, and if $X \in \mathcal{A}(\phi c, \phi d; k_0)$ and $Y \in \mathcal{A}(\phi c', \phi d'; k_0)$, then either $i = j$ or $X \wedge Y$.*

Proof. This is just Lemma 10.4 of [Bo8] applied to the particular case of $\mathbb{P}(\Sigma)$. □

We will fix $h = h_0$ (as determined below), and hence k_0, depending only on $\xi(\Sigma)$. We then abbreviate $\mathcal{A}(a, b) = \mathcal{A}(a, b; k_0)$.

Definition. We say that $\phi : Q \longrightarrow \mathbb{P}(\Sigma)$ is *non-degenerate* if $\mathcal{A}(\phi c, \phi d) \neq \varnothing$ for every side, c, d, of Q.

We also want to discuss cubes in $\mathbb{P}^*(\Sigma)$.

Suppose that Q is a ξ_0-cube in $\mathbb{P}^*$. Recall that we have maps $\theta^*_X : \mathbb{P}^* \longrightarrow \mathbb{G}^*(X)$ and $\psi^*_X : \mathbb{P}^* \longrightarrow \mathbb{P}^*(X)$ for $X \in \mathcal{U}X_N$. Given $i \in \{1, \dots, \xi_0\}$, let a, b be an ith side of Q. We write

$$A_i(Q) = \{X \in \mathcal{UX}_N \mid \theta^*_X a \neq \theta^*_X b\},$$
$$B_i(Q) = \{X \in \mathcal{UX}_N \mid \psi^*_X a \neq \psi^*_X b\},$$
$$C_i(Q) = \{X \in \mathcal{UX}_N \mid \theta^*_X|[a, b] \text{ is injective}\},$$
$$D_i(Q) = \{X \in \mathcal{UX}_N \mid \psi^*_X|[a, b] \text{ is injective}\}.$$

Note that any two such sides are parallel, so the above are well defined independently of the choice of a, b. Since θ^*_X factors through ψ^*_X (via χ^*_X), we clearly have $C_i(Q) \subseteq A_i(Q) \subseteq B_i(Q)$, and $C_i(Q) \subseteq D_i(Q) \subseteq B_i(Q)$. We also note that any complexity-1 surface in $D_i(Q)$ lies in $C_i(Q)$.

Definition. Given two ξ_0-cubes $Q, Q' \subseteq \mathbb{P}^*(\Sigma)$, we say that Q' is *smaller than* Q (or that Q is *bigger than* Q') if $Q' \subseteq \text{hull}(Q)$.

Note that, in this case, we can label the sides of Q and Q' so that any ith side of Q' is parallel to a pair of points lying in the interval given by any ith side of Q. We also have $B_i(Q') \subseteq B_i(Q)$, $C_i(Q') \supseteq C_i(Q)$ and $D_i(Q') \supseteq D_i(Q)$.

Note that by Lemma 6.7 of [Bo8], there is some h_0, depending only on $\xi(\Sigma)$, such that for any cube, Q, there are h_0-quasimorphisms, $\phi_\zeta : Q \longrightarrow \mathbb{P}(\Sigma)$, such that $\phi_\zeta x \to x$ for all $x \in Q$ (in the sense that the $\mathcal{Z}$-sequence $(\phi_\zeta x)_\zeta$ corresponds to the point x). Note that ϕ_ζ will be non-degenerate for almost all ζ.

The following is now an immediate consequence of Lemma 5.1:

Lemma 5.2. *If $X \in A_i(Q)$ and $Y \in A_j(Q)$, then either $i = j$, or $X \wedge Y$.*

Proof. Suppose $i \neq j$. Let $\phi_\zeta : Q \longrightarrow \mathbb{P}(\Sigma)$ be as described above. In the notation of Lemma 5.1, we have $X_\zeta \in \mathcal{A}(\phi c, \phi d; k_0)$ and $Y_\zeta \in \mathcal{A}(\phi c', \phi d'; k_0)$ for almost all ζ, and so $X_\zeta \wedge Y_\zeta$. □

In what follows, we will split the discussion into "odd" and "even" cases. The odd case is somewhat simpler to describe, and we will deal mainly with that case first.

5.2. **The odd case.**

For the moment, we will assume that Σ is odd, unless otherwise stated. In statements of lemmas (in particular Lemmas 5.7 to 5.11) 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.

Lemma 5.3. *Suppose that Σ is odd, and that $\phi : Q \longrightarrow \mathbb{P}(\Sigma)$, is a non-degenerate quasimorphism. Then there is a good multicurve, τ, and for each $i \in \{1, \dots, \xi_0\}$, there is some $Y_i \in \mathcal{X}_N(\tau)$ such that if c, d is any ith side of Q, then $\mathcal{A}(\phi c, \phi d) = \{Y_i\}$.*

Proof. By the definition of non-degeneracy, we can find, for each i, some $Y_i \in \mathcal{A}(\phi c, \phi d)$, where c, d is an ith side of Q. By Lemma 5.1, the Y_i are pairwise disjoint. Since they all have complexity at least 1, they all have complexity exactly 1, and we see that there is a good multicurve, τ, with $\mathcal{X}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$. It now follows by Lemma 5.1 again that if c', d' is any ith side of Q, then $\mathcal{A}(\phi c', \phi d') \subseteq \{Y_i\}$, so $\mathcal{A}(\phi c', \phi d') = \{Y_i\}$. □

Note that by the distance formula (given as Theorem 2.3 here), we see that if c, d is an ith side of Q, then $\rho(\phi c, \phi d)$ agrees to within linear bounds with $\sigma_{Y_i}(\phi c, \phi d)$ (since the only contribution to the formula comes from Y_i).

Given $i \in \{1, \dots, \xi_0\}$, choose any ith side, c, d, of Q, and set $\mathcal{A}(\phi, i) = \mathcal{A}(\phi c, \phi d)$. (The choice of side is not important.)

We revert to the terminology of "curves" and "standard curves" etc. as at the end of Section 4.

Lemma 5.4. *Suppose that Σ is odd, and that $Q \subseteq \mathbb{P}^*(\Sigma)$ 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 $Y_i \in \mathcal{UX}_N(\tau)$ such that $A_i(Q) = C_i(Q) = \{Y_i\}$.*

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

Proof. By Lemma 6.7 of [Bo8], we have a $\mathcal{Z}$-sequence of h_0-quasimorphisms, $\phi_\zeta : Q \longrightarrow \mathbb{P}(\Sigma)$, with $\phi_\zeta x \to x$ for all $x \in Q$. By Lemma 5.3 here, 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 $\tau_\zeta \to \tau$, $Y_{i,\zeta} \to Y_i$ and with $\mathcal{UX}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$. As observed after Lemma 5.3, if c, d is an ith side of Q, then $\rho_\zeta(\phi_\zeta c, \phi_\zeta d)$ agrees with $\sigma_{Y_{i,\zeta}}(\phi_\zeta c, \phi_\zeta d)$ to within linear bounds. Passing to the limit, we see that since $\rho^*(c, d) \neq 0$, we have $\sigma^*_{Y_i}(c, d) \neq 0$, i.e. $\theta^*_{Y_i} c \neq \theta^*_{Y_i} d$, so $Y_i \in A_i(Q)$. It now follows 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. □

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

Recall that if $\tau \subseteq \mathcal{U}\mathbb{G}^0(\Sigma)$ is a good multicurve, we have associated a convex subset $T^*(\tau) \subseteq \mathbb{P}^*(\Sigma)$.

Lemma 5.5. *Suppose that Σ is odd. Suppose that $\tau \subseteq \mathcal{U}\mathbb{G}^0(\Sigma)$ is a good multicurve, and that $Q \subseteq T^*(\tau)$ is a ξ_0-cube. Then $\tau(Q) = \tau$.*

Proof. Note that $\mathrm{hull}(Q) \subseteq T^*(\tau)$, and we can identify this set with $\prod_i \mathbb{P}^*(Y_i)$ via υ^*_τ, where $Y_1, \dots, Y_{\xi_0}$ are the complementary components of τ. Thus, $\mathrm{hull}(Q) \equiv \prod_i I_i$, where each I_i is a non-trivial interval in $\mathbb{P}^*(Y_i)$. In particular, we see that

$\theta^*_{Y_i}|I_i$ is injective, and so $Y_i \in C_i(Q)$. It follows by Lemma 5.4 that $C_i(Q) = \{Y_i\}$, and so $\tau(Q) = \tau$. □

Note that it follows for an odd surface that if τ and τ' are good multicurves and $T^*(\tau) \cap T^*(\tau')$ contains a ξ_0-cube, then $\tau = \tau'$. (In particular if $T^\infty(\tau) = T^\infty(\tau') \neq \varnothing$, then $\tau = \tau'$.)

Lemma 5.6. *If Σ is odd, and $Q \subseteq \mathbb{P}^*(\Sigma)$ is a ξ_0-cube, then $Q \subseteq T^*(\tau(Q))$.*

Proof. The proof follows exactly as with Lemma 12.6 of [Bo8], given that $\omega^*_{\tau(Q)} : \mathrm{hull}(Q) \longrightarrow T^*(\tau)$ is injective. □

We now restrict to $\mathbb{P}^\infty(\Sigma) \subseteq \mathbb{P}^*(\Sigma)$.

We recall the following definition from Section 4 of [Bo8].

Definition. A *tree-product*, T, in $\mathbb{P}^\infty(\Sigma)$ is a convex subset which is median isomorphic to a direct product of ξ_0 non-trivial rank-1 median algebras. It is *maximal* if it is not contained in any strictly larger tree-product.

Note that T is bilipschitz equivalent to a direct product of $\mathbb{R}$-trees. (This follows since $\mathbb{P}^\infty(\Sigma)$ is bilipschitz equivalent to a median metric. In the induced metric, T is isometric to a direct product of $\mathbb{R}$-trees in the l^1-metric.)

The closure of a tree-product is also a tree-product. (One way to see this is to note that the completion of an l^1 product of metric spaces is the l^1 product of the completions of the factors. Moreover, it is well known that the completion of an $\mathbb{R}$-tree is an $\mathbb{R}$-tree.) Therefore, any maximal tree-product is closed.

Note that in the above terminology, if $\xi_0 \geq 2$, then any closed subset of $\mathbb{P}^\infty(\Sigma)$ homeomorphic to a direct product of ξ_0 furry $\mathbb{R}$-trees is a tree-product (by Proposition 4.8).

The following few statements (5.7 to 5.11) will be valid for both odd and even surfaces. We restrict proofs for the moment to the odd case.

Lemma 5.7. *Suppose that $T \subseteq \mathbb{P}^\infty(\Sigma)$ is a tree-product. Then there is a good multicurve $\tau \subseteq \mathcal{UG}^0(\Sigma)$ such that $T \subseteq T^\infty(\tau)$.*

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 5.6, $x \in Q' \subseteq T^*(\tau) \cap \mathbb{P}^\infty(\Sigma) = T^\infty(\tau)$. Thus $T \subseteq T^\infty(\tau)$. □

Note that (by Lemma 5.5), if Σ is odd, then τ is unique.

Corollary 5.8. *If τ is a good multicurve, then $T^\infty(\tau)$ is either empty or a maximal tree-product.*

Proof. (For Σ odd.) Suppose $T^\infty(\tau) \neq \varnothing$. By Lemma 4.7, $T^\infty(\tau)$ is a tree-product. Let $T \supseteq T^\infty(\tau)$ be a larger tree-product. By Lemma 5.7, $T \subseteq T^\infty(\tau')$ for some τ'. By the remark after Lemma 5.5, it follows that $\tau = \tau'$. Therefore $T = T^\infty(\tau)$. □

Putting the above together, we see that a (closed) subset of $\mathbb{P}^\infty(\Sigma)$ is a maximal tree-product if and only if it has the form $T^\infty(\tau)$ for some good multicurve τ such that $T^\infty(\tau) \neq \varnothing$. In particular, each factor must be a furry tree. Moreover, by Proposition 4.8, a closed subset of $\mathbb{P}^\infty(\Sigma)$, homeomorphic to a product of ξ_0 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 $\mathbb{P}^\infty(\Sigma)$). We deduce (in the odd case):

Lemma 5.9. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi_0(\Sigma) = \xi_0(\Sigma')$ and that $f : \mathbb{P}^\infty(\Sigma) \longrightarrow \mathbb{P}^\infty(\Sigma')$ 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 $f(T^\infty(\tau)) = T^\infty(\tau')$.*

To recover information back in $\mathbb{P}(\Sigma)$, we use the results of Section 14 of [Bo8]. The general argument follows almost exactly as with the proof of Lemma 14.5 of that paper (where a similar statement was proven for flats in the marking complex).

Lemma 5.10. *Suppose that Σ and Σ' are compact orientable surfaces with $\xi_0(\Sigma) = \xi_0(\Sigma') \geq 2$. Suppose that $\phi : \mathbb{P}(\Sigma) \longrightarrow \mathbb{P}(\Sigma')$ is a quasi-isometry. Then, given any good multicurve, τ, in Σ, there is a good multicurve τ' in Σ' such that $\mathrm{hd}(T(\tau'), \phi T(\tau))$ is bounded above by some constant depending only on $\xi(\Sigma)$ and the parameters of ϕ.*

Proof. (If Σ and Σ' are both odd.) To simplify the account, we first show that there is such a bound for a given map ϕ. To this end, $\mathcal{F}$ be the set of subsets of the form $T(\tau)$ as τ ranges over all good multicurves, and let $\mathcal{E} = \{\phi(F) \mid F \in \mathcal{F}\}$. We need to verify that these collections satisfy the conditions (S1)–(S3) laid out in Section 14 of [Bo8].

Now (S1) simply says that each $E \in \mathcal{E}$ is coarsely connected (that is, some uniform neighbourhood is connected). But this is clear, since it is the quasi-isometric image of a connected graph.

By Lemma 14.4 of [Bo8], Property (S2) holds provided the family $\mathcal{F}$ is "linearly divergent". In our case, linear divergence is precisely the conclusion of Lemma 4.5 here.

As in Section 14 of [Bo8], Property (S3) is an immediate consequence of the fact that if $\tau \subseteq \mathcal{UG}^0$ is a good multicurve, then there is some good multicurve, $\tau' \subseteq \mathcal{UG}^0$, with $\phi^\infty(T^\infty(\tau)) \subseteq T^\infty(\tau')$. (See the discussion before Lemma 14.5 of [Bo8].)

We can thus apply Lemma 14.3 of [Bo8]. This tells us that for each $E \in \mathcal{E}$, there is some $F \in \mathcal{F}$ with $\mathrm{hd}(E, F)$ bounded. In other words, this shows that the conclusion holds for a particular quasi-isometry ϕ.

To show that k only depends on $\xi(\Sigma)$ and the quasi-isometry parameters, we consider simultaneously all quasi-isometries, ϕ, of $\mathbb{P}(\Sigma)$ to $\mathbb{P}(\Sigma')$ with given parameter bounds, and let $\mathcal{E}$ be the set of all images $\phi(F)$ now allowing ϕ also to

vary among such maps. There is no change to (S1) or (S2), and property (S3) still holds for the same reason. This now gives a uniform bound as required. □

Recall that $T(\tau)$ is, up to bounded distance, the image of a quasi-isometric embeddings, $\upsilon_\tau : \mathcal{T}(\tau) \longrightarrow \mathbb{P}(\Sigma)$, where $\mathcal{T}(\tau) = \prod_{X \in \mathcal{X}_N(\tau)} \mathbb{P}(X)$. We can elaborate on Lemma 5.10 to conclude that the factors of $T(\tau)$ and $T(\tau')$ are also coarsely preserved:

Proposition 5.11. *Suppose that Σ, Σ', ϕ are as in Lemma 5.10. Then, given any good multicurve, τ, in Σ, there is a good multicurve τ' in Σ', a bijection, $\pi : \mathcal{X}_N(\tau) \longrightarrow \mathcal{X}_N(\tau')$, and a quasi-isometry, $\phi_X : \mathbb{P}(X) \longrightarrow \mathbb{P}(\pi(X))$ for each $X \in \mathcal{X}_N(\tau)$, such that the maps $\upsilon_\tau \circ (\prod_X \phi_X)$ and $\phi \circ \upsilon_{\tau'} : \mathcal{T}(\tau) \longrightarrow \mathbb{P}(\Sigma')$ agree up to bounded distance. The bound, and the parameters of the maps ϕ_X, depend only on $\xi(\Sigma)$ and the parameters of ϕ.*

Here, of course, $\mathcal{X}_N(\tau')$ denotes a collection of subsurfaces of Σ'.

Proof. (If Σ and Σ' are both odd.) By Lemma 5.10, we see that $\phi | T(\tau)$ is a bounded distance from a quasi-isometry from $T(\tau)$ to $T(\tau')$. Thus, via the quasi-isometric embedding υ_τ and $\upsilon_{\tau'}$, we get a quasi-isometry $\hat{\phi} : \mathcal{T}(\tau) \longrightarrow \mathcal{T}(\tau')$. Now each of the factors of $\mathcal{T}(\tau)$ and $\mathcal{T}(\tau')$ is a bushy hyperbolic space (in this case, a quasitree). 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 there one would need to observe, in addition, that each of the factors admits a cobounded isometric action. In any case, bushy is all that is really required for their argument. See [Bo5] for further discussion of this.) □

5.3. **The even case.**

We move on to consider the even case.

Recall that in this case, a "good multicurve", τ, is of one of the types (T1) or (T2) described in Section 4. In case (T2), we denote the complexity-2 component of $\Sigma \setminus \tau$ by $W(\tau)$.

We will also need to refer to two other classes of multicurve:

(T1′a): $\Sigma \setminus \tau$ has exactly one $S_{0,3}$ component, and exactly ξ_0 other components, all of complexity 1.
(T1′b): $\Sigma \setminus \tau$ has exactly ξ_0 components, all of complexity 1, one of which is an $S_{0,4}$ with two boundary components identified to a single component of τ (so as to give an $S_{1,2}$).

We write (T1′) for the union of cases (T1′a) and (T1′b). Again, $|\mathcal{X}_N(\tau)| = \xi_0$ in this case. Note that we can always remove a component of a type (T1′) multicurve so as to give a good multicurve of type (T1).

(Expanding on Lemma 4.3, it's not hard to see that any multicurve τ, with $|\mathcal{X}_N(\tau)| = \xi_0$ is of exactly one of the types, (T0), (T1), (T1′) or (T2).)

Again, we begin by describing quasicubes. Suppose that Q is a ξ_0-cube, and that $\phi : Q \longrightarrow \mathbb{P}(\Sigma)$ is a non-degenerate h_0-quasimorphism. Recall that, given $i \in \{1, \dots, \xi_0\}$, we choose any ith side, c, d, of Q, and set $\mathcal{A}(\phi, i) = \mathcal{A}(\phi c, \phi d)$.

We suppose that Σ is even, and Q is a ξ_0-cube.

Lemma 5.12. *Suppose that $\phi : Q \longrightarrow \mathbb{P}(\Sigma)$ is a non-degenerate quasimorphism. Then there is a uniquely determined multicurve, $\tau = \tau(\phi)$, such that we can write $\mathcal{X}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$, so that exactly one of the following hold.*
(U1): τ is of type (T1), and for all $i \in \{1, \dots, \xi_0\} \setminus \{i_0\}$, $\mathcal{A}(\phi, i) = \{Y_i\}$, where i_0 is the index for which $Y_{i_0} = W(\tau)$. Moreover, the elements of $\mathcal{A}(\phi, i_0)$ together fill $W(\tau)$.
(U1′): τ is of type (T1′), and $\mathcal{A}(\phi, i) = \{Y_i\}$ for all $i \in \{1, \dots, \xi_0\}$.
(U2): τ is of type (T2), and $\mathcal{A}(\phi, i) = \{Y_i\}$ for all $i \in \{1, \dots, \xi_0\}$.

Proof. By non-degeneracy, $\mathcal{A}(\phi, i) \neq \varnothing$ for all i. Also, by Lemma 5.2, all elements of $\mathcal{A}(\phi, i)$ are disjoint from all elements of $\mathcal{A}(\phi, j)$ whenever $i \neq j$. Since all these surfaces have complexity at least 1, there can be at most one index, say i_0, for which $\mathcal{A}(\phi, i_0)$ does not consist of a single complexity-1 surface. If such an index exists, the surfaces in $\mathcal{A}(\phi, i_0)$ fill some complexity-2 surface, Y_{i_0}. Write Y_i for the unique element of $\mathcal{A}(\phi, i) = \{Y_i\}$ when $i \neq i_0$. Let τ be the union of all relative boundary components of all the Y_i in Σ. Then, $\mathcal{X}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$, and we see we are in the situation of (U1).

If there is no such index, i_0, then we can write $\mathcal{A}(\phi, i) = \{Y_i\}$ for all $i \in \{1, \dots, \xi_0\}$, where each Y_i has complexity-1. We again write τ for the union of all relative boundary curves, and we see that in this case, we are either in case (U1′) or case (U2). □

Note that (as in the odd case) if c, d is an ith side of Q with $\mathcal{A}(\phi, i) = Y_i$, then $\rho(\phi c, \phi d)$ agrees up to linear bounds with $\sigma_{Y_i}(\phi c, \phi d)$. In the case where $i = i_0$ in (U1), we get instead that $\rho(\phi c, \phi d)$ agrees to within linear bounds with $\rho_{W(\tau)}(\phi c, \phi d)$. This follows by a similar argument. All contributions to the distance formula (given as Theorem 2.3 here) in $\mathbb{P}(\Sigma)$ come from subsurfaces of X, and so agree up to linear bounds, with the same formula in $\mathbb{P}(W(\tau))$.

We now again pass to $\mathbb{P}^*(\Sigma)$. We need the following general facts.

Lemma 5.13. *Let $X, Y \in \mathcal{X}_N$. Suppose $a, b, c, d \in \mathbb{P}^*(\Sigma)$, 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 $\theta^*_Y b \neq \theta^*_Y c$. Then either $X = Y$ or $X \wedge Y$.*

Proof. This is an immediate consequence of Lemma 11.2 of [Bo8]. (Here $\mathbb{P}(\Sigma)$ satisfies all the hypotheses of the space $\mathcal{M}$ there. The assumptions on a, b, c, d, X, Y are identical.) □

Corollary 5.14. *Suppose $Q \subset \mathbb{P}^*(\Sigma)$ 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 $X \wedge Y$.*

Proof. Let a, d be an ith side of Q. By the definition of $A_i(Q)$, $\theta^*_Y a \neq \theta^*_Y d$. Since $\theta^*_X | [a, d]$ is injective, $\theta^*_X a, \theta^*_X b, \theta^*_X c, \theta^*_X d$ are all distinct. By continuity of θ^*_Y, we can find points $b, c \in [a, d] \setminus \{a, d\}$ with $\theta^*_Y b \neq \theta^*_Y c$. We now apply Lemma 5.13. □

Lemma 5.15. *Suppose that $Q \subset \mathbb{P}^*(\Sigma)$ is a ξ_0-cube. Then there is a uniquely determined multicurve, $\tau = \tau(Q)$, such that we can write $\mathcal{UX}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$, so that exactly one of the following hold.*
(V1): τ is of type (T1), and for all $i \in \{1, \dots, \xi_0\} \setminus \{i_0\}$, $A_i(Q) = C_i(Q) = \{Y_i\}$, where i_0 is the index for which $Y_{i_0} = W(\tau)$. Moreover, $C_{i_0}(Q) \subset \{W(\tau)\}$ and $D_{i_0}(Q) \cap \mathcal{UX}_N(\tau) = \{W(\tau)\}$.
(V1′): τ is of type (T1′), and $A_i(Q) = C_i(Q) = \{Y_i\}$ for all $i \in \{1, \dots, \xi_0\}$.
(V2): τ is of type (T2), and $A_i(Q) = C_i(Q) = \{Y_i\}$ for all $i \in \{1, \dots, \xi_0\}$.

Proof. We proceed similarly as with the proof Lemma 5.4. Let $\phi_\zeta : Q \longrightarrow \mathbb{P}(\Sigma)$ be a sequence of non-degenerate h_0-quasimorphisms converging on $Q \subset \mathbb{P}^*(\Sigma)$. We can assume that the type of $\tau(\phi_\zeta)$ is constant. By Lemma 5.12, this means that each ϕ_ζ is exactly as described by (U1), (U2) or (U1′), according to this type.

Suppose first that the type of each $\tau(\phi_\zeta)$ is (T2). Write $\mathcal{X}_N(\tau) = \{Y_{1,\zeta}, \dots, Y_{\xi_0,\zeta}\}$. Let $\tau \subset \mathcal{X}_N(\tau)$ be the limiting multicurve, and let Y_i be the limit of $(Y_{i,\zeta})_\zeta$. Thus, $\mathcal{UX}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$. As with Lemma 5.4, we see that $A_i(Q) = C_i(Q) = \{Y_i\}$, and we are in case (V2).

If the type is (T1′), we similarly get the a limiting τ, with $\mathcal{X}_N(\tau) = \{Y_1, \dots, Y_{\xi_0}\}$, and with $Y_i \in C_i(Q)$. Now, if $Y \in A_i(Q) \setminus \{Y_i\}$, then by Lemma 5.2, it must be disjoint from each Y_j, with $j \neq i$. Therefore $Y \pitchfork Y_i$ giving a contradiction to Corollary 5.14. We conclude that $A_i(Q) = C_i(Q) = \{Y_i\}$.

Now suppose that the type is (T1). Let τ_0 be the limiting multicurve. We can write $\mathcal{X}_N(\tau_0) = \{Y_1, \dots, Y_{\xi_0}\}$, where Y_i is the limit of $(Y_{i,\zeta})_\zeta$. Let c, d be an i_0th side of Q. Now, $\rho(\phi_\zeta c, \phi_\zeta d)$ agrees to within linear bounds with $\rho_{W(\tau_0)}(\phi_\zeta c, \phi_\zeta d)$, and so it follows that $\psi^*_{W(\tau_0)} c \neq \psi^*_{W(\tau_0)} d$.

In fact, we claim that $\psi^*_{W(\tau_0)} | [c, d]$ is injective. To see this, suppose that $e, f \in [c, d]$, with $e \neq f$. Let $Q' \subset \mathbb{P}^*(\Sigma)$ be the cube with i_0th side parallel to e, f, and all other sides parallel to those of Q. Thus, Q' is a smaller cube. Now $A_i(Q') = C_i(Q') = A_i(Q) = C_i(Q) = \{Y_i\}$ for all $i \neq i_0$. By Lemma 5.2 applied to Q', we see that all elements of $A_{i_0}(Q')$ are subsurfaces of $W(\tau_0)$. Applying what we already know to Q' in place of Q, we see we must either be in the situation of the previous paragraph, so that there is some complexity-1 subsurface of $W(\tau_0)$ lying in $C_{i_0}(Q)$; or else we are in the situation described in the present paragraph, so that $\psi^*_{W(\tau_0)} e \neq \psi^*_{W(\tau_0)} f$. Either way, we must have $\psi^*_{W(\tau_0)} e \neq \psi^*_{W(\tau_0)} f$. This proves the claim. In other words, we have shown that $W(\tau_0) \in D_i(Q)$.

As observed above, each element of $C_{i_0}(Q)$ is a subsurface of $W(\tau_0)$. Suppose that $C_{i_0}(Q)$ does not contain any complexity-1 surface. Then $C_{i_0}(Q) \subseteq \{W(\tau_0)\}$. Moreover, since any complexity-1 surface in $D_{i_0}(Q)$ would also lie in $C_{i_0}(Q)$, we get that $D_{i_0}(Q) = \{W(\tau_0)\}$. Setting $\tau = \tau_0$, and $Y_{i_0} = W(\tau)$, we see that we are in case (V1).

Suppose instead that there is some complexity-1 subsurface, $Y \in C_{i_0}(Q)$. Applying Lemma 5.2 and Corollary 5.14 (similarly as for the type (T1′) case above) we see that $C_{i_0}(Q) = \{Y\}$. If $W(\tau_0)$ is an $S_{0,5}$, then Y is an $S_{0,4}$. If $W(\tau_0)$ is an $S_{1,2}$, then Y is either an $S_{1,1}$, or else an $S_{0,4}$ with two boundary curves identified. In all cases, the relative boundary of Y in $W(\tau_0)$ is a single curve, β. Let $\tau = \tau_0 \cup \beta$. Set $Y_{i_0} = Y \in \mathcal{UX}_N(\tau)$. We are now back in case (V1′). □

Definition. We will say that a ξ_0-cube, $Q \subseteq \mathbb{P}^*(\Sigma)$, is of *type (1)* if $\tau(Q)$ is of type (T1) or of type (T1′) (i.e. Q is of type (V1) or (V1′) described by Lemma 5.15). We will say that it is of *type (2)* if $\tau(Q)$ is of type (T2).

Lemma 5.16. *Suppose that Σ is even and that $Q \subseteq \mathbb{P}^*(\Sigma)$ is a ξ_0-cube. Then $Q \subseteq T^*(\tau(Q))$.*

Proof. This now follows exactly as with Lemma 5.6. □

Suppose that $Q' \subseteq Q$ is a smaller cube. As noted earlier, $C_i(Q) \subseteq C_i(Q')$. In fact, we now see that the only way that these can differ is if $\tau(Q')$ is of type (T1′) and $\tau(Q)$ is of type (T1), obtained by deleting a single curve. It follows that $\tau(Q') = \tau(Q)$ except in the above situation. In particular, Q, Q' are either both of type (1), or else they are both of type (2).

We now restrict attention to $\mathbb{P}^\infty(\Sigma) \subseteq \mathbb{P}^*(\Sigma)$.

Suppose that $T \subseteq \mathbb{P}^\infty(\Sigma)$ 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 $\theta^\infty_X|\Delta_i : \Delta_i \longrightarrow \mathbb{G}^\infty(X)$ and $\psi^\infty_X|\Delta_i : \Delta_i \longrightarrow \mathbb{P}^\infty(X)$.

Lemma 5.17. *Suppose $T \subseteq \mathbb{P}^\infty(\Sigma)$ is a tree-product. Then either all ξ_0-cubes in T are of type (1) or else they are all of type (2).*

Proof. Suppose $Q_0, Q_1 \subseteq T$ are ξ_0-cubes. Then there are ξ_0-cubes Q'_0, Q'_1, Q'', with Q_0, Q'' bigger than Q'_0, and with Q_1, Q'' bigger than Q'_1. (It is readily checked that, given any two non-trivial intervals, I_0, I_1, in an $\mathbb{R}$-tree, there are non-trivial intervals, $I'_0 \subseteq I_0$, $I'_1 \subseteq I_1$, such that $I'_0 \cup I'_1$ is contained in an interval I''. One then performs this construction separately on each factor of T.) It follows that each of these cubes has the same type. □

We refer to a tree-product as *type (1)* or *type (2)* accordingly.

The type (2) case is simpler, and similar to the situation for an odd surface. In particular, following the same argument as with Lemma 5.5, we see that if τ is of type (T2), and $Q \subseteq T^*(\tau)$, then $\tau(Q) = \tau$. Moreover, similarly as with Lemma 5.7, we see that if $T \subseteq \mathbb{P}^\infty(\Sigma)$ is a type (2) tree-product, then there is a type

(T2) multicurve, τ, with $T \subseteq T^\infty(\tau)$. As with Corollary 5.8 in the odd case, we conclude:

Lemma 5.18. *If τ is a type (T2) multicurve, then $T^\infty(\tau)$ is either empty or a maximal tree-product. Moreover, every maximal tree-product of type (2) has this form.*

The type (1) case is a bit more complicated. We have:

Lemma 5.19. *Suppose that Σ is even, and that $T \subseteq \mathbb{P}^\infty(\Sigma)$ is a type (1) tree-product. Then there is a type (T1) 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 $\theta^\infty_{Y_i}|\Delta_i$ is injective, and such that $\psi^\infty_{W(\tau)}|\Delta_{i_0}$ is injective.*

(Note that we are allowing the possibility that there may be some (unique) complexity-1 surface $Y \prec W(\tau)$ with $\theta^\infty_Y|\Delta_{i_0}$ injective.)

Proof. Suppose first that all of the ξ_0-cubes in T are of type (V1′) (as in Lemma 5.17). If Q_0, Q are ξ_0-cubes in T with Q bigger than Q_0, then we see that $\tau(Q) = \tau(Q_0)$, as in Lemma 5.18. It follows that $\tau(Q) = \tau_0$, say, 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 $\mathcal{UX}_N(\tau)$ consistently as $Y_1, \dots, Y_{\xi_0}$, so that Y_i the ith 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 ith side $\{a, b\}$, say. Since $Y_i \in C_i(Q)$, we get $\theta^\infty_{Y_i} a \neq \theta^\infty_{Y_i} b$. In other words, $\theta^\infty_{Y_i}|\Delta_i$ is injective. We can now (somewhat artificially) remove a component of τ_0 to give a type (T1) multicurve, τ. If Y_{i_0} is the element lying in $W(\tau)$, then since $\theta^\infty_{Y_{i_0}}|\Delta_{i_0}$ is injective, so is $\psi^\infty_{W(\tau)}|\Delta_{i_0}$.

We can therefore assume that there is a type (V1) cube, $Q \subseteq T$. Let $\tau = \tau(Q)$. We claim that if $i \neq i_0$, then $\theta^\infty_{Y_i}|\Delta_i$ is injective. For suppose $a, b \in \Delta_i$ are distinct. Let Q' be the ξ_0-cube with ith 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 (T1) with $\tau(Q') = \tau$, and so $C_i(Q') = \{Y_i\}$. In particular, $\theta^\infty_{Y_i} a \neq \theta^\infty_{Y_i} b$ as claimed. Finally, we claim that $\psi^\infty_{W(\tau)}|\Delta_{i_0}$ is injective. For suppose that $c, d \in \Delta_{i_0}$ are distinct. Let Q'' be any ξ_0-cube with i_0th side parallel to $\{c, d\}$. If $i \neq i_0$, then (since $\theta^\infty_{Y_i}|\Delta_i$ is injective) we have $Y_i \in C_i(Q'')$. It follows that Q'' is either of type (U1) with $\tau(Q'') = \tau$, or type (U1′) with $\tau(Q'') \supseteq \tau$. Either way we get $\psi^\infty_{W(\tau)} c \neq \psi^\infty_{W(\tau)} d$ as required. □

It follows that any cube $Q \subseteq T$ has $\tau(Q) \supseteq \tau$. Thus, by Lemma 5.16, we have $Q \subseteq T^\infty(\tau(Q)) \subseteq T^\infty(\tau)$. Since any point of T lies in such a cube, we have $T \subseteq T^\infty(\tau)$. We deduce:

Lemma 5.20. *If τ is a type (T1) multicurve, then $T^\infty(\tau)$ is either empty or a tree-product. Moreover, every maximal tree-product of type (1) has this form.*

Since good multicurves are precisely those of type (T1) or (T2), we deduce Corollary 5.8 in the even case.

We can now proceed to prove Lemmas 5.9 and 5.10 and Proposition 5.11 in the even case, hence in general.

5.4. **Application to the quasi-isometric rigidity of the pants graphs.**

We aim to prove Theorem 1.6, and derive some other consequences to be used later.

We begin by elaborating on Lemma 5.20. Note that if τ is a type (T1) multicurve, we can always find another type (T1) multicurve, τ', such that $\tau \cup \tau'$ is a type (T1′) multicurve. (To see this, let $\beta \in \tau$ be a boundary component of $W(\tau)$, and let $\gamma \prec W(\tau)$ be any curve which cuts off an $S_{0,3}$ with β as another boundary curve. We set $\tau' = (\tau \setminus \beta) \cup \gamma$. Note that this uses our assumption that $\xi(\Sigma) \geq 3$, hence at least 4 in the even case.) In particular, $T^\infty(\tau \cup \tau') \subseteq T^\infty(\tau) \cap T^\infty(\tau')$ is a tree-product. On the other hand, if τ is of type (T2), and τ' is any good multicurve different from τ, then $T^\infty(\tau) \cap T^\infty(\tau')$ cannot contain any ξ_0-cube. Now any topological embedding of a ξ_0-ball into $\mathbb{P}(\Sigma)$ must contain a ξ_0-cube. (This is because any compact subset thereof lies inside a "cubulated" set: see Proposition 4.3 of [Bo8]. The statement there assumed that the image was closed and concluded that the whole ball was cubulated. But as noted at the end of Section 3 of that paper, one does not need to assume closed for the weaker statement given here to hold. A more direct argument when the ball is contained in a tree-product is given in [Bo5].) It follows that $T^\infty(\tau) \cap T^\infty(\tau')$ cannot contain any topological ξ_0-ball. The distinction between type (1) and type (2) is therefore detectable from the topology of $\mathbb{P}^\infty(\Sigma)$.

More formally, in either the odd or even case, we say that a maximal tree-product is *isolated* if its intersection with any other maximal tree-product does not contain any topological ξ_0-ball. We can now complete the proof of Theorem 1.6.

Proof of Theorem 1.6. We see that Σ is odd if and only if all maximal tree-products in $\mathbb{P}^\infty(\Sigma)$ are isolated. The topology of $\mathbb{P}^\infty(\Sigma)$ therefore determines the parity of $\xi(\Sigma)$. Also $\xi_0(\Sigma)$ is determined (for example, as the locally compact dimension of $\mathbb{P}^\infty(\Sigma)$, or as the maximal rank of a tree-product). These facts together determine $\xi(\Sigma)$. □

For the following statement, we can drop the assumption that $\xi(\Sigma) \geq 3$.

Proposition 5.21. *Suppose that Σ, Σ' are compact surfaces with $\mathbb{P}(\Sigma)$ quasi-isometric to $\mathbb{P}(\Sigma')$, 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 $\mathbb{P}(\Sigma)$ is hyperbolic but not a quasitree (see the remark at the end of Section 2). If $\xi(\Sigma) = 1$, then $\mathbb{P}(\Sigma)$ is a quasitree. This deals with all cases. □

Note that this is equivalent to the first part of Theorem 1.3. We will give a stronger statement later, see Theorem 7.8.

We also note that we can refine Lemma 5.10 and Proposition 5.11 so as to distinguish the type of a good multicurve, τ, in terms of the coarse structure of $T(\tau)$ in the quasi-isometry type $\mathbb{P}(\Sigma)$. Thus, in Lemma 5.10, τ' must have the same type as τ. This follows from the above discussion of $\mathbb{P}^\infty(\Sigma)$. (In the proof of Lemma 5.10 in the even case, we can take $\mathcal{F}$ be the set of all tree-products of type (T1), and ignore those of type (T2).) In fact, it also follows more directly, since τ is of type (T0) or (T2) if and only if all factors of $T(\tau)$ are quasitrees.

Lemma 5.22. *In Lemma 5.10 and Proposition 5.11, if τ is of type (T0) or (T1), then τ' is also of type (T0) or (T1).*

Proof. We now know that Σ and Σ' are of the same parity. By definition, (T0) occurs only in the odd case, and (T1) only in the even case. It therefore suffices to distinguish (T1) from (T2) in the even case, which is achieved in the above discussion. □

We will eventually show that Σ and Σ' are equal (Theorem 7.8). Therefore, retrospectively, we will see that the type of τ' is preserved.

6. The arrangement of product regions

In this section, we give an account of how maximal product regions are arranged in $\mathbb{P}(\Sigma)$. This will enable us to identify certain subsurfaces of Σ in terms of the coarse geometry. It reduces the question of quasi-isometric rigidity to a combinatorial question, which we will discuss further in Section 7. The discussion applies to both the odd and even cases, though with a few differences.

Given a multicurve, τ, write $G(\tau) \leq \mathrm{Map}(\Sigma)$ for the subgroup which preserves setwise each component of τ and each component of $\Sigma \setminus \tau$. If $X \in \mathcal{X}$, write $G(X) \leq \mathrm{Map}(\Sigma)$ for the subgroup supported on X. (Note that $G(\tau)$ is a direct product of the groups $G(X)$ as X ranges over $\mathcal{X}_A(\tau) \cup \mathcal{X}_N(\tau)$.)

We recall the notation used in Section 15 of [Bo8] as follows.

We write $\mathcal{B} = \mathcal{B}(\mathbb{P}(\Sigma))$ for the set of subsets of $\mathbb{P}(\Sigma)$ defined up to finite Hausdorff distance. If $A, B \in \mathcal{B}$, 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 $\mathcal{B}$.

Any subgroup, $G \leq \mathrm{Map}(\Sigma)$, determines an element $B(G) \in \mathcal{B}$, namely the class of any orbit of G in $\mathbb{P}(\Sigma)$. We will abbreviate $B(\tau) = B(G(\tau))$ and $B(X) = B(G(X))$. Note that $G(\tau)$ acts coboundedly on $T(\tau)$, and so $B(\tau)$ is just the class of $T(\tau)$. Also, if X is an annulus, then $B(X)$ is just the class of bounded subsets. (Of course, the groups $G(X)$ do not act discretely here, but that does not affect our discussion.)

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 $G(X) \leq G(\tau)$ and so $B(X) \leq B(\tau)$.

Lemma 6.1. *If $X, Y \in \mathcal{X}_N$ with $B(X) \leq B(Y)$, then $X \preceq Y$.*

Proof. Let $a \in \mathbb{P}(\Sigma)$ be any complete multicurve 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 $\beta = \theta_X\alpha \in \mathbb{G}^0(X)$. Let $g \in G(X)$ be any pseudoanosov in X. Then $\sigma_X(\beta, g^n\beta) \to \infty$ as $n \to \infty$. Thus, $\sigma_X(a, g^n a) \to \infty$. 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, $\sigma_X(g^n a, G(Y)a) \to \infty$. But $\rho(g^n a, G(Y)a)$ is linearly bounded below by its projection in $\mathbb{G}(X)$, and so $\rho(g^n a, G(Y)a) \to \infty$, contradicting the assumption that $B(X) \leq B(Y)$. □

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 multicurve, τ, of either type (T0) or type (T1) 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 it is either an $S_{1,1}$ or else it is an $S_{0,4}$ and each component of the complement is odd and meets X in exactly one curve. If Σ is even, then X is admissible if and only if it is one of the following four possibilities: (1) an $S_{1,1}$; (2) an $S_{0,4}$ with all but one of the complementary components odd and all meeting X in a single curve; (3) an $S_{0,4}$ with all complementary components odd, one meeting X in two curves, and each of the others meeting X is a single curve; or (4) an $S_{1,2}$ or $S_{0,5}$ with all complementary components odd and meeting X in a single curve. One can give a similar description of compatibility (though we won't need this here).

Note that maximal compatible sets are in bijective correspondence to type (T0) or (T1) multicurves — they are precisely those of the form $\mathcal{X}_N(\tau)$.

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

Lemma 6.2. *Suppose that Σ, Σ' are compact surfaces with $\xi(\Sigma) = \xi(\Sigma') \geq 3$. Suppose that $\phi : \mathbb{P}(\Sigma) \longrightarrow \mathbb{P}(\Sigma')$ 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 subsurfaces of Σ'. It preserves the complexity of the subsurface. Moreover, a set, $\mathcal{Y}$, of admissible subsurfaces of Σ is compatible, if and only if their images, $\pi\mathcal{Y}$, are compatible in Σ'.*

Proof. Let $X \subseteq \Sigma$ be an admissible subsurface of Σ, i.e. $X \in \mathcal{X}_N(\tau)$ for some multicurve, τ, of type (T0) or (T1). By Lemma 5.10, there is some multicurve, τ', in Σ' with $\mathrm{hd}(T(\tau'), \phi T(\tau)) < \infty$. Moreover, by Proposition 5.11, there is some $X' \in \mathcal{X}_N(\tau')$ so that the factor of $T(\tau')$ corresponding to X' is a finite Hausdorff distance from the ϕ-image of the factor corresponding to X in $T(\tau)$. In other words, $B(X') = \phi(B(X))$. By Lemma 5.22, τ' is also of type (T0) or (T1), and so

X' is an admissible subsurface of Σ'. By Lemma 6.1, X' is uniquely determined, and we set $\pi X = X'$. Also, by Lemma 6.1, π is injective. If X, Y are compatible, then we can take $X, Y \in \mathcal{X}_N(\tau)$, and so $\pi X, \pi Y$ are distinct elements of $\mathcal{X}_N(\tau')$, hence compatible. Clearly this construction is invertible, so π is a bijection, as claimed. Note that an admissible subsurface, X, has complexity 1 if and only if it does not properly contain another admissible surface, and this is detectable, again using Lemma 6.1. (Alternatively, X has complexity 1 if and only if $B(X)$ is a quasitree.) Therefore π preserves complexity. □

Note that, if τ is a (T0) or (T1) multicurve, then by construction, the map π of Lemma 6.2, when restricted to $\mathcal{X}_N(\tau)$, agrees with the map π given by Proposition 5.11.

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 of its 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*. (We elaborate on this terminology in Section 7.) We write $C_1 = C_1(\Sigma) \subseteq \mathbb{G}^0(\Sigma)$ for the set of 1-separating curves.

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 6.3. *With the hypotheses of Lemma 6.2, X is terminal in Σ if and only if πX is terminal in Σ'.*

Proof. We show that 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 an $S_{0,4}$ 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 is 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. This gives a contradiction.

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

Therefore, π restricts to a bijection from $C_1(\Sigma)$ to $C_1(\Sigma')$. Let $\mathbb{G}_1(\Sigma)$ be the full subgraph of the curve graph $\mathbb{G}(\Sigma)$ with vertex set $C_1(\Sigma)$.

Lemma 6.4. *The map π gives rise to an isomorphism from $\mathbb{G}_1(\Sigma)$ to $\mathbb{G}_1(\Sigma')$.*

Proof. In other words, π preserves adjacency. Note that two curves in $\mathbb{G}_1(\Sigma)$ are adjacent if and only if the complexity-1 surfaces, X, Y, which they bound are disjoint. By the observation prior to Lemma 6.3, this is equivalent to saying that X, Y are compatible. By Lemma 6.2 this is equivalent to saying that πX and πY are compatible. This is in turn equivalent to saying that the relative boundaries of πX and πY are adjacent in $\mathbb{G}_1(\Sigma')$. □

In particular, this shows that if $\mathbb{P}(\Sigma)$ and $\mathbb{P}(\Sigma')$ are quasi-isometric, then $\mathbb{G}_1(\Sigma)$ and $\mathbb{G}_1(\Sigma')$ are isomorphic.

We will elaborate on this in Section 7, and explain how, for most surfaces at least, it can be used to deduce quasi-isometric rigidity of the pants graph.

7. Variations on the curve graph

Most of the content of this section is combinatorial. We will reduce the combinatorial rigidity of the graph, $\mathbb{G}_1(\Sigma)$, introduced at the end of Section 6, to that of the "strongly separating curve graph" (here denoted $\mathbb{G}_{1+}(\Sigma)$) as investigated in [Bo6]. The application to quasi-isometric rigidity of the pants graphs will be discussed at the end (see Theorems 7.8 and 7.9).

Let Σ be a compact surface with $\xi(\Sigma) \geq 3$. Recall that $\mathrm{Map}(\Sigma)$ acts cofinitely on the curve graph, $\mathbb{G}(\Sigma)$. Given a subset, $C \subseteq \mathbb{G}^0(\Sigma)$, we write $\mathbb{G}(\Sigma, C)$ for the full subgraph of $\mathbb{G}(\Sigma)$ with vertex set C. Clearly, if C is $\mathrm{Map}(\Sigma)$-invariant, then $\mathrm{Map}(\Sigma)$ acts on $\mathbb{G}(\Sigma, C)$.

Definition. We say that $\mathbb{G}(\Sigma, C)$ is *rigid* if any automorphism of $\mathbb{G}(\Sigma, C)$ is induced by an element of $\mathrm{Map}(\Sigma)$.

It is natural to ask for which $\mathrm{Map}(\Sigma)$-invariant subsets, C, the graph is $\mathbb{G}(\Sigma, C)$ rigid. It was shown in [Iv, Ko, L] that $\mathbb{G}(\Sigma)$ itself is rigid for all but finitely many Σ (see also [S]). In fact, $\mathbb{G}(S_{g,p})$ is rigid if $g \geq 2$ or ($g = 1$ and $p \geq 2$) or ($g = 0$ and $p \geq 5$).

Various other cases are known. For example the rigidity of (most of) the non-separating curve graphs was established in [Ir] (though that is not directly relevant to the present paper). Of particular interest here is when C is the set of all separating curves. In this case, we write $\mathbb{G}_s(\Sigma) = \mathbb{G}(\Sigma, C)$. 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 $\mathbb{G}_s(\Sigma) = \mathbb{G}(\Sigma)$, and this case was dealt with independently in [Ko] and [L]. Therefore, combining these we get:

Theorem 7.1. [Ko, L, BreM, Ki] *$\mathbb{G}_s(S_{g,p})$ is rigid if $g \geq 3$ or ($g = 2$ and $p \geq 2$) or ($g = 1$ and $p \geq 3$) or ($g = 0$ and $p \geq 5$).*

Note that, under the above conditions, we see that the isomorphism class of the graph $\mathbb{G}(\Sigma, C)$ determines Σ — since it determines $\mathrm{Map}(\Sigma)$ up to isomorphism, hence Σ.

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 $\kappa(\gamma) = \min\{\xi(X^-(\gamma)), \xi(X^+(\gamma))\}$.

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 it cuts off an $S_{0,3}$. Similarly it is 1-separating if it cuts off an $S_{1,1}$ or $S_{0,4}$. (This accords with the terminology in Section 6.) In this case, we write $F(\gamma)$ for the surface cut off by γ. If $\xi(\Sigma) \geq 4$, this is uniquely determined (and any ambiguity will be unimportant otherwise).

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

Recall that $\mathbb{G}_1(\Sigma)$ was the graph introduced at the end of Section 6, where we saw that it is determined by the coarse geometry of $\mathbb{P}(\Sigma)$. 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 $\mathbb{G}_1(\Sigma)$. It can however be recovered, at least in most cases, as we will show in Lemma 7.3.

In what follows we write $\mathbb{G}^c(\Sigma)$ for the complementary graph of $\mathbb{G}(\Sigma)$; that is, with the same vertex set and complementary edge set. If C is any set of curves, we write $\mathbb{G}^c(\Sigma, C)$ for the full subgraph of $\mathbb{G}^c(\Sigma)$ with vertex set C. Clearly this is complementary to $\mathbb{G}(\Sigma, C)$. We write $\mathbb{G}^c_1(\Sigma) = \mathbb{G}^c(\Sigma, C_1)$.

Given a subsurface X of Σ, we write $P(X) = \{\gamma \in C_1 \mid \gamma \prec X\}$. Clearly, $P(X)$ is invariant under the subgroup, $G(X)$, of $\mathrm{Map}(\Sigma)$, supported on X.

Definition. We say that X is *big* if $\xi(X) \geq 2$ and $P(X) \neq \varnothing$.

(Note that the latter condition is redundant if X has only one relative boundary component in Σ.)

In other words, X is big precisely if it properly contains a terminal subsurface of Σ (as defined in Section 6) — note that the terminal subsurface cut off by any element of $P(X)$ is properly contained in X.

One sees easily (applying the group, $G(X)$), that if X is big, then $P(X)$ is infinite, and the elements of $P(X)$ fill X. Moreover, $\mathbb{G}^c(\Sigma, P(X))$ is connected (of diameter 2).

Definition. 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. (Here the transverse orientation points from X^- to X^+.)

Given a subset $P \subseteq C_1$, write $L(P)$ for the set of elements of $C_1 \setminus P$ which are adjacent in $\mathbb{G}^c_1(\Sigma)$ to some element of P. By a *division* of $\mathbb{G}^c_1(\Sigma)$, we mean an ordered pair, $\underline{P} = (P^-, P^+)$, of disjoint infinite subsets, $P^-, P^+ \subseteq C_1$, such that

$\mathbb{G}^c(\Sigma, P^-)$ and $\mathbb{G}^c(\Sigma, P^+)$ 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 $P^\mp$. (We should imagine the curves in P^- and P^+ as lying in two complementary surfaces, X^+ and X^-, as we describe in the next lemma.)

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

Proof. Let $\underline{X}$ be a division of Σ. We have already observed that $P^\pm$ is infinite and that $\mathbb{G}^c(\Sigma, P^\pm)$ 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 $\mathbb{G}_1^c(\Sigma)$.

Conversely, suppose that $\underline{P}$ is a division of $\mathbb{G}_1^c(\Sigma)$. Let $X^\pm$ be the subsurface of Σ filled by the curves of $P^\pm$. Since $\mathbb{G}^c(\Sigma, P^\pm)$ 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^+). We need to show that these are complementary, in the sense that $\Sigma \setminus (X^- \cup X^+)$ consists only of annuli.

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 as claimed. 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 $\underline{P} = \underline{P}(\underline{X})$.

It is clear from the construction that this gives a bijection as required. □

We have therefore shown that we can "see" divisions of Σ in terms of the structure of $\mathbb{G}_1(\Sigma)$.

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 $\underline{X} \leq \underline{Y}$ and $\underline{X} \neq \underline{Y}$. Clearly this defines a partial order on the set of divisions. Similarly, if $\underline{P}, \underline{Q}$, are divisions of $\mathbb{G}^c(\Sigma)$, 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 7.2.

By definition, we are only allowed to "divide" Σ into big subsurfaces. However, we will also want to include terminal subsurfaces in this. To this end we make the following definition.

Definition. A *slice* of Σ is an ordered pair, $\underline{X} = (X^-, X^+)$, where each of X^- and X^+ is either a terminal 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 a terminal $S_{0,4}$.

Again, this can be recognised from $\mathbb{G}_1(\Sigma)$. 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 $\gamma \prec Y^-$ and the sign is $-$. Similarly $\underline{Y} < \underline{X}$ says that $\gamma \prec Y^+$ with sign $+$.

Using slices, we can now distinguish elements of C_{1HT} and C_{4HS}, at least if $\xi(\Sigma) \geq 6$. To this end, 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 $\mathbb{G}_1(\Sigma)$. For example, the statement that $\gamma \prec X^+$ is the same as saying that either $\underline{X}$ corresponds to $\underline{P}$ with $\gamma \in P^+$, or else $\underline{X}$ corresponds to $(\beta, -)$, where $\beta \in C_1$ is a curve disjoint from γ.

Lemma 7.3. *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 $F(\gamma) \subseteq U$. (Recall that $F(\gamma)$ is the complexity-1 subsurface bounded by γ.) 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 β_1 and to β_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 both of 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 $\xi(\Sigma) \geq 6$, $\Sigma \setminus F(\gamma)$ has complexity at least 4. 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, and each with connected relative boundary. 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 $W^+ \cap Y^-$, with disjoint interiors and 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, $\underline{X}_1 < \underline{X}_2 < \cdots < \underline{X}_m$, of length m.

Lemma 7.4. *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, $F_1, F_2, \ldots, F_{2g+p-2}$, 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 $i = 1, \ldots, 2g+p-3$. This gives a chain $\underline{X}_1 < \underline{X}_2 < \cdots < \underline{X}_{2g+p-3}$. Conversely, given any chain $\underline{X}_1 < \underline{X}_2 < \cdots < \underline{X}_m$, 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^- or X_m^+) must be an $S_{0,4}$, and this must accommodate two $S_{0,3}$'s. Similarly, if $g = 0$, we lose 2, since then both extreme surfaces will be $S_{0,4}$'s. □

Let us summarise what we have so far detected in terms of $\mathbb{G}_1(\Sigma)$. Note first that if $\xi(\Sigma) \leq 2$ then $\mathbb{G}_1(\Sigma) = \varnothing$, and if $\xi(\Sigma) = 3$, then $\mathbb{G}_1(\Sigma)$ 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, via Lemma 7.3, as we have already noted. Since we also know $m(\Sigma)$ by Lemma 7.4, we can also determine p.

We have shown:

Lemma 7.5. *Suppose that Σ, Σ' are compact orientable surfaces with $\mathbb{G}_1(\Sigma)$ isomorphic to $\mathbb{G}_1(\Sigma')$. 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.*

Lemmas 6.2, 6.4 and 7.5 together give us Theorem 1.3.

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 $\{S_{2,2}, S_{1,5}, S_{0,8}\}$ respectively). In fact, one can distinguish $S_{2,1}$ from $S_{1,4}$ and $S_{0,7}$ as we observe after Theorem 7.8.

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 (since at least one of X^- or X^+ is complexity-1, hence terminal, by the definition of a slice). We therefore suppose that $\xi(\Sigma) \geq 6$. 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 (again from the definition). But now we can detect $\text{genus}\, X^\pm$ as the maximal number of disjoint C_{1HT} curves contained in $X^\pm$. This shows that we can recognise simple slices as claimed.

We therefore have a means of describing (1+)-separating curves in Σ in terms of $\mathbb{G}_1(\Sigma)$. 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. (They are disjoint if and only if we can choose the transverse orientation so that the slices are nested.) In other words, we can reconstruct the whole of $\mathbb{G}_{1+}(\Sigma)$ from $\mathbb{G}_1(\Sigma)$.

We have shown:

Lemma 7.6. *Suppose that Σ, Σ' are compact orientable surfaces, and that $\xi(\Sigma) = \xi(\Sigma') \geq 4$. Then any isomorphism from $\mathbb{G}_1(\Sigma)$ to $\mathbb{G}_1(\Sigma')$ extends to an isomorphism from $\mathbb{G}_{1+}(\Sigma)$ to $\mathbb{G}_{1+}(\Sigma')$.*

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

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

This is taken to imply that $\xi(\Sigma) \geq 4$. (Otherwise $\mathbb{G}_{1+}(\Sigma)$ is either empty or a discrete set of points.) Note that if $p \leq 1$, then $\mathbb{G}_{1+}(\Sigma) = \mathbb{G}_s(\Sigma)$, and so applying the result of [BreM, Ki], given as Theorem 7.1 here, we see that $S_{g,0}$ and $S_{g,1}$ are of rigid type if $g \geq 3$. Also in [Bo6], it is shown that Σ is of rigid type if $g(\Sigma) + p(\Sigma) \geq 7$. We remark that [Mc] gives a different proof for $S_{0,p}$ for $p \geq 8$.

In summary, this shows:

Theorem 7.7. *If $g(\Sigma) + p(\Sigma) \geq 7$, or if $g(\Sigma) \geq 3$ and $p(\Sigma) \leq 1$, then Σ is of rigid type.*

(We suspect that this holds whenever $\mathbb{G}_{1+}(\Sigma)$ is non-empty and connected, i.e. when $\xi(\Sigma) \geq 4$ and $\Sigma \neq S_{2,1}$.)

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

Theorem 7.8. *Suppose that Σ, Σ' are compact orientable surfaces with $\mathbb{P}(\Sigma)$ quasi-isometric to $\mathbb{P}(\Sigma')$. Then $\xi(\Sigma) = \xi(\Sigma')$. Moreover, if $\xi(\Sigma) = \xi(\Sigma') \geq 6$, then Σ is homeomorphic to Σ'.*

Proof. Note that if $\xi(\Sigma) = 1$ if and only if $\mathbb{P}(\Sigma)$ is a quasitree; $\xi(\Sigma) = 2$ if and only if $\mathbb{P}(\Sigma)$ is hyperbolic and not a quasitree; and $\xi(\Sigma) \geq 3$ if and only if $\mathbb{P}(\Sigma)$ is not hyperbolic. We can therefore assume that $\xi(\Sigma) \geq 3$. As observed at the end of Section 6, we then have that $\mathbb{G}_1(\Sigma)$ and $\mathbb{G}_1(\Sigma')$ are isomorphic, and so the statement then follows by Lemma 7.5. □

A bit more can be said when $\xi(\Sigma) = 4$. Note that $\mathbb{G}_{1+}(S_{0,7})$ and $\mathbb{G}_{1+}(S_{1,4})$ are both connected. (The latter can be seen for example by a simple argument using the fact that $\mathbb{G}_s(S_{1,4})$ is connected.) However, $\mathbb{G}_{1+}(S_{2,1}) = \mathbb{G}_s(S_{2,1})$ is disconnected. (It retracts onto $\mathbb{G}_s(S_{2,0})$ deleting the puncture, and the latter graph has no edges.) Therefore $\mathbb{P}(S_{2,1})$ is not quasi-isometric to either $\mathbb{P}(S_{0,7})$ or $\mathbb{P}(S_{1,4})$.

It is well known that $\mathbb{P}(S_{1,2})$ is isomorphic to $\mathbb{P}(S_{0,5})$ and that $\mathbb{P}(S_{2,0})$ is isomorphic to $\mathbb{P}(S_{0,6})$. Indeed both isomorphisms are natural. (The isomorphisms arise from the hyperelliptic involutions on $S_{1,2}$ and on $S_{2,0}$, respectively. See, for example, Lemma 2.1 of [L], which shows that the respective curve graphs are isomorphic. One can check that the construction also applies to the pants graphs. Indeed, it is not hard to reconstruct the pants graph from the curve graph.)

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

Regarding quasi-isometric rigidity, we can now show:

Theorem 7.9. *Suppose that Σ is a compact orientable surface of rigid type, and that $\phi : \mathbb{P}(\Sigma) \longrightarrow \mathbb{P}(\Sigma)$ is a quasi-isometry. Then there is some $h \in \mathrm{Map}(\Sigma)$ such that if $a \in \mathbb{P}(\Sigma)$, then $\rho(\phi a, ha) \leq k$, where k depends only on $\xi(\Sigma)$ and the parameters of the quasi-isometry, ϕ.*

Note that Theorems 7.7 and 7.9 together imply Theorem 1.4.

Proof of Theorem 7.9. The map π given by Lemma 6.2 determines an automorphism of $\mathbb{G}_1(\Sigma)$. Therefore, by Lemma 7.6, after applying some element of $\mathrm{Map}(\Sigma)$, we can assume this to be the identity on $\mathbb{G}_1(\Sigma)$. In other words, if Y is any terminal subsurface of Σ, we have $\pi Y = Y$. But now if X is any admissible subsurface of Σ, each component of $\Sigma \setminus X$ is filled by terminal subsurfaces of Σ, in the sense that any curve in $\Sigma \setminus X$ must cross, or be contained in, some terminal subsurface. (Possibly it is itself a terminal subsurface.) Now these subsurfaces determine X uniquely, and so it follows that X must be fixed by π. In other words, π must be the identity, and so $\phi B(X) = B(X)$ for every admissible subsurface, X, of Σ.

Now suppose that τ is a good multicurve in Σ. Lemma 5.10 gives us a good multicurve τ' in Σ such that $\mathrm{hd}(T(\tau'), \phi(T(\tau)))$ is finite and bounded above in

terms of $\xi(\Sigma)$ and the parameters of ϕ. Also (as observed after Lemma 6.2) the map, π, given by Lemma 6.2, when restricted to $\mathcal{X}_N(\tau)$, agrees with the map π given by Proposition 5.11. 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 $\mathrm{hd}(T(\tau), \phi(T(\tau)))$ is uniformly bounded above for all good multicurves τ.

The remainder of the proof follows as with that of Theorem 5.12 of [Bo8]. Note that if τ, τ' are good multicurves which fill Σ, then by Lemma 4.6 here, $T(\tau)$ and $T(\tau')$ uniformly diverge, and so any point a bounded distance from both gets moved a bounded distance by ϕ. But this applies to all points of $\mathbb{P}(\Sigma)$, since $\mathrm{Map}(\Sigma)$ acts coboundedly on $\mathbb{P}(\Sigma)$. Again the bound depends only on $\xi(\Sigma)$ and the parameters of ϕ. □

We remark that our argument does not give any constructive means of determining k from the input parameters.

Note that $\mathbb{W}(\Sigma)$ is undefined (or trivial) for $S_{0,0}, S_{0,1}, S_{0,2}, S_{0,3}$. For $S_{0,4}, S_{1,0}, S_{1,1}$ it is quasi-isometric to an infinite-valence tree (or the Farey graph), and is certainly not quasi-isometrically rigid in these cases. The cases left unresolved by the present paper are therefore:

$$S_{0,5},\ S_{0,6},\ S_{1,2},\ S_{1,3},\ S_{1,4},\ S_{1,5},\ S_{2,0},\ S_{2,1},\ S_{2,2},\ S_{2,3},\ S_{2,4},\ S_{3,2},\ S_{3,3},\ S_{4,2}.$$

The pairs $\{S_{0,5}, S_{1,2}\}$ and $\{S_{0,6}, S_{2,0}\}$ are essentially equivalent, since as we have noted above, they have naturally isomorphic pants graphs (identified via the hyperelliptic involutions on $S_{1,2}$ and $S_{2,0}$ respectively). It follows that the respective Weil-Petersson geometries are naturally quasi-isometric. Moreover, the hyperelliptic involutions are central in the respective mapping class groups, and the quasi-isometry can be assumed equivariant modulo quotienting by this involution.

There are therefore essentially twelve cases left open.

The case of $\mathbb{W}(S_{0,5}) \equiv \mathbb{W}(S_{1,2})$ is different in nature, since this space hyperbolic, as we have already noted. One could probably deal with a few more cases by elaborating on the arguments of [Bo6], though a complete answer is likely to require some new ideas.

References

[Ao] T.Aougab, *Uniform hyperbolicity of the graphs of curves* : Geom. Topol. **17** (2013) 2855–2875.

[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.Hedlíková, *Median algebras* : Discrete Math. **45** (1983) 1–30.

[Be] J.A.Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space* : Geom. Topol. **10** (2006) 1523–1578.

[BeDS] J.A.Behrstock, C.Druţu, M.Sapir, *Median structures on asymptotic cones and homomorphisms into mapping class groups* : Proc. London Math. Soc. **3** (2011) 503–554.

[BeHS1] J.Behrstock, M.F.Hagen, A.Sisto, *Hierarchically hyperbolic spaces II: Combination theorems and the distance formula* : Pacific J. Math. **299** (2019) 257–338.

[BeHS2] J.A.Behrstock, M.F.Hagen, A.Sisto, *Quasiflats in hierarchically hyperbolic spaces* : preprint, 2017, posted at arXiv:1704.04271.

[BeKMM] J.A.Behrstock, B.Kleiner, Y.Minsky, L.Mosher, *Geometry and rigidity of mapping class groups* : Geom. Topol. **16** (2012) 781–888.

[BeM1] J.A.Behrstock, Y.N.Minsky, *Dimension and rank for mapping class groups* : Ann. of Math. **167** (2008) 1055–1077.

[BeM2] J.A.Behrstock, Y.N.Minsky, *Centroids and the rapid decay property in mapping class groups* : J. London Math. Soc. **84** (2011) 765–784.

[Bo1] B.H.Bowditch, *Coarse median spaces and groups* : Pacific J. Math. **261** (2013) 53–93.

[Bo2] B.H.Bowditch, *Embedding median algebras in products of trees* : Geom. Dedicata **170** (2014) 157–176.

[Bo3] B.H.Bowditch, *Uniform hyperbolicity of the curve graphs* : Pacific J. Math. **269** (2014) 269–280.

[Bo4] B.H.Bowditch, *Some properties of median metric spaces* : Groups, Geom. Dyn. **10** (2016) 279–317.

[Bo5] B.H.Bowditch, *Quasi-isometric maps between direct products of hyperbolic spaces* : Internat. J. Algebra and Computation **26** (2016) 619–633.

[Bo6] B.H.Bowditch, *Rigidity of the strongly separating curve graph* : Michigan Math. J. **65** (2016) 813–832.

[Bo7] B.H.Bowditch, *Large-scale rank and rigidity of the Teichmüller metric* : J. Topol. **9** (2016) 985–1020.

[Bo8] B.H.Bowditch, *Large-scale rigidity properties of the mapping class groups* : Pacific J. Math. **293** (2018) 1–73.

[Bo9] B.H.Bowditch, *Quasiflats in coarse median spaces* : preprint, Warwick (2019).

[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.

[BroMar] J.Brock, D.Margalit, *Weil-Petersson isometries via the pants complex* : Proc. Amer. Math. Soc. **135** (2007) 795–803.

[BroMas] J.Brock, H.Masur, *Coarse and synthetic Weil-Petersson geometry: quasi-flats, geodesics and relative hyperbolicity* : Geom. Topol. **12** (2008) 2453–2495.

[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.

[ClRS] M.Clay, K.Rafi, S.Schleimer, *Uniform hyperbolicity of the curve graph via surgery sequences* : Algebr. Geom. Topol. **14** (2014) 3325–3344.

[DrK] C.Druţu, M.Kapovich, *Geometric group theory* : Amer. Math. Soc. Colloq. Publ. No. 63, American Mathematical Society (2018).

[DyP] A.Dyubina, I.Polterovich, *Explicit constructions of universal $\mathbb{R}$-trees and asymptotic geometry of hyperbolic spaces* : Bull. London Math. Soc. **33** (2001) 727–734

[EMR1] A.Eskin, H.Masur, K.Rafi, *Large scale rank of Teichmüller space* : Duke Math. J. **116** (2017) 1517–1572.

[EMR2] A.Eskin, H.Masur, K.Rafi, *Rigidity of Teichmüller space* : Geom. Topol. **22** (2018) 4259–4306.

[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).

[HaT] A.Hatcher, W.Thurston, *A presentation for the mapping class group* : Topology **19** (1980) 221–237.

[HeP] S.Hensel, P.Przytycki, *The ending lamination space of the five-punctured sphere is the Nöbeling curve* : J. Lond. Math. Soc. **84** (2011) 103–119.

[HePW] S.Hensel, P.Przytycki, R.C.H.Webb, *1-slim triangles and uniform hyperbolicity for arc graphs and curve graphs* : J. Eur. Math. Soc. **17** (2015) 755–762.

[Ir] E.Irmak, *Complexes of nonseparating curves and mapping class groups* : Michigan Math. J. **54** (2006) 81-110.

[Is] J.R.Isbell, *Median algebra* : 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.

[KlL] 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 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.

[Mc] A.McLeay, *Normal subgroups of the braid groups and the metaconjecture of Ivanov* : prepint 2018, posted at arXiv:1801.05209.

[NWZ] G.A.Niblo, N.Wright, J.Zhang, *A four point characterisation for coarse median spaces* : Groups Geom. Dyn. **13** (2019) 939–980.

[R] M.A.Roller, *Poc-sets, median algebras and group actions, an extended study of Dunwoody's construction and Sageev's theorem* : Habilitationsschrift, Regensberg, 1998.

[S] K.J.Shackleton, *Combinatorial rigidity in curve complexes and mapping class groups* : Pacific J. Math. **230** (2007) 217–232.

[VaW] L.van den Dries, A.J.Wilkie, *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, *Uniform bounds for bounded geodesic image theorems* : J. Reine Angew. Math. **709** (2015) 219–228.

[Wo] S.A.Wolpert, *Geometry of the Weil-Petersson completion of Teichmüller space* : in "Surveys in differential geometry, Vol. VIII" International Press. (2003) 357–393.

[Y] S.-K.Yeung, *Quasi-isometry of metrics on Teichmüller spaces* : Int. Math. Res. Not. **4** (2005) 239–255.

[Z1] R.Zeidler, *Coarse median structures on groups* : Masterarbeit, University of Vienna, 2013.

[Z2] R.Zeidler, *Coarse median structures and homomorphisms from Kazhdan groups* : Geom. Dedicata **180** (2016) 49–68.

Mathematics Institute, University of Warwick, Coventry, CV4 7AL, Great Britain.