

GEOMETRY AND RIGIDITY OF MAPPING CLASS GROUPS

JASON BEHRSTOCK, BRUCE KLEINER, YAIR MINSKY, AND LEE MOSHER

ABSTRACT. We study the large scale geometry of mapping class groups $\mathcal{MCG}(S)$, using hyperbolicity properties of curve complexes. We show that any self quasi-isometry of $\mathcal{MCG}(S)$ (outside a few sporadic cases) is a bounded distance away from a left-multiplication, and as a consequence obtain quasi-isometric rigidity for $\mathcal{MCG}(S)$, namely that groups quasi-isometric to $\mathcal{MCG}(S)$ are equivalent to it up to extraction of finite-index subgroups and quotients with finite kernel. (The latter theorem was proved by Hamenstädt using different methods).

As part of our approach we obtain several other structural results: a description of the tree-graded structure on the asymptotic cone of $\mathcal{MCG}(S)$; a characterization of the image of the curve complex projections map from $\mathcal{MCG}(S)$ to $\prod_{Y \subseteq S} \mathcal{C}(Y)$; and a construction of Σ -hulls in $\mathcal{MCG}(S)$, an analogue of convex hulls.

CONTENTS

| | |
|-----------------------------------|----|
| 1. Introduction | 1 |
| 2. Preliminaries | 8 |
| 3. Product regions and cubes | 23 |
| 4. Consistency theorem | 37 |
| 5. Σ -hulls | 43 |
| 6. Contractibility and homology | 47 |
| 7. Separation properties | 52 |
| 8. Local finiteness for manifolds | 63 |
| 9. Germs and orthants | 75 |
| 10. Finishing the proofs | 87 |
| References | 95 |

1. Introduction

In this paper we investigate the coarse geometry of the mapping class group $\mathcal{MCG}(S)$ of an oriented finite-type surface S , with our main goal being a proof of quasi-isometric rigidity for the group. Along the way we develop a number of tools which we hope will be of independent interest.

Partially supported by NSF grants DMS-0812513, DMS-0701515, DMS-0504019, DMS-0706799.

The following classification theorem for quasi-isometries applies for all but a standard collection of exceptional cases (definitions are in §2, and complete statements of these theorems handling the exceptions are in §10.)

Theorem 1.1. (Classification of Quasi-Isometries) *If S has complexity at least 2 and is not a two-holed torus, quasi-isometries of $\mathcal{MCG}(S)$ are uniformly close to isometries induced by left-multiplication.*

Note that an immediate consequence of this theorem is that, barring a few exceptional cases, $\mathcal{MCG}(S)$ is isomorphic to its own quasi-isometry group, i.e. the group of quasi isometries modulo those that are finite distance from the identity (see Corollary 10.1).

Theorem 1.1 also implies the following rigidity property (see Section 10 for an alternate statement of this result which also covers the complexity 2 cases):

Theorem 1.2. (Quasi-Isometric Rigidity) *If S has complexity greater than 2 and Γ is a finitely generated group quasi-isometric to $\mathcal{MCG}(S)$ then there is a homomorphism from Γ to $\mathcal{MCG}(S)$, with finite kernel and finite-index image.*

Theorem 1.2 was proved by Hamenstädt in [26]. The two proofs have a similar flavor in broad outline, although the underlying machinery supporting the outline is different. Mosher-Whyte [44] previously established the theorem for once-punctured surfaces of genus at least 2, using quite different methods.

The study of coarse properties of groups can be said to have started with Milnor and Švarc [42, 41, 54]. This advanced considerably with Stallings' theorem on ends of groups [52], the Mostow rigidity theorem [47], and Gromov's theorem on groups of polynomial growth [22]. Gromov [24] proposed an ambitious program for the study and classification of groups by their coarse geometric properties, which continues to guide current research. One branch of this program is the classification of groups up to quasi-isometry, i.e. determining up to isomorphism (or commensurability) the groups in a quasi-isometry class. In the last twenty years much progress has been made on different cases of this program:

- Groups quasi-isometric to lattices in semi-simple Lie groups (Gromov, Sullivan, Tukia, Casson-Jungreis, Gabai, Hinkkanen, Pansu, Cannon-Cooper, Schwartz, Farb-Schwartz, Kleiner-Leeb, Eskin [9, 10, 14, 20, 21, 23, 24, 25, 29, 35, 48, 51, 53, 55]).
- Groups quasi-isometric to Euclidean buildings (Kleiner-Leeb [35]).
- Groups quasi-isometric to Fuchsian buildings (Bourdon-Pajot, Xie [8, 57]).
- Graphs of n -dimensional Poincaré duality groups, for some fixed n (Farb-Mosher, Whyte, Mosher-Sageev-Whyte, Papasoglu [18, 19, 45, 46, 49, 56]).
- Certain solvable semi-direct products $\mathbb{R}^k \rtimes \mathbb{R}^l$ (Eskin-Fisher-Whyte, Dymarz, Peng [13, 15, 50]).

Research in mapping class groups has long been motivated by a drive to find analogies with lattices. The list of such results is intriguing but incomplete: see for example McCarthy [40] and Ivanov [30] (a Tits-alternative), Ivanov [32] and Farb-Masur [17] (non-arithmeticity and superrigidity), Harvey [28] (analogy to Bruhat-Tits buildings), Harer [27] (homology stability), Farb-Lubotzky-Minsky [16] (rank 1 behavior) and Andersen [1] (failure of property (T)). The expectation has been that quasi-isometric rigidity should hold for these groups, but the standard tools have not previously found purchase.

Structural results

Our analysis builds on the geometry of curve complexes of surfaces developed in [38, 39, 2, 7], but introduces a number of new ideas and techniques, which have some further applications (for example the rapid-decay property for $\mathcal{MCG}(S)$, Behrstock-Minsky [6]). Among these are a theorem characterizing the curve-complex projection image of $\mathcal{MCG}(S)$; a coarse convex hull construction in $\mathcal{MCG}(S)$; and a theory of “jets,” which analyzes infinitesimal directions in the asymptotic cone of $\mathcal{MCG}(S)$ in order to control separation properties. We summarize these below (precise statements appear later in the paper):

Consistency Theorem. To the isotopy class of each essential subsurface $W \subseteq S$ is associated a δ -hyperbolic complex $\mathcal{C}(W)$, its curve complex (or twist complex if W is an annulus). There is a coarse-Lipschitz projection $\pi_W : \mathcal{MCG}(S) \rightarrow \mathcal{C}(W)$, and combining these over all W we obtain the *curve complex projections map*

$$\Pi : \mathcal{MCG}(S) \rightarrow \prod_W \mathcal{C}(W).$$

The Quasidistance Formula of [39] (see Theorem 2.8 below) shows that this map is, in a limited sense, like a quasi-isometric embedding to the ℓ^1 metric on the product.

Behrstock established *consistency inequalities* (see §4) that are satisfied by the image of Π , and in this paper we prove that these give a coarse characterization of the image:

Theorem 4.3 *The consistency inequalities give a necessary and sufficient condition for a point in $\prod_W \mathcal{C}(W)$ to be near the image of Π .*

This theorem makes it possible to state and analyze many constructions in $\mathcal{MCG}(S)$ simply in terms of what they look like in the projection image, where the hyperbolicity of the $\mathcal{C}(W)$ ’s can be exploited.

Σ -hulls. An important construction that follows from the Consistency Theorem is a coarse substitute for a convex hull in $\mathcal{MCG}(S)$. In a δ -hyperbolic space, the union of all geodesics joining a set of points is quasi-convex; let us call this simply the “hull” of the set. For a finite subset $A \subset \mathcal{MCG}(S)$ (and

suitable $\epsilon > 0$) we describe a set $\Sigma_\epsilon(A) \subset \mathcal{MCG}(S)$, which is exactly the set of points that project in each factor of $\prod_W \mathcal{C}(W)$ to within ϵ of the hull of the image of A . This notion of hull satisfies a number of useful properties.

Proposition 5.2 and Lemma 5.4. *Fixing the cardinality of A , $\Sigma_\epsilon(A)$ depends in a coarse-Lipschitz way on A with respect to the Hausdorff metric. Its diameter is controlled by $\text{diam}(A)$, and it admits a coarse retraction $\mathcal{MCG}(S) \rightarrow \Sigma_\epsilon(A)$, which itself has Lipschitz dependence on A .*

These hulls, and their rescaling limits in the asymptotic cone, give us a way to build and control singular chains in the cone, which we use a la Kleiner-Leeb [35] to describe top-dimensional flats using local homology arguments.

A different application of Σ -hulls will appear in Behrstock-Minsky [6] where $\mathcal{MCG}(S)$ is shown to have the rapid decay property of Haagerup/Jolissaint.

Jets and separation. In Behrstock-Minsky [7], it was shown that any pair of distinct points in the asymptotic cone of $\mathcal{MCG}(S)$ could be separated by an ultralimit of subsets associated to mapping class groups of subsurfaces. This enabled an inductive argument to compute the compact topological dimension of the cone. Here we refine our understanding of the components of the complement of such a set, introducing the notion of a *jet*. A jet is a sequence of geodesic segments in curve complexes of subsurfaces of S (modulo an ultrafilter), such that asymptotic behavior of projections to these segments determines the division of the complement into connected components. In the outline below we motivate this notion via an analogy to δ -hyperbolic spaces.

As an application of these ideas we also pause in Section 7.4 to give a characterization of the tree-graded structure (a la Druţu-Sapir [12]) of the cone of $\mathcal{MCG}(S)$. This characterization has been used in the recent proof that any finitely generated group with Property (T) has at most finitely many pairwise non-conjugate homomorphisms into a mapping class group (Behrstock-Druţu-Sapir [5]).

Outline of the proof

In this section we discuss the broad structure of the paper and the proof of Theorem 1.1 (the proof of Theorem 1.2 from 1.1 is fairly standard and will be given in Section 10).

Definitions and other preliminaries will be given in Section 2.

Coarse preservation of Dehn twist flats. In order to control a quasi-isometry $f: G \rightarrow G$ of any group, we wish to identify structures in G which are robust enough to be preserved by f , and intricate enough that they can only be preserved in the obvious ways. In the case of $\mathcal{MCG}(S)$, these structures are

(*maximal*) *Dehn twist flats*, which are cosets of maximal-rank free-abelian subgroups generated by Dehn twists. Later we will actually work with equivalent subsets of the *marking complex* $\mathcal{M}(S)$, which is our geometric model of choice for $\mathcal{MCG}(S)$.

Theorem 10.3 states that a quasi-isometry $f: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(S)$ *coarsely preserves* the set of Dehn twist flats. That is, the image of any such flat is within finite Hausdorff distance of another such flat, with the bound depending only on the quality of the quasi-isometry.

Once Theorem 10.3 is established, we can apply known results to prove Theorem 1.1. The coarse permutation of flats induces an automorphism of the curve complex of S which by a theorem of Ivanov–Korkmaz–Luo [33, 36, 37] is induced by some mapping class $\phi \in \mathcal{MCG}(S)$ (when $S = S_{1,2}$ this is not quite right but we ignore this for now). This gives us the desired element of $\mathcal{MCG}(S)$ and it is then not hard to show that left-multiplication by ϕ is uniformly close to the quasi-isometry f .

Preservation of asymptotic Dehn twist flats. Theorem 10.3 will be proven, following Kleiner–Leeb [35], by reduction to the *asymptotic cone* of $\mathcal{MCG}(S)$. The asymptotic cone, which we denote \mathcal{M}_ω , is a limit obtained by unbounded rescaling of the word metric on $\mathcal{MCG}(S)$. Extracting this limit requires the choice of an ultrafilter, although our results hold for any choice — see §2.3 for details.

A quasi-isometry of $\mathcal{MCG}(S)$ converges after rescaling to a bilipschitz homeomorphism of \mathcal{M}_ω , and Dehn twist flats limit to bilipschitz-embedded copies of Euclidean space. Thus our goal in this context is to show that these asymptotic Dehn twist flats are permuted by the limiting map:

Theorem 10.2. *Except when S is a one-holed torus or four-holed sphere, any homeomorphism $f: \mathcal{M}_\omega(S) \rightarrow \mathcal{M}_\omega(S)$ permutes the Dehn twist flats in $\mathcal{M}_\omega(S)$.*

The proof of Theorem 10.2 will take up most of the paper.

The reduction of Theorem 10.3 to Theorem 10.2 will be discussed in detail in Section 10.2. Let us now sketch the proof of Theorem 10.2.

Structure of $\mathcal{MCG}(S)$ via hyperbolicity. We begin in Sections 3–5 by refining the tools developed in Masur–Minsky [38, 39] and Behrstock [3] to study the coarse structure of $\mathcal{MCG}(S)$ using properties of curve complexes.

In Section 3, we analyze the structure of subsets of $\mathcal{MCG}(S)$ that come, in a simple way, from restrictions on some of the coordinates of the subsurface projection maps. In particular what we call a *product region* is a set $\mathcal{Q}(\mu)$ corresponding to all markings of S that contain a fixed partial marking μ . Coarsely this is the same as a coset in $\mathcal{MCG}(S)$ of a stabilizer of the partial marking. The

factors in the product structure are indexed by the different components of the subsurface of S on which the partial marking μ is undefined. In this language a Dehn twist flat is defined as a set $\mathcal{Q}(\mu)$ where μ is a pants decomposition whose curves are unmarked by any choice of transversal. Choosing transversals gives one real-valued parameter for each curve in μ , making a Dehn twist flat quasi-isometric to $\mathbf{R}^{\xi(S)}$, where $\xi(S)$ is the number of components of a pants decomposition of S .

A *cube* is a special subset of a product region which is in fact quasi-isometric to a Euclidean cube in a way compatible with the product structure. We show that product regions and cubes are quasi-isometrically embedded subsets of $MCG(S)$, generalizing a result of [39]. In Lemmas 3.3 and 3.7 we analyze the sets, called *junctions*, along which two of these regions are close — junctions are generalizations of the coarse intersections of quasiconvex sets in a hyperbolic space.

In Section 4, we prove the Consistency Theorem 4.3, which characterizes the image of the curve complexes projection map. This theorem can be applied in many cases to supplant the use of the *hierarchy paths* from Masur–Minsky [39]. Although useful, these paths are technical to define and to work with. Hence, for the most part we have avoided using them and Theorem 4.3 is one of the main tools that allows us to do this.

In Section 5, we use Theorem 4.3 to define and study Σ -hulls.

Local homology via hulls. In Section 6 we use Σ -hulls in order to study the local homology properties of the asymptotic cone. The coarse properties established in Proposition 5.2 and Lemma 5.4 imply, in the cone, that ultralimits of Σ -hulls are contractible and have controlled geometry (Lemmas 6.1 and 6.2). This allows us to use them to build singular chains which Σ -compatible, meaning that each simplex is contained in the Σ -hull of its 0-skeleton. With this, and the results of Behrstock–Minsky [7], we prove local homological dimension bounds, a result originally established by Hamenstädt [26]. We also obtain Corollary 6.9, an analogue of a result of Kleiner–Leeb [35], which controls embedded top-dimensional manifolds in the cone. Σ -compatible chains will also be crucial in Sections 7 and 8.

Separation via jets. In Section 7, we refine the results of [3] and [7] to analyze separation properties of the asymptotic cone.

One can consider the δ -hyperbolic case in order to describe the basic intuition behind these separation arguments. If X is δ -hyperbolic and $\{g_n\}$ is a sequence of geodesic segments of lengths going to infinity, we obtain an ultralimit g_ω in the asymptotic cone X_ω , which may be a geodesic segment or a point (assume the latter, for simplicity). Nearest-point projection $\pi_n : X \rightarrow g_n$ yields a relation on $X_\omega \setminus g_\omega$: say that two ultralimits x_ω and y_ω are equivalent if for

representative sequences x_n and y_n , the sequence of distances $d(\pi_n(x_n), \pi_n(y_n))$ is ω -almost everywhere bounded.

It is a nice exercise to check directly from δ -hyperbolicity that this gives a well-defined equivalence relation, whose equivalence classes are open, and hence these classes are separated from each other by g_ω . In particular if x_n and y_n are always projected to opposite ends of g_n and neither x_ω nor y_ω lies in g_ω then they are separated by it. When g_ω is a point we call the sequence $\{g_n\}$ a *microscopic jet*.

Theorem 7.2 gives an analogous statement for $\mathcal{MCG}(S)$, where δ -hyperbolicity of individual curve complexes is exploited in a similar way, and separating sets are not points but product regions.

We will also have need to think about the setting where, appealing again to our δ -hyperbolic analogy, one of our points x_ω or y_ω may lie in g_ω . For this we introduce a finer analysis of what we call macroscopic jets with either *gradual* or *sudden growth*, and prove a suitable separation result in Theorem 7.7. In this case we show that certain of the components are acyclic, and this is where Σ -hulls come into the proof.

Section 7.4 is a digression in which we use these ideas to characterize the *pieces* of the *tree-graded structure* of the asymptotic cone of $\mathcal{MCG}(S)$, in the sense of Druţu-Sapir [12]. Although this is not needed for the rest of the proof, it is a structural fact which follows directly from our techniques and is likely of independent interest.

Finiteness for manifolds in the asymptotic cone. In Section 8 we apply the foregoing results to prove a local finiteness theorem for manifolds in the asymptotic cone.

Theorem 8.5 shows that the Σ -hull of a finite set in a top-dimensional manifold in the cone is always contained in a finite number of cubes. Most of the work is done in Theorem 8.1, which uses the separation theorems to control which subsurface projections of the finite set can grow without bound. This allows us to control the structure of paths connecting points in the set which behave in an efficient way with respect to their curve complex projections, e.g., hierarchy paths.

Finally, Theorem 8.7 states that any top-dimensional manifold is, locally at any point, contained in a finite number of cubes. This uses the results of Section 6 — in particular Theorem 8.5 and a triangulation argument allow us to approximate any sphere in the manifold as the boundary of a chain supported in finitely many cubes, and Corollary 6.9 implies the ball in the manifold bounded by the sphere is therefore contained (except for a small error near the boundary) in these cubes as well.

Orthant complex. In Section 9 we use the finiteness theorem to study the local structure of manifolds in the asymptotic cone, reducing it to a combinatorial question about the *complex of orthants*, which is the complex of germs of cubes with a corner at a given point \mathbf{x} . The starting point, using Theorem 8.7, is the fact that the germ of any top-dimensional manifold at \mathbf{x} is equal to a finite collection of orthants. This allows us to characterize the structure of the complex of orthants using purely topological properties, and in particular (Corollary 9.8) to characterize the germs of Dehn-twist flats in the cone. This means that any homeomorphism of the cone must permute the germs of Dehn-twist flats. A simple local-to-global argument gives the proof of Theorem 10.2.

2. Preliminaries

In Section 2.1 we review the foundations of curve complexes and marking complexes. In Section 2.2 we review projection maps between curve complexes of subsurfaces, and how they are used to obtain a quasidistance formula for the marking complex. The main references are [39] and [7]. In Section 2.3 we review asymptotic cones.

2.1. Curves and markings

2.1.1. Basic definitions. A *finite type surface* X is an oriented surface homeomorphic to a closed surface minus a finite set of points. The missing points are in one-to-one correspondence with the ends of X , and these are referred to as the *punctures* of X . If X is connected we denote $X = S_{g,n}$ where g is the genus and n the number of punctures, and we quantify the *complexity* of $S_{g,n}$ by $\xi(S_{g,n}) = \max\{3g - 3 + n, 0\}$. All surfaces in this paper will be of finite type.

Throughout the paper we will consider a single connected “ambient” surface S such that $\xi(S) \geq 2$ — so S is not a sphere with ≤ 4 punctures nor a torus with ≤ 1 puncture. We will also consider subsurfaces of S for which $\xi \geq 1$, as well as subannuli of S .

The (*extended*) *mapping class group* of S is the group

$$\mathcal{MCG}(S) = \text{Homeo}(S) / \text{Homeo}_0(S)$$

where $\text{Homeo}(S)$ is the group of homeomorphisms of S , and $\text{Homeo}_0(S)$ is the normal subgroup of homeomorphisms isotopic to the identity. We will often implicitly consider isotopy classes, i.e. $\text{Homeo}_0(S)$ -orbits, of various objects such as subsurfaces and simple closed curves. When this relation is explicit we will denote it by $\stackrel{i}{\cong}$.

2.1.2. Curves. An *essential curve* on a finite type surface X is an embedded circle γ such that if X is not an annulus then no component of $X - \gamma$ is a disc or a once-punctured disc, and if X is an annulus then γ is a core of X .

An *essential curve system* on X is a nonempty collection C of finitely many pairwise disjoint essential curves.

If X is connected then a curve system C on X is called a *pants decomposition* if each component of $X - C$ is a three-punctured sphere, a *pair of pants*. A nonempty pants decomposition exists on X if and only if $\xi(X) \geq 1$, in which case its number of curves is $\xi(X)$.

Given any finite type surface X , all pairwise nonisotopic, maximal, essential curve systems on X have the same cardinality, a number denoted $r(X)$ and called the *rank* of X , equal to the sum of the ξ -values of the components of X plus the number of annulus components. When X is an essential subsurface of S (see below), this number $r(X)$ equals the locally compact dimension of certain subsets of the asymptotic cone of $\mathcal{MCG}(S)$ associated to X ; see Lemma 3.5.

Given two essential curve systems C, C' , we may always isotope one of them so that they are in *efficient position*, which means that C, C' are transverse and no component of $S - (C \cup C')$ has closure which is a *bigon*, a nonpunctured disc whose boundary consists of an arc of C and an arc of C' . We say that C and C' *overlap* if, after putting them in efficient position, the intersection is nonempty. The concept of overlap will be generalized below.

2.1.3. The lattice of essential subsurfaces. An *essential subsurface* of S is a subsurface $Y \subset S$ with the following properties.

- Y is a union of (not necessarily all) complementary components of an essential curve system C . Denote $C \cap \bar{Y}$ by ∂Y , the boundary curves of Y .
- No two components of Y are isotopic — equivalently, no two annulus components are isotopic.
- Each nonannulus component of Y has $\xi \geq 1$, equivalently, no component is a 3-punctured sphere.

Essential subsurfaces of S are identified when they are isotopic in S . Note that two isotopic essential subsurfaces need not be ambient isotopic, for instance the complement of a single essential curve c is isotopic to but not ambient isotopic to the complement of an annulus neighborhood of c .

Given an essential subsurface X of S , let $\Gamma(X)$ denote the set of isotopy classes in S of simple closed curves that are essential in X . Note that a boundary curve of X has isotopy class in $\Gamma(X)$ if and only if it is isotopic to the core of an annulus component of X . Because we have excluded 3-holed spheres, $\Gamma(X)$ is empty if and only if X is empty. Note that $r(X)$ equals the maximum cardinality of a curve system in S whose elements are in $\Gamma(X)$.

On the set of essential subsurfaces define a relation $X \textcircled{=} Y$, read “ X is an essential subsurface of Y ”, to mean $\Gamma(X) \subset \Gamma(Y)$.

Lemma 2.1. *The relation \subseteq is a partial order on the set of isotopy classes of essential subsurfaces of S (including \emptyset). In particular, X is isotopic to Y if and only if $\Gamma(X) = \Gamma(Y)$. Moreover there exist binary operations \circledast , and \circledcirc , called essential union and essential intersection, which have the following properties:*

- (1) $X \circledcirc Y$ is the unique \subseteq -minimal essential subsurface Z such that $X \subseteq Z$ and $Y \subseteq Z$.
- (2) $X \circledast Y$ is the unique \subseteq -maximal essential subsurface Z such that $Z \subseteq X$ and $Z \subseteq Y$.

In other words we have a lattice whose partial order is \subseteq and whose meet and join operations are \circledast and \circledcirc , respectively.

Proof. To define these operations it is helpful to fix a complete hyperbolic metric on S . Every essential curve has a unique geodesic representative. Every connected essential subsurface X which is not an annulus is represented by the appropriate component of the complement of the union of the geodesic representatives of ∂X . Every essential subannulus is represented by the geodesic representative of its core. We call this the *geodesic representative* of a connected essential subsurface. Note that disjoint components of an essential subsurface have disjoint geodesic representatives, even when annuli are involved.

Now we can see that $\Gamma(X)$ determines X as follows. If $C \subset \Gamma(S)$, then for any finite subset of C we can take a regular neighborhood of the union of geodesic representatives, fill in disk or punctured-disk components of the complement, and obtain an essential subsurface. Any exhaustion of C by finite sets gives an increasing sequence of such subsurfaces, which must therefore eventually stabilize up to isotopy. This uniquely determines an essential subsurface which we call $\text{Fill}(C)$. One easily shows $X \stackrel{i}{=} \text{Fill}(\Gamma(X))$ provided $\Gamma(X) \neq \emptyset$, that is, if X is not a pair of pants. It follows immediately that \subseteq is a partial order.

Let us now show that \circledcirc is defined. Given X and Y , let $Z = \text{Fill}(C)$ where $C = \Gamma(X) \cup \Gamma(Y)$. Any curve in C either overlaps some other curve in C , in which case it is essential in a nonannulus component of Z , or it does not, in which case it is the core of an annulus component of Z and again essential. Therefore $\Gamma(X) \cup \Gamma(Y) \subseteq \Gamma(Z)$, so that $X \subseteq Z$ and $Y \subseteq Z$. Z is minimal with respect to this property because if Z' is a competitor then every finite subset of C is realized geodesically in the geodesic representative of Z' , and hence $Z \subseteq Z'$. Uniqueness follows from the fact that \subseteq is a partial order. We therefore set $X \circledcirc Y = Z$.

In fact we note that \circledcirc is defined for arbitrary collections $\{X_i\}$, merely by letting $C = \cup \Gamma(X_i)$. Now we can obtain $X \circledast Y$ satisfying (3) by taking the essential union of $\{Z : Z \subseteq X \text{ and } Z \subseteq Y\}$. \square

Here are a few remarks on the proof.

Notice that $X \Subset Y$ if and only if the geodesic representative of each component of X is pointwise contained in the geodesic representative of some component of Y . This is in turn equivalent to saying that each component of X is isotopic to an essential subsurface of a component of Y (where we allow an annulus to be an essential subsurface of itself).

It is helpful to notice that $\Gamma(X \circledast Y) = \Gamma(X) \cap \Gamma(Y)$. This is because any element γ in $\Gamma(X) \cap \Gamma(Y)$ is the core of an essential annulus A in both, hence $A \Subset X \circledast Y$ by (2), so $\gamma \in \Gamma(X \circledast Y)$. The other direction follows from the fact that $X \circledast Y$ is essentially contained in both X and Y .

We also define the *essential complement* X^c to be the maximal essential subsurface Z whose geodesic representative is disjoint from that of X . More concretely X^c is the union of complementary components of X that are not 3-holed spheres, together with an annulus for each component of ∂X that is not isotopic into an annulus of X . (This definition agrees with that in Behrstock-Minsky [7]). Note that essential complement does not behave like a true lattice theoretic complement operator, in that $(X^c)^c$ need not be isotopic to X , and $X \circledast X^c$ is usually not S ; for example, if X is a regular neighborhood of a pants decomposition on S then $X^c = \emptyset$.

On the other hand, the essential complement does satisfy the following easily verified formula, which in the asymptotic cone will allow us to make sense of codimension:

Proposition 2.2. *For each essential subsurface $X \subset S$ we have*

$$r(X) + r(X^c) = r(S) = \xi(S)$$

□

2.1.4. Curve complex. We associate a simplicial complex $\mathcal{C}(Y)$ with each connected surface Y with $\xi(Y) \geq 1$, as well as for each essential subannulus of our ambient surface S . For $\xi(Y) \geq 1$, the vertex set $\mathcal{C}_0(Y)$ of $\mathcal{C}(Y)$ is $\Gamma(Y)$, the isotopy classes of essential curves, and for $\xi(Y) > 1$, k -simplices correspond to sets of $k + 1$ vertices with disjoint representatives. Hence $\dim \mathcal{C}(Y) = \xi(Y) - 1$. When $\xi(Y) = 1$, we place an edge between any two vertices whose geometric intersection number is the smallest possible on Y (1 when $Y = S_{1,1}$ and 2 when $Y = S_{0,4}$). See Harvey [28], Ivanov [31] and Masur-Minsky [38].

If Y is a connected essential subsurface of S and $\xi(Y) \geq 1$ then the inclusion $Y \hookrightarrow S$ naturally induces an embedding $\mathcal{C}_0(Y) \hookrightarrow \mathcal{C}_0(S)$, whose image we identify with $\mathcal{C}_0(Y)$. If furthermore $\xi(Y) \geq 2$ then the embedding $\mathcal{C}_0(Y) \hookrightarrow \mathcal{C}_0(S)$ extends to a simplicial embedding $\mathcal{C}(Y) \hookrightarrow \mathcal{C}(S)$, whose image we identify with $\mathcal{C}(Y)$.

As in [39], we define $\mathcal{C}(Y)$ for an essential annulus $Y \subset S$ by considering the annular cover of S to which Y lifts homeomorphically, and which has a natural compactification to a compact annulus \mathcal{A}_Y (inherited from the usual

compactification of the universal cover \mathbb{H}^2). Define an *essential arc* in \mathcal{A}_Y to be an embedded arc with endpoints on different components of $\partial\mathcal{A}_Y$. We define $\mathcal{C}(Y)$ to be the graph whose vertices $\mathcal{C}_0(Y)$ are isotopy classes rel endpoints of essential arcs in \mathcal{A}_Y , with an edge for each pair of distinct vertices represented by essential arcs with disjoint interiors.

Note that if Y, Y' are isotopic essential annuli then $\mathcal{C}(Y), \mathcal{C}(Y')$ are *the same* complex.

Given an essential curve γ in S we let $\mathcal{C}(\gamma)$ denote $\mathcal{C}(Y)$ for any essential annulus $Y \subset S$ with core curve isotopic to γ .

The mapping class of the Dehn twist about γ acts naturally on $\mathcal{C}(\gamma)$ as follows: choose the twist τ_γ to be supported on an annulus neighborhood Y of γ , lift τ_γ through the covering map $\text{int}(\mathcal{A}_Y) \rightarrow S$ to a homeomorphism $\tilde{\tau}_\gamma: \mathcal{A}_Y \rightarrow \mathcal{A}_Y$ that is supported on the preimage of Y , and let this lift act on the essential arcs in \mathcal{A}_Y . The following properties of this action are easy; for details see [39].

Lemma 2.3. *For any essential curve γ in S and any vertex $v \in \mathcal{C}_0(\gamma)$ the orbit map $n \mapsto \tau_\gamma^n(v)$ is a quasi-isometry between \mathbb{Z} and $\mathcal{C}(\gamma)$. The action of the infinite cyclic group $\langle \tau_\gamma \rangle$ on $\mathcal{C}(\gamma)$ has a fundamental domain of diameter 2.*

We will need to use the main result of [38]:

Theorem 2.4. *For each surface Y with $\xi(Y) \geq 1$, the curve complex $\mathcal{C}(Y)$ is an infinite diameter δ -hyperbolic metric space, with respect to the simplicial metric.*

2.1.5. Markings and partial markings. We define markings and the marking complex for any connected surface Y with $\xi \geq 1$, as well as for any essential subsurface of S , including those which are disconnected and/or have some annulus component. We also define partial markings. (In [39], partial markings are called markings, and markings are called complete markings).

Suppose that Y is connected and $\xi(Y) \geq 1$. A *partial marking* $\mu = (\text{base}(\mu), t)$ on Y consists of a simplex $\text{base}(\mu)$ in $\mathcal{C}(Y)$ together with a choice of element $t(b) \in \mathcal{C}_0(b)$, which we call a *transversal*, for some (possibly none) of the vertices $b \in \text{base}(\mu)$; by convention we allow the empty set \emptyset as a partial marking of Y . If $t(b)$ is defined then we say that b is *marked (by μ)*, otherwise b is *unmarked (by μ)*. A *marking* (sometimes *full marking*) is a maximal partial marking, one for which $\text{base}(\mu)$ is a pants decomposition and every $b \in \text{base}(\mu)$ is marked. Given two partial markings $\mu = (\text{base}(\mu), t)$ and $\mu' = (\text{base}(\mu'), t')$, we write $\mu \subset \mu'$ to mean that $\text{base}(\mu) \subset \text{base}(\mu')$ and, for each $b \in \text{base}(\mu)$, b is marked by μ only if it is marked by μ' in which case $t(b) = t'(b)$; we also write $\mu \stackrel{i}{\equiv} \mu'$ to mean $\mu \subset \mu'$ and $\mu' \subset \mu$.

Consider now a connected essential subsurface Y of S . If $\xi(Y) \geq 1$ then we have defined above markings and partial markings of Y . If Y is an annulus

then a marking of Y is just a vertex of $\mathcal{C}(Y)$, and a *partial marking* of Y is either a marking of Y or \emptyset .

Finally, given an arbitrary essential subsurface Y of S , a marking of Y simply means a choice of marking on each component of Y .

2.1.6. The marking complex. We define the marking complex of any connected surface Y with $\xi \geq 1$, and of any essential subsurface of S .

First, given an essential subannulus Y of S , define the marking complex of Y to be $\mathcal{M}(Y) = \mathcal{C}(Y)$.

Suppose now that Y is connected and $\xi(Y) \geq 1$. The vertices of $\mathcal{M}(Y)$ are just the markings of Y . To define the edges we first need this notion: if b and c are overlapping essential curves, we denote $\pi_b(c) \in \mathcal{C}(b)$ to be the set of lifts of c to essential arcs in the annular cover associated to b . The diameter of this set is bounded in $\mathcal{C}(b)$, with a bound depending only on the topology of Y . The map π_b is an example of a subsurface projection map, defined below in a more general setting.

Edges in $\mathcal{M}(Y)$ correspond to *elementary moves* among markings on Y , which come in two flavors: *twist moves* and *flip moves*. To define them, consider a marking μ on Y and a curve $b \in \text{base}(\mu)$. A marking μ' is said to be obtained from μ by a *twist move* about b if $\text{base}(\mu) = \text{base}(\mu')$, μ, μ' have the same transversals to each curve in $\text{base}(\mu) \setminus \{b\} = \text{base}(\mu') \setminus \{b\}$, and the transversals $t(b)$ in μ and $t'(b)$ in μ' satisfy $d_{\mathcal{C}(b)}(t, t') \leq 2$. A marking $\mu' = (\text{base}(\mu'), t')$ is said to be obtained from μ by a *flip move* along b if there exists $b' \in \text{base}(\mu')$ such that $\text{base}(\mu) \setminus \{b\} = \text{base}(\mu') \setminus \{b'\}$, $\text{Fill}(b, b')$ is a one-holed torus or 4-holed sphere W such that $d_{\mathcal{C}(W)}(b, b') = 1$, $d_{\mathcal{C}(b)}(\pi_b(b'), t(b)) \leq 2$, and $d_{\mathcal{C}(b')}(\pi_{b'}(b), t'(b')) \leq 2$.

Finally, for any essential subsurface Y of S with components Y_1, \dots, Y_n , define $\mathcal{M}(Y)$ to be the 1-skeleton of the cartesian product $\mathcal{M}(Y_1) \times \dots \times \mathcal{M}(Y_n)$ with the usual CW-product structure. To put it another way, the vertices of $\mathcal{M}(Y)$ are the markings of Y , and there is an edge between two markings μ, μ' of Y if and only if μ, μ' are isotopic outside of a certain component Y_i , and the restrictions of μ, μ' to Y_i are connected by an edge of $\mathcal{M}(Y_i)$.

The marking complex is locally infinite because of the structure of transversals, but it is still quasi-isometric to a locally finite complex, as follows. Recall from [39] that a *clean marking* is a marking $\mu = (\text{base}(\mu), t)$ with the following properties: for each $b \in \text{base}(\mu)$, $t(b)$ is $\pi_b(c)$ where $c = c(b)$ is an essential curve in the component F of $Y \setminus (\text{base}(\mu) \setminus \{b\})$ containing b ; and the curves b and c have minimal nonzero intersection number in F . The complex of clean markings is a connected, $\mathcal{MCG}(Y)$ -invariant subcomplex of $\mathcal{M}(Y)$ whose vertices are the clean markings.

In fact the clean marking complex is what is usually referred to as the marking complex, see e.g., [3] and [7]. Because the full complex is more convenient for our purposes, we record this quasi-isometry:

Proposition 2.5. *The marking complex $\mathcal{M}(Y)$ is quasi-isometric to $\mathcal{MCG}(Y)$ and to the subcomplex of clean markings. More precisely, for each $\mu_0 \in \mathcal{M}(Y)$ the orbit map $\phi \mapsto \phi(\mu_0)$ is a quasi-isometry from $\mathcal{MCG}(Y)$ to $\mathcal{M}(Y)$.*

In particular $\mathcal{M}(Y)$ is connected, which may not have been obvious from the definition.

Sketch of proof. As noted in [39], the clean marking complex is locally finite, the action of $\mathcal{MCG}(Y)$ is properly discontinuous and cocompact, and so by the Milnor-Švarc lemma [42] [54] the orbit map is a quasi-isometry between $\mathcal{MCG}(Y)$ and the complex of clean markings. The inclusion of the complex of clean markings into $\mathcal{M}(Y)$ is an $\mathcal{MCG}(Y)$ -equivariant quasi-isometry, because for each marking there is a clean marking within a uniformly bounded distance C by Lemma 2.3, and for each edge of $\mathcal{M}(Y)$ connecting two markings μ_0, μ_1 , if μ'_0, μ'_1 are two clean markings within distance C of μ_0, μ_1 respectively then the distance between μ'_0, μ'_1 in the clean marking complex is uniformly bounded — this is checkable explicitly from the definition in [39] of the edges allowed between clean markings. \square

Remark on notation. In any context where $\mathcal{C}(Y)$ is under consideration the essential subsurface Y is assumed to be connected, whereas when $\mathcal{M}(Y)$ is being considered then Y can be disconnected.

2.1.7. Overlap. We define a symmetric binary relation of *overlap* for objects on S , denoted \pitchfork , as follows.

We have already defined overlap of two essential curve systems, in Section 2.1.2.

Overlap of an essential curve $\gamma \subset S$ and an essential subsurface $Y \subset S$, denoted $\gamma \pitchfork Y$ and $Y \pitchfork \gamma$, means that γ cannot be isotoped into the complement of Y . Equivalently, after isotoping γ to intersect ∂Y efficiently, the intersection $\gamma \cap Y$ is either a non-boundary-parallel essential curve in Y (the core of an annulus component is not allowed) or a nonempty pairwise disjoint union of *essential arcs in \bar{Y}* , each a properly embedded arc $\alpha \subset \bar{Y}$ that is not homotopic rel endpoints into ∂Y .

Given an essential curve system C and an essential subsurface Y , define $C \pitchfork Y$ to mean that there exists a component γ of C such that $\gamma \pitchfork Y$.

Overlap of two essential subsurfaces $X, Y \subset S$, denoted $X \pitchfork Y$, means that $\partial Y \pitchfork X$ and $\partial X \pitchfork Y$. Equivalently, after X, Y are isotoped so that $\partial X, \partial Y$ intersect efficiently, some component of $\bar{Y} \cap \partial X$ is an essential curve or arc in Y and some component of $\partial Y \cap \bar{X}$ is an essential curve or arc in X .

It is also equivalent to say that neither of Y or X is isotopic into the other and no matter how Y, X are chosen in their isotopy classes their intersection is nonempty. Notice for example that if $\partial X \cap \partial Y$ then $X \cap Y$, but the converse can fail.

Overlap of an essential subsurface Y and a partial marking μ , denoted $\mu \cap Y$ and $Y \cap \mu$, means that either $\text{base}(\mu) \cap Y$ or there exists $b \in \text{base}(\mu)$ such that b is marked by μ and some component of Y is an annulus neighborhood b .

2.1.8. Open subsurface and support. Given a partial marking $\mu = (\text{base}(\mu), t)$ on S , its *open subsurface*, denoted $\text{open}(\mu) = \text{open}_S(\mu)$, is defined to be the essential union of all subsurfaces Z such that $Z \not\cap \mu$. Equivalently, $\text{open}(\mu)$ is the largest essential subsurface which does not overlap μ . One can also describe it explicitly as the union of the components F of $S - \text{base}(\mu)$ such that $\xi(F) \geq 1$, and the annuli (if any) homotopic to the *unmarked* $b \in \text{base}(\mu)$. Note that each boundary curve of $\text{open}(\mu)$ is isotopic to a curve of $\text{base}(\mu)$. We usually drop the subscript S in the notation $\text{open}_S(\mu)$, unless we want to emphasize the surface in which the operation takes place, as in the proof of Lemma 3.3.

The *support* of a partial marking μ of S , denoted $\text{supp}(\mu) = \text{supp}_S(\mu)$, is defined to be $\text{open}(\mu)^c$, the essential complement of $\text{open}(\mu)$. The support of μ does not always behave as might at first be expected: for example, if no transversals are defined in μ then $\text{supp}(\mu) = \emptyset$, even if $\mu \neq \emptyset$.

We note several properties of $\text{open}(\mu)$ and $\text{supp}(\mu)$. Let $\mu = \mu^u \cup \mu^m$ where μ^u consists of the unmarked curves of $\text{base}(\mu)$ and μ^m consists of the marked curves and their transversals.

- Each component of $\partial \text{supp}(\mu)$ is isotopic to a component of $\text{base}(\mu)$.
- μ^m restricts to a (full) marking of each component of $\text{supp}(\mu)$.
- $\text{supp}(\mu)$ is characterized up to isotopy as the maximal essential subsurface of S with respect to the previous two properties.
- The curves μ^u are precisely the cores of the annulus components of $\text{open}(\mu)$.

2.1.9. Dehn twist flats. Given a pants decomposition μ on S , the set of markings whose base is μ is denoted $\mathcal{Q}(\mu) \subset \mathcal{M}(S)$, and is called the *Dehn twist flat* corresponding to μ . The terminology “flat” comes from the fact, proved in [7] Lemma 2.1, that $\mathcal{Q}(\mu)$ is a quasi-isometrically embedded copy of $\mathbf{R}^{\xi(S)}$; see also Proposition 3.1 (1) below, which generalizes this to allow μ to be any partial marking.

We state here without proof two other ways to view Dehn twist flats, in two other models of the quasi-isometric geometry of $\mathcal{MCG}(S)$: the word metric on $\mathcal{MCG}(S)$; and the thick part of Teichmüller space.

Under the quasi-isometry $\mathcal{M}(S) \leftrightarrow \mathcal{MCG}(S)$ of Proposition 2.5, a Dehn twist flat $\mathcal{Q}(\mu)$ corresponds uniformly to a left coset of a maximal-rank free-abelian subgroup generated by Dehn twists. To be precise, consider the finite set of $\mathcal{MCG}(S)$ orbits of pants decompositions, and choose one representative μ_1, \dots, μ_K from each orbit. There is a constant $A \geq 0$ such that for each pants decomposition $\mu = \Phi \cdot \mu_k$, the image of $\mathcal{Q}(\mu)$ in $\mathcal{MCG}(S)$ has Hausdorff distance $\leq A$ from the left coset $\Phi T(\mu_k)$ where $T(\mu_k)$ is the rank $\xi(S)$ free abelian group generated by Dehn twists about the components of μ_k .

Another quasi-isometric model of $\mathcal{MCG}(S)$ is the thick part of the Teichmüller space $\mathcal{T}(S)$. By Margulis' Lemma there is a constant $\epsilon_0 > 0$ independent of S such that in any hyperbolic structure on S , the set of simple closed geodesics of length $\leq \epsilon_0$ is pairwise disjoint, and every other simple closed geodesic has length $\geq 2\epsilon_0$. The thick part $\mathcal{T}_{\text{thick}}(S)$ is defined to be the set of hyperbolic structures whose shortest closed curve has length $\geq \epsilon_0$. The action of $\mathcal{MCG}(S)$ on $\mathcal{T}_{\text{thick}}(S)$ is properly discontinuous and cocompact and so, by the Švarc-Milnor lemma, there is a quasi-isometry $\mathcal{M}(S) \leftrightarrow \mathcal{T}_{\text{thick}}(S)$, with respect to any equivariant proper geodesic metric on $\mathcal{T}_{\text{thick}}(S)$. Applying Fenchel–Nielsen coordinates one sees that $\mathcal{T}_{\text{thick}}(S)$ is a *manifold-with-corners*, locally modelled on the closed orthant $\{x \in \mathbf{R}^{2\xi(S)} \mid x_i \geq 0 \text{ for all } i\}$. Each curve family C of cardinality j corresponds to a codimension- j facet $F(C)$ consisting of hyperbolic structures on which the curves of C are precisely the curves of length ϵ_0 . Note that for a pants decomposition μ , the Fenchel-Nielsen length coordinates in the facet $F(\mu)$ are all fixed to be ϵ_0 , and so the remaining coordinates are just the twists around the curves of μ , making $F(\mu)$ homeomorphic to $\mathbf{R}^{\xi(S)}$. Putting this all together, there is a constant $A \geq 0$ such that for each pants decomposition μ , the image of $\mathcal{Q}(\mu)$ in $\mathcal{T}_{\text{thick}}(S)$ has Hausdorff distance $\leq A$ from $F(\mu)$.

2.2. Projections

2.2.1. Subsurface projections. Following [39], [3], and [7], given a surface S and an essential subsurface Y we shall define several projection maps which take curves and markings in S to curves and markings in Y . When the target is the marking complex $\mathcal{M}(Y)$ we use $\pi_{\mathcal{M}(Y)}$ for the projection, and when the target is the curve complex $\mathcal{C}(Y)$ we use $\pi_{\mathcal{C}(Y)}$ or more briefly π_Y . Because these definitions require choices, for example choosing a vertex among a set of vertices, formally speaking we define the image of each map to be the set of all choices. However in all cases the map is *coarsely well-defined* (see Lemma 2.10), which means that the set of choices has uniformly bounded diameter. In practice we may sometimes abuse terminology and treat these maps as if the image of a point is a point. We may also abuse distance notation between the images of two points, confusing minimum distance, Hausdorff distance, and distance between any two representative elements of the images (see below under “Notation for various distances”), because those quantities

all differ by a uniformly bounded amount. Furthermore, these subsurface projection maps are coarsely Lipschitz (Lemma 2.11) and they behave well with respect to composition (Lemma 2.12).

Projecting curves to (sets of) curves. Suppose that $Y \Subset S$ is connected and not an annulus. If $\gamma \in \mathcal{C}(S)$, we define $\pi_{\mathcal{C}(Y)}(\gamma)$ to be the set of vertices of $\mathcal{C}(Y)$ obtained from essential arcs or curves of intersection of γ with Y , by the process of surgery along ∂Y . To be more precise, put γ in efficient position with respect to ∂Y , choose a component α of $\gamma \cap Y$, and consider a component of the boundary of a regular neighborhood of $\alpha \cup \partial Y$; the set of all essential curves in Y obtained in this way is $\pi_{\mathcal{C}(Y)}(\gamma)$.

If Y is an annulus, we let $\pi_{\mathcal{C}(Y)}(\gamma)$ be the set of vertices of $\mathcal{C}(Y)$ obtained as lifts of γ to the annular cover associated to Y . This operation was denoted π_b in the earlier discussion of marking complexes, where b is the core of Y .

Note in both cases that $\pi_{\mathcal{C}(Y)}(\gamma)$ is nonempty if and only if $\gamma \pitchfork Y$. See Lemma 2.10 for coarse well-definedness of $\pi_{\mathcal{C}(Y)}$.

Notation: we often write π_Y for any projection map whose target is $\mathcal{C}(Y)$. When the target needs to be emphasized we write $\pi_{\mathcal{C}(Y)}$.

The *bounded geodesic projection theorem* from [39] will be important for us:

Theorem 2.6. *Let Y be a connected essential subsurface of S and let g be a geodesic segment in $\mathcal{C}(S)$ such that $v \pitchfork Y$ for every vertex v of g . Then*

$$\text{diam}_{\mathcal{C}(Y)}(g) \leq B$$

where B is a constant depending only on $\xi(S)$.

Projecting (partial) markings to curves. We define a projection $\pi_{\mathcal{C}(Y)}(\mu) \subset \mathcal{C}(Y)$ for a partial marking μ in S as follows: When $\text{base}(\mu) \pitchfork Y$ we let $\pi_{\mathcal{C}(Y)}(\mu)$ be the union of $\pi_{\mathcal{C}(Y)}(b)$ over all $b \in \text{base}(\mu)$. When Y is an annulus neighborhood of a marked $b \in \text{base}(\mu)$ then we define $\pi_{\mathcal{C}(Y)}(\mu) = t(b)$. In all other cases $\pi_{\mathcal{C}(Y)}(\mu) = \emptyset$.

Projecting (partial) markings to (partial) markings. If μ is a partial marking in S and Y an essential subsurface, we will define a partial marking $\pi_{\mathcal{M}(Y)}(\mu)$ in Y . If μ is a marking of S then $\pi_{\mathcal{M}(Y)}(\mu)$ will be a marking of Y , so that we will obtain a coarse Lipschitz map $\pi_{\mathcal{M}(Y)}: \mathcal{M}(S) \rightarrow \mathcal{M}(Y)$ (see Lemma 2.11).

Write $\mu = (\text{base}(\mu), t)$. If $Y \not\pitchfork \mu$ then $\pi_{\mathcal{M}(Y)}(\mu) = \emptyset$. From now on we may assume $Y \pitchfork \mu$. If Y is disconnected we can view partial markings as tuples of partial markings in the components and define $\pi_{\mathcal{M}(Y)}$ componentwise. So we may also assume that Y is connected.

When Y is an annulus let $\pi_{\mathcal{M}(Y)}(\mu)$ denote any choice of element of $\pi_{\mathcal{C}(Y)}(\mu)$, recalling that $\mathcal{M}(Y) = \mathcal{C}(Y)$.

If $\xi(Y) \geq 1$, let b be any choice of element in $\pi_{\mathcal{C}(Y)}(\mu)$, let A be an annulus with core b , and let Y_b be the union of A with its essential complement in Y . Now inductively define

$$\pi_{\mathcal{M}(Y)}(\mu) = b \cup \pi_{\mathcal{M}(Y_b)}(\mu)$$

where the second term of the union is interpreted as a union over the components of Y_b . Note that, at the bottom of the induction, the annulus case provides transversals for all the base curves of $\pi_{\mathcal{M}(Y)}(\mu)$ that overlap μ .

There are choices at each stage of this construction, but when μ is a marking the final output is coarsely well-defined, as proved in [3]. See Lemma 2.10 for a statement.

For a partial marking μ , recall that $\text{open}(\mu)$ is the unique maximal essential subsurface Y such that $\mu \not\cap Y$, or equivalently such that $\pi_{\mathcal{M}(Y)}(\mu)$ is empty. The following lemma characterizes the (relative) open subsurface of the projection of a partial marking, as the maximal subsurface that doesn't overlap the marking:

Lemma 2.7. *If Y is an essential subsurface of S and μ a partial marking in S , then*

$$\text{open}_Y(\pi_{\mathcal{M}(Y)}(\mu)) = \bigcup \{Z \subseteq Y : Z \not\cap \mu\}.$$

Proof. Let $\mu' = \pi_{\mathcal{M}(Y)}(\mu)$. Every base curve in the inductive construction of μ' is either a base curve of μ itself, or an element of a subsurface projection of μ into some subsurface of Y . The induction terminates when the complementary subsurfaces of the base have no more overlap with μ , and when the base curves are either marked by μ' or are base curves of μ that have no transversals. It follows that $\text{open}_Y(\mu')$ does not overlap μ .

Conversely, let $Z \subseteq Y$ be a subsurface that does not overlap μ . If Z is an annulus around $b \in \text{base}(\mu)$ then b is unmarked by μ . Otherwise Z is disjoint from all vertices of $\pi_{\mathcal{C}(Y)}(\mu)$. Hence in the first step of the construction of μ' , the subsurface Z does not overlap the chosen base curve. Continuing by induction, Z does not overlap μ' . Hence $Z \subseteq \text{open}_Y(\mu')$.

We have shown that $\text{open}_Y(\mu')$ is among the set of subsurfaces of Y that don't overlap μ , and that every subsurface of Y that doesn't overlap μ is essentially contained in $\text{open}_Y(\mu')$. Hence the two sides are equal. \square

Notation for various distances: In a metric space M , when p, q are subsets of M we use $d_M(p, q)$ to denote *minimum distance*, meaning the infimum of the distance between an element of p and an element of q . We also use $d_{H,M}(p, q)$ to denote *Hausdorff distance*, the infimum of all $\epsilon \geq 0$ such that p is contained in the ϵ -neighborhood of q and q is contained in the ϵ -neighborhood of p . When the context is clear we abbreviate $d_{H,M}(p, q)$ to $d_H(p, q)$.

Given an essential subsurface $Y \subset S$, and any objects a and b in the domain of $\pi_{\mathcal{C}(Y)}$ or of $\pi_{\mathcal{M}(Y)}$ we denote

$$d_{\mathcal{C}(Y)}(a, b) = d_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(Y)}(a), \pi_{\mathcal{C}(Y)}(b))$$

and

$$d_{\mathcal{M}(Y)}(a, b) = d_{\mathcal{M}(Y)}(\pi_{\mathcal{M}(Y)}(a), \pi_{\mathcal{M}(Y)}(b))$$

as long as the right hand side makes sense, for example when a, b are full markings in the latter case. When the context is clear we often abbreviate $d_{\mathcal{C}(Y)}$ to d_Y .

2.2.2. Quasidistance formula. Knowing that $\mathcal{MCG}(S)$ is quasi-isometric to the marking complex $\mathcal{M}(S)$, we can study the asymptotic geometry of $\mathcal{MCG}(S)$ by having a useful quasidistance formula on $\mathcal{M}(S)$, which is provided by the following.

Given two numbers $d \geq 0$, $A \geq 0$, denote the truncated distance

$$\{d\}_A = \begin{cases} d & \text{if } d \geq A \\ 0 & \text{otherwise} \end{cases}$$

Given $r, s \geq 0$, $K \geq 1$, $C \geq 0$ we write $r \stackrel{K,C}{\approx} s$ to mean that $\frac{1}{K}s - C \leq r \leq Ks + C$. We also write $r \approx s$ to mean that $r \stackrel{K,C}{\approx} s$ for some *constants of approximation* K, C which are usually specified by the context, and we similarly write $r \lesssim s$ to mean $r \leq Ks + C$.

The following result is proved in [39].

Theorem 2.8. *There exists a constant $A_0 \geq 0$ depending only on the topology of S such that for each $A \geq A_0$, and for any $\mu, \mu' \in \mathcal{M}(S)$ we have the estimate*

$$d_{\mathcal{M}(S)}(\mu, \mu') \approx \sum_{Y \subset S} \{d_{\mathcal{C}(Y)}(\mu, \mu')\}_A$$

and the constants of approximation depend on A and on the topology of S .

The constant A in this theorem is usually called the *threshold constant*.

Remark: In summations and other expressions with index Y as in the above theorem, the convention will be that the index set consists of one representative Y in each isotopy class of *connected* essential subsurfaces, perhaps with some further restriction on the isotopy class; see for example Proposition 3.1(2) below.

There are two ways that Theorem 2.8 is applied. First, we can raise the threshold with impunity, which can make some terms drop out in a way that the remaining terms are more easily described. Second, if each term in the sum is replaced by another term differing by at most a uniform constant $C \geq 0$ then, after raising the threshold above $2C$, we may make the replacement at

the cost of a multiplicative factor of at most 2. As an example we have the following:

Corollary 2.9. *For any r there exists t such that for any $\mu, \nu \in \mathcal{M}(S)$, if $d_{\mathcal{C}(W)}(\mu, \nu) \leq r$ for all $W \subseteq S$, then $d_{\mathcal{M}(W)}(\mu, \nu) \leq t$.*

The proof is simply to raise the threshold above r , so that the right hand side of the quasidistance formula becomes 0. In fact a more careful look at the machinery of [39] yields $t = O(r^{\xi(S)})$.

2.2.3. Basic properties. We conclude this section with a brief summary of some of the basic properties of projections:

Lemma 2.10. *Subsurface projections are coarsely well defined:*

- *The diameter of $\pi_{\mathcal{C}(Y)}(x)$, where x is a curve or marking in S , is uniformly bounded.*
- *Similarly, the diameter of all possible choices in the construction of $\pi_{\mathcal{M}(Y)}(\mu)$, for $\mu \in \mathcal{M}(S)$, is uniformly bounded.*

Lemma 2.11. *Subsurface projections are coarsely Lipschitz in the following sense:*

- *If $x, y \in \mathcal{C}(S)$ with $d(x, y) = 1$ and both $x \pitchfork Y$ and $y \pitchfork Y$, then $\text{diam}(\pi_{\mathcal{C}(Y)}(x) \cup \pi_{\mathcal{C}(Y)}(y))$ is uniformly bounded.*
- *Similarly $d_{\mathcal{M}(Y)}(\mu, \nu)$ is uniformly bounded for any μ, ν in $\mathcal{M}(S)$ with $d(\mu, \nu) = 1$.*

In Lemmas 2.10 and 2.11, “uniformly bounded” means bounded by a constant depending only on $\xi(S)$. In fact, for the $\mathcal{C}(Y)$ bounds of these lemmas, there is a bound of 3 when Y is an annulus and of 2 when $\xi(Y) > 1$ — see [39] for details.

Lemma 2.12. *Subsurface projections for nested subsurfaces are coarsely composable in the following sense. Let $X, Y \subset S$ be essential subsurfaces such that $X \Subset Y$.*

- *For any $\gamma \in \mathcal{C}_0(S)$, the curve γ overlaps X if and only if γ overlaps Y and $\pi_Y(\gamma)$ contains at least one element α which overlaps X . In this case, π_X is coarsely equivalent to $\pi_X \circ \pi_Y$, meaning that*

$$\text{diam}_{\mathcal{C}(X)}(\pi_X(\gamma) \cup \pi_X(\pi_Y(\gamma)))$$

is uniformly bounded.

- *Similarly, if $\mu \in \mathcal{M}(S)$ then*

$$d_{\mathcal{M}(X)}(\pi_{\mathcal{M}(X)}(\mu), \pi_{\mathcal{M}(X)}(\pi_{\mathcal{M}(Y)}(\mu)))$$

is uniformly bounded.

We remark that for all three of these lemmas, the statements for curve complex projections are elementary from the definitions, and the statements for marking projections follow easily from the quasidistance formula.

2.3. Asymptotic cones

The asymptotic cone of a metric space is a way to encode the geometry of that space as seen from arbitrarily large distances. We will discuss this construction and the notation we will be using (see [11, 22] for further details).

To start, we recall that a (*non-principal*) *ultrafilter* is a finitely additive probability measure ω defined on the power set of the natural numbers, which takes values only 0 or 1, and for which every finite set has zero measure. If two sequences coincide on a set of indices whose ω -measure is equal to 1 then they they are said to be ω -*equivalent*, or to coincide ω -a.e. or ω -a.s. We will have a tendency, especially later in the paper, to abbreviate the terminology by speaking of ω -equivalent objects as being “equal” or “the same”.

The *ultraproduct* of a sequence of sets X_n is the quotient $\prod_n X_n / \sim$ of the cartesian product identifying two sequences $(x_n), (y_n)$ if they are ω -equivalent. We will often use the notation \bar{X} for the ultraproduct, and we use \bar{x} or $\langle x_n \rangle$ for the ω -equivalence class of a sequence (x_n) , also called its *ultraproduct equivalence class*.

In a topological space X , the *ultralimit* of a sequence of points (x_n) is x , denoted $x = \lim_\omega x_n$, if for every neighborhood U of x the set $\{n : x_n \in U\}$ has ω -measure equal to 1. Ultralimits are unique when they exist, moreover two ω -equivalent sequences in X have the same ultralimit. In this language the Bolzano-Weierstrass Theorem says that when X is compact every sequence has an ultralimit.

The *ultralimit* of a sequence of based metric spaces $(X_n, x_n, \text{dist}_n)$ is defined as follows: for $\bar{y}, \bar{z} \in \bar{X}$, we define $\text{dist}(\bar{y}, \bar{z}) = \lim_\omega \text{dist}_n(y_n, z_n)$, where the ultralimit is taken in the compact set $[0, \infty]$. We then let

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

where we define $\bar{y} \sim \bar{y}'$ if $\text{dist}(\bar{y}, \bar{y}') = 0$. Clearly dist makes this quotient into a metric space called the *ultralimit* of the X_n .

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

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

For the image of \bar{y} in the asymptotic cone, the *rescaled ultralimit* we use the notation either y_ω or \mathbf{y} . Given a sequence of subsets $A_n \subset X$ we use the

notation A_ω to denote the subset of $\text{Cone}_\omega(X, (x_n), s_n)$ consisting of all y_ω for sequences (y_n) such that $y_n \in A_n$ for all n .

The rescaling limit works equally well for a sequence $(X_n, x_n, \text{dist}_n)$ in which the X_n are not all the same metric space, and so we call $\lim_\omega(X_n, x_n, \text{dist}_n/s_n)$ the asymptotic cone of the sequence.

Convention: For the remainder of the paper we fix a non-principal ultrafilter ω . Usually also the scaling sequence $s_n \rightarrow \infty$ and the basepoint μ_0 for $\mathcal{M}(S)$ are implicitly fixed, and we write $\mathcal{M}_\omega = \mathcal{M}_\omega(S)$ to denote an asymptotic cone of $\mathcal{M}(S)$ with respect to these choices. Any choice of scaling sequence and basepoint will do, but in the last section we will need the flexibility of varying the choice of the scaling sequence (s_n) . Further, properties of linear or sub-linear growth of a non-negative function $f(n)$ are always taken with respect to the choice of ω and (s_n) , that is, we say $f(n)$ has *linear growth* if $0 < \lim_\omega \frac{f(n)}{s_n} < \infty$ and *sublinear growth* if $\lim_\omega \frac{f(n)}{s_n} = 0$.

Note that since $\mathcal{M}(S)$ is quasi-isometric to the group $\mathcal{MCG}(S)$ with any finitely generated word metric, and since the isometry group of $\mathcal{MCG}(S)$ acts transitively, the asymptotic cone is independent of the choice of basepoint.

Any essential connected subsurface W inherits a basepoint $\pi_{\mathcal{M}(W)}(\mu_0)$, canonical up to bounded error by Lemma 2.11, and we can use this to define the asymptotic cone $\mathcal{M}_\omega(W)$ of its marking complex $\mathcal{M}(W)$. For a disconnected subsurface $W = \sqcup_{i=1}^k W_i$ we have $\mathcal{M}(W) = \prod_{i=1}^k \mathcal{M}(W_i)$ and we may similarly construct $\mathcal{M}_\omega(W)$ which can be identified with $\prod_{i=1}^k \mathcal{M}_\omega(W_i)$ — this follows from the general fact that the process of taking asymptotic cones commutes with finite products. Note that for an annulus A we have defined $\mathcal{M}(A) = \mathcal{C}(A)$ which is quasi-isometric to \mathbb{Z} , so $\mathcal{M}_\omega(A)$ is bilipschitz equivalent (isometric) to \mathbb{R} .

For a sequence (W_n) of subsurfaces we can similarly form the ultraproduct of $(\mathcal{M}(W_n))$, which we denote by $\mathcal{M}(\overline{W})$, where $\overline{W} = \langle W_n \rangle$. The asymptotic cone of this sequence with the inherited basepoints is denoted $\mathcal{M}_\omega(\overline{W})$. We also let \overline{S} denote the constant sequence (S, S, \dots) so that $\mathcal{M}(\overline{S})$ is the ultraproduct of $(\mathcal{M}(S), \dots)$ and $\mathcal{M}_\omega(\overline{S})$ is the same as $\mathcal{M}_\omega(S)$.

Any sequence in a finite set A is ω -a.e. constant: given $(a_n \in A)$ there is a unique $a \in A$ such that $\omega(\{n : a_n = a\}) = 1$. It follows that the ultralimit \bar{a} of this sequence is naturally identified with a . For example if (W_n) is a sequence of essential subsurfaces of S then the topological type of W_n is ω -a.e. constant, so we call this the topological type of \overline{W} . Similarly the topological type of the pair (S, W_n) is ω -a.e. constant. We can moreover interpret expressions like $\overline{U} \subset \overline{W}$ to mean $U_n \subset W_n$ for ω -a.e. n , and so on. Note that $\mathcal{M}_\omega(\overline{W})$ can be identified with $\mathcal{M}_\omega(W)$, where W is a surface homeomorphic to W_n for ω -a.e. n ; this identification is not natural, however, because it depends up to ω -equivalence on choosing a homeomorphism between W and each W_n .

For two sequences of sets (A_n) and (B_n) and a sequence of functions $f_n: A_n \rightarrow B_n$, passing to the ultralimit gives rise to a single function $\bar{f}: \bar{A} \rightarrow \bar{B}$, and \bar{f} determines f_n up to ω -equivalence in the ultraproduct of the sequence $(B_n^{A_n})$. We can therefore think of a sequence of projection maps $\pi_{\mathcal{M}(W_n)}: \mathcal{M}(S) \rightarrow \mathcal{M}(W_n)$ as a single map

$$\pi_{\mathcal{M}(\bar{W})}: \mathcal{M}(\bar{S}) \rightarrow \mathcal{M}(\bar{W})$$

Upon rescaling, this map descends to a map of the asymptotic cones, the *rescaled ultralimit of the projection maps*

$$\pi_{\mathcal{M}_\omega(\bar{W})}: \mathcal{M}_\omega(S) \rightarrow \mathcal{M}_\omega(\bar{W}).$$

Note by Lemma 2.11 that this is a Lipschitz map. This sort of notation will be used heavily in Section 8.

3. Product regions and cubes

In this section we define and study subsets of the marking complex obtained by holding fixed one part of the surface and varying the rest. These will be called *product regions*, because of the product structure described in Lemma 3.1. A special case of a product region is a *Dehn twist flat*. We will also consider particular subsets of product regions called *cubes*, which are in fact naturally quasi-isometrically parametrized by cubes in Euclidean space.

The metric relation between a pair of product regions or cubes will be described in terms of *junctions*, the part of each set which comes closest to the other. Later, when we pass to the asymptotic cone, these junctions will become intersections, and will be important in understanding the structure of orthants in the cone.

3.1. Product regions

For each partial marking μ of S , define its associated *product region*

$$\mathcal{Q}(\mu) = \{\mu' \in \mathcal{M}(S) : \mu \subset \mu'\},$$

that is, the set of all (full) markings that extend μ .

For example, if μ is a curve system with no transversals then $\mathcal{Q}(\mu)$ is the collection of all markings whose base curves contain μ . In this case $\mathcal{Q}(\mu)$ is uniformly quasi-isometric to a left coset of a subgroup stabilizing a certain curve system, just as Dehn twist flats are uniformly quasi-isometric to left cosets of certain maximal rank Dehn twist subgroups, as explained in Section 2.1.9.

In particular, if μ is a pants decomposition with no transversals then $\mathcal{Q}(\mu)$ is a Dehn twist flat, and so our notation is consistent with the notation for Dehn twist flats introduced in Section 2.1.1.

Product structure on $\mathcal{Q}(\mu)$. An element $\mu' \in \mathcal{Q}(\mu)$ is specified by choosing, for each component Y of $\text{open}(\mu)$, a marking on Y which we denote $\mu' \upharpoonright Y$, and which may be identified with $\pi_{\mathcal{M}(Y)}(\mu')$. Hence $\mathcal{Q}(\mu)$ is naturally identified with $\mathcal{M}(\text{open}(\mu))$, which is a product

$$\prod_{Y \in |\text{open}(\mu)|} \mathcal{M}(Y)$$

where $|Z|$ denotes the set of components of Z .

With this in mind, given two partial markings μ, ν on S , define the *extension of μ by ν* to be the partial marking given by:

$$\mu \upharpoonright \nu = \mu \cup \pi_{\mathcal{M}(\text{open}(\mu))}(\nu).$$

If $\pi_{\mathcal{M}(\text{open}(\mu))}(\nu)$ is a full marking in $\text{open}(\mu)$ then $\mu \upharpoonright \nu \in \mathcal{Q}(\mu)$. This applies for example when ν itself is a full marking on S , in which case item (3) of the following proposition tells us that $\mu \upharpoonright \nu$ is (coarsely) the closest point projection of ν onto $\mathcal{Q}(\mu)$.

The following result generalizes the case considered in [7] where μ was a curve system without transversals.

Proposition 3.1. *Let $\xi(S) \geq 1$ and let μ be a partial marking of S .*

(1) *The map*

$$\mathcal{M}(\text{open}(\mu)) = \prod_{Y \in |\text{open}(\mu)|} \mathcal{M}(Y) \quad \rightarrow \quad \mathcal{M}(S)$$

induced by the identification with $\mathcal{Q}(\mu)$ is a quasi-isometric embedding, with constants depending only on the topology of S .

(2) *There is a constant A_0 depending only on the topology of S such that for each $A \geq A_0$, and for each $x \in \mathcal{M}(S)$, the minimum distance $d_{\mathcal{M}(S)}(x, \mathcal{Q}(\mu))$ from x to $\mathcal{Q}(\mu)$ in $\mathcal{M}(S)$ is estimated by*

$$d_{\mathcal{M}(S)}(x, \mathcal{Q}(\mu)) \approx \sum_{Y \uparrow \mu} \left\{ \left\{ d_{\mathcal{C}(Y)}(x, \mu) \right\} \right\}_A$$

where the constants of approximation depend on A and on the topology of S .

(3) *Moreover, again with uniform constants, for each $x \in \mathcal{M}(S)$*

$$d_{\mathcal{M}(S)}(x, \mathcal{Q}(\mu)) \approx d_{\mathcal{M}(S)}(x, \mu \upharpoonright x).$$

As a consequence of (1), combined with the ordinary quasidistance formula for components of $S \setminus \text{base}(\mu)$, if we let Δ be the subset of $\text{base}(\mu)$ consisting of those curves for which no transversal is defined then we have a quasi-isometry

$$\mathcal{Q}(\mu) \approx \mathbb{Z}^\Delta \times \prod_{\substack{Y \in |\text{open}(\mu)| \\ \xi(Y) \geq 1}} \mathcal{M}(Y)$$

As another example of (1), given an essential subsurface W , we have

$$\mathcal{Q}(\partial W) \approx \mathcal{M}(W) \times \mathcal{M}(W^c)$$

As an example of (2), if μ is a full marking then $\mathcal{Q}(\mu) = \{\mu\}$, $\text{open}(\mu) = \emptyset$, $\text{supp}(\mu) = S$, and (2) is just the ordinary quasidistance formula to μ .

Proof of Proposition 3.1. Suppose that F is an essential subsurface of S that is *not* essentially contained in $\text{open}(\mu)$, and recall that $\mu \pitchfork F$. It follows that for any $\mu', \mu'' \in \mathcal{Q}(\mu)$ each of $\pi_{\mathcal{C}(F)}(\mu')$, $\pi_{\mathcal{C}(F)}(\mu'')$ is within a uniformly bounded distance of $\pi_{\mathcal{C}(F)}(\mu)$ — this is a consequence of coarse well-definedness of $\pi_{\mathcal{C}(F)}$, Lemma 2.10. We can therefore make the $\pi_{\mathcal{C}(F)}$ term drop out of the quasidistance formula for $d(\mu', \mu'')$ by raising the threshold. The remaining terms can be collected to give the sums of the quasidistance formulas for the projections of μ', μ'' to the components of $\text{open}(\mu)$, and we obtain (1) as an immediate consequence.

Now consider a connected essential subsurface $Y \subset S$ such that $Y \pitchfork \mu$. For each $\mu' \in \mathcal{Q}(\mu)$, from coarse well-definedness (Lemma 2.10) it follows that $\pi_{\mathcal{C}(Y)}(\mu)$ and $\pi_{\mathcal{C}(Y)}(\mu')$ are within uniformly bounded distance of each other. Therefore, for each $x \in \mathcal{M}(S)$ we have $d_{\mathcal{C}(Y)}(x, \mu) \stackrel{1,C}{\approx} d_{\mathcal{C}(Y)}(x, \mu')$ where C depends only on the topology of S . The terms $d_{\mathcal{C}(Y)}(x, \mu)$ comprise the right side of (2) and the terms $d_{\mathcal{C}(Y)}(x, \mu')$ are among the terms in the quasidistance formula for $d_{\mathcal{M}(S)}(x, \mu')$, so after raising the threshold by $2C$ we obtain the \gtrsim direction of (2).

Next, let $\mu' = \mu \rfloor x$. For each Y , if $\mu \not\pitchfork Y$ then $d_{\mathcal{C}(Y)}(x, \mu')$ is bounded above by a constant depending only on the topology of S , because by Lemma 2.12 $\pi_{\mathcal{C}(Y)}(\mu')$ is within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(\text{open}(\mu))}(\mu'))$, which equals $\pi_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(\text{open}(\mu))}(x))$, which is itself within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(x)$. We can then raise the threshold above this constant, so that all of these terms drop out of the quasidistance formula for $d_{\mathcal{M}(S)}(x, \mu')$, leaving only the terms where $\mu \pitchfork Y$. This proves the \lesssim direction of (2), as well as (3). \square

Junctures. Let \mathcal{F} be a family of subsets of a metric space \mathcal{M} . We say that \mathcal{F} *has junctures* if the following holds: for any X, Y in \mathcal{F} there exist $E(X, Y) \subset X$ and $E(Y, X) \subset Y$, both members of \mathcal{F} as well, such that:

- (1) The Hausdorff distance $d_H(E(X, Y), E(Y, X))$ is finite.
- (2) If $x \in X, y \in Y$ then

$$d(x, y) \gtrsim d(x, E(X, Y)) + d_H(E(X, Y), E(Y, X)) + d(E(Y, X), y).$$

with constants of approximation being uniform over the family \mathcal{F} .

The sets $E(X, Y)$ and $E(Y, X)$ are called the *junctures* of X and Y . Note that the junctures are “parallel” in the sense not just of the bound on Hausdorff

distance, but the inequality the other way $d(x, y) \gtrsim d_H(E(X, Y), E(Y, X))$ which by (2) holds for all $x \in E(X, Y)$ and $y \in E(Y, X)$.

The motivating example of a family having junctures is the family of geodesics (finite or infinite) in a δ -hyperbolic space. Here the implicit constants depend on δ . This example has the feature that for any X, Y either $E(X, Y)$ and $E(Y, X)$ are points, or $d_H(E(X, Y), E(Y, X)) \approx 0$. Junctures for the family $\mathcal{Q}(\mu)$ will not have this feature.

Junctures are examples of the general concept of “coarse intersection”, as we now explain. In a metric space \mathcal{M} , given subsets $A, B \subset \mathcal{M}$, we say that the *coarse intersection of A, B is well-defined* if there exists $R \geq 0$ such that any two elements of the collection of subsets $\{\mathcal{N}_r(A) \cap \mathcal{N}_r(B) \mid r \geq R\}$ have finite Hausdorff distance. For any subset $C \subset \mathcal{M}$ which has finite Hausdorff distance from any one of these sets, we also say that the coarse intersection of A and B is *represented by* the set C . We can define coarse intersection of a finite number of sets in the same way.

If the family \mathcal{F} of subsets of \mathcal{M} has junctures then the junctures are representatives of coarse intersections. To prove this, choose $K \geq 1$, $C \geq 0$ to be constants of approximation as in the definition of junctures and let $L = 1 + \frac{2}{K}$. For any $X, Y \in \mathcal{F}$, let $D = d_H(E(X, Y), E(Y, X))$, and suppose that $r \geq LD$. We have $E(X, Y), E(Y, X) \subset \mathcal{N}_r(X) \cap \mathcal{N}_r(Y)$ because $r \geq D$. If $q \in \mathcal{N}_r(X) \cap \mathcal{N}_r(Y)$ then we may choose $x \in X$, $y \in Y$ such that $d(x, q), d(q, y) \leq r$, and we have

$$\begin{aligned} 2r &\geq d(x, y) \\ &\geq K(d(x, E(X, Y)) + D + d(E(Y, X), y)) + C \\ \frac{2r}{K} - D &\geq d(x, E(X, Y)) \\ d(q, E(X, Y)) &\leq d(q, x) + d(x, E(X, Y)) \\ &\leq Lr - D \end{aligned}$$

and similarly for $d(q, E(Y, X))$ (notice that $Lr - D \geq 0$ because $r \geq LD \geq \frac{D}{L}$). We record this as:

Lemma 3.2. *If \mathcal{M} is a metric space and \mathcal{F} is a family of subsets of \mathcal{M} having junctures then the coarse intersection of any two elements of \mathcal{F} is well-defined and is represented by their junctures. More precisely, there exists $L \geq 1$, depending only on the constants of approximation, such that for all $X, Y \in \mathcal{F}$, letting $D = d_H(E(X, Y), E(Y, X))$, for all $r \geq LD$ we have*

$$E(X, Y) \cup E(Y, X) \subset \mathcal{N}_r(X) \cap \mathcal{N}_r(Y) \subset \mathcal{N}_{Lr-D}(E(X, Y)) \cap \mathcal{N}_{Lr-D}(E(Y, X))$$

□

Junctures for the family $\{\mathcal{Q}(\mu)\}$.

Lemma 3.3. *The family of subsets $\mathcal{Q}(\mu) \subset \mathcal{M}(S)$ has junctures: for any partial markings μ_0, μ_1 of S , the junctures for $\mathcal{Q}(\mu_0)$ and $\mathcal{Q}(\mu_1)$ are*

$$\begin{aligned} E_{01} &= E(\mu_0, \mu_1) = \mathcal{Q}(\mu_0 \rfloor \mu_1) \subset \mathcal{Q}(\mu_0) \\ E_{10} &= E(\mu_1, \mu_0) = \mathcal{Q}(\mu_1 \rfloor \mu_0) \subset \mathcal{Q}(\mu_1) \end{aligned}$$

More precisely we have:

- (1) *The subsurfaces $\text{open}(\mu_0 \rfloor \mu_1)$, $\text{open}(\mu_1 \rfloor \mu_0)$, and $\text{open}(\mu_0) \circledast \text{open}(\mu_1)$ are all isotopic. Let $\text{open}(\mu_0, \mu_1)$ denote a surface in this isotopy class.*
- (2) *The Hausdorff distance $d_H(E_{01}, E_{10})$ in $\mathcal{M}(S)$ is estimated by*

$$d_H(E_{01}, E_{10}) \approx d_{\mathcal{M}(\text{supp}(\mu_0, \mu_1))}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0)$$

where we denote $\text{supp}(\mu_0, \mu_1) = \text{open}(\mu_0, \mu_1)^c$.

- (3) *For $x_i \in \mathcal{Q}(\mu_i)$ we have*

$$\begin{aligned} d_{\mathcal{M}(S)}(x_0, x_1) &\approx d_{\mathcal{M}(S)}(x_0, E_{01}) + d_{\mathcal{M}(S)}(x_1, E_{10}) \\ &\quad + d_{\mathcal{M}(\text{supp}(\mu_0, \mu_1))}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0) + d_{\mathcal{M}(\text{open}(\mu_0, \mu_1))}(x_0, x_1) \end{aligned}$$

Proof. Part (1) follows (by symmetry) from the general identity

$$\text{open}(\mu \rfloor \nu) \stackrel{i}{=} \text{open}(\mu) \circledast \text{open}(\nu) \tag{3.1}$$

for any two partial markings. The subsurface $Z = \text{open}(\mu \rfloor \nu)$ is the maximal essential subsurface that does not overlap $\mu \rfloor \nu$, hence Z does not overlap μ so $Z \subsetneq \text{open}(\mu)$. Also, Z does not overlap $\pi_{\mathcal{M}(\text{open}(\mu))}(\nu)$, so by Lemma 2.7, Z does not overlap ν and therefore $Z \subsetneq \text{open}(\nu)$. We conclude that $\text{open}(\mu \rfloor \nu) \subsetneq \text{open}(\mu) \circledast \text{open}(\nu)$.

Conversely letting $X = \text{open}(\mu) \circledast \text{open}(\nu)$, from Lemma 2.7 it follows that $X \subset \text{open}_{\text{open}(\mu)}(\pi_{\mathcal{M}(\text{open}(\mu))}(\nu))$ and so X does not overlap $\mu \rfloor \nu$. We conclude that $\text{open}(\mu) \circledast \text{open}(\nu) \subsetneq \text{open}(\mu \rfloor \nu)$, and (3.1) follows.

The proofs of (2) and (3) will be applications of the quasidistance formula. Note that, now that we know that $\mu_0 \rfloor \mu_1$ and $\mu_1 \rfloor \mu_0$ have the same support surface $\text{supp}(\mu_0, \mu_1)$, the distance between these markings in the complex $\mathcal{M}(\text{supp}(\mu_0, \mu_1))$ is defined, so that (2) makes sense.

To obtain the inequality \gtrsim in (2), consider any term in the quasidistance formula for $d_{\mathcal{M}(\text{supp}(\mu_0, \mu_1))}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0)$, indexed by $Y \subset \text{supp}(\mu_0, \mu_1)$. This term is within uniform distance of $d_{\mathcal{C}(Y)}(x, x')$ for any $x \in E_{01}$ and $x' \in E_{10}$, since x contains $\mu_0 \rfloor \mu_1$ and x' contains $\mu_1 \rfloor \mu_0$. Hence this term contributes to a lower bound for the quasidistance formula for $d_{\mathcal{M}(S)}(x, x')$. As before, raising the threshold eliminates the effect of the additive errors.

To prove the inequality \lesssim of (2), note that each $x \in E_{ij}$ contains $\mu_i \rfloor \mu_j$. If we replace this part of x by $\mu_j \rfloor \mu_i$, holding the part $x \upharpoonright \text{open}(\mu, \nu)$ constant, we obtain a point $x' \in E_{ji}$. For any $Y \subset S$ which does *not* index the quasidistance formula for the right hand side of (2), the term $d_{\mathcal{C}(Y)}(x, x')$ is uniformly bounded, as we see by enumerating cases. If Y essentially intersects $\partial \text{open}(\mu, \nu)$ then

$\pi_{\mathcal{C}(Y)}(x)$ and $\pi_{\mathcal{C}(Y)}(x')$ are uniformly close to $\pi_{\mathcal{C}(Y)}(\partial \text{open}(\mu, \nu))$. If Y does not essentially intersect $\partial \text{open}(\mu, \nu)$ then Y is isotopic into $\text{open}(\mu, \nu)$ or its complement. If Y is isotopic into the complement of $\text{open}(\mu, \nu)$ then either Y is an annulus component of $\text{open}(\mu, \nu)$ or Y is an index for the right hand side of (2). If Y is an annulus component of $\text{open}(\mu, \nu)$, or if Y is an essential subsurface of a component of $\text{open}(\mu, \nu)$, then $\pi_{\mathcal{C}(Y)}(x)$ and $\pi_{\mathcal{C}(Y)}(x')$ are within uniformly bounded distance of the projection of $x \mid \text{open}(\mu, \nu) = x' \mid \text{open}(\mu, \nu)$. This exhausts all cases. By raising the threshold, it follows that $d_{\mathcal{M}(S)}(x, x')$ reduces to the right hand side of (2), proving the inequality \lesssim .

To prove (3), let $x_i \in \mathcal{Q}(\mu_i)$ for $i = 0, 1$. We just need to check that each term in the quasidistance formula for $d_{\mathcal{M}(S)}(x_0, x_1)$ contributes to one of the four summands on the right hand side.

The first summand $d_{\mathcal{M}(S)}(x_0, E_{01})$, by Proposition 3.1(2), is estimated by

$$d_{\mathcal{M}(S)}(x_0, E_{01}) \approx \sum_{Y \pitchfork \mu_0 \rfloor \mu_1} \left\{ \left\{ d_{\mathcal{C}(Y)}(x_0, \mu_0 \rfloor \mu_1) \right\} \right\}_A$$

However, note that if $Y \pitchfork \mu_0$ then $d_{\mathcal{C}(Y)}(x_0, \mu_0 \rfloor \mu_1) \approx 1$ since both markings contain μ_0 . On the other hand, if $Y \pitchfork \mu_0 \rfloor \mu_1$ and $Y \not\pitchfork \mu_0$ then $Y \textcircled{\subset} \text{open}(\mu_0)$ and $Y \pitchfork \mu_1$ by Lemma 2.7; and the converse holds as well. Therefore by raising the threshold A we get

$$d_{\mathcal{M}(S)}(x_0, E_{01}) \approx \sum_{\substack{Y \textcircled{\subset} \text{open}(\mu_0) \\ Y \pitchfork \mu_1}} \left\{ \left\{ d_{\mathcal{C}(Y)}(x_0, \mu_0 \rfloor \mu_1) \right\} \right\}_A$$

Each term $d_{\mathcal{C}(Y)}(x_0, \mu_0 \rfloor \mu_1)$ is within a uniformly bounded distance of $d_{\mathcal{C}(Y)}(x_0, \mu_1)$, by Proposition 2.12; it follows that by raising the threshold above twice this bound, at the cost of another multiplicative factor of 2, we get

$$d_{\mathcal{M}(S)}(x_0, E_{01}) \approx \sum_{\substack{Y \textcircled{\subset} \text{open}(\mu_0) \\ Y \pitchfork \mu_1}} \left\{ \left\{ d_{\mathcal{C}(Y)}(x_0, \mu_1) \right\} \right\}_A \quad (3.2)$$

We obtain a similar expression for the second summand:

$$d_{\mathcal{M}(S)}(x_1, E_{10}) \approx \sum_{\substack{Y \textcircled{\subset} \text{open}(\mu_1) \\ Y \pitchfork \mu_0}} \left\{ \left\{ d_{\mathcal{C}(Y)}(x_1, \mu_0) \right\} \right\}_A \quad (3.3)$$

The third summand is given by

$$d_{\mathcal{M}(\text{supp}(\mu_0, \mu_1))}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0) \approx \sum_{Y \textcircled{\subset} \text{supp}(\mu_0, \mu_1)} \left\{ \left\{ d_{\mathcal{C}(Y)}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0) \right\} \right\}_A$$

If $Y \textcircled{\subset} \text{supp}(\mu_0, \mu_1)$ and $Y \not\pitchfork \mu_0$ then $Y \pitchfork \mu_1$ and both $\pi_{\mathcal{C}(Y)}(\mu_0 \rfloor \mu_1)$ and $\pi_{\mathcal{C}(Y)}(\mu_1 \rfloor \mu_0)$ are within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(\mu_1)$, so these terms may be dropped by raising the threshold. Similarly, if $Y \textcircled{\subset} \text{supp}(\mu_0, \mu_1)$ and $Y \not\pitchfork \mu_1$ then $Y \pitchfork \mu_0$ and these terms may be dropped.

If $Y \not\subseteq \text{supp}(\mu_0, \mu_1)$, and if $Y \pitchfork \mu_0$ and $Y \pitchfork \mu_1$, then $Y \pitchfork \partial \text{supp}(\mu_0, \mu_1)$; for if not, Y would be isotopic to the complement of $\text{supp}(\mu_0, \mu_1)$ and so to overlap μ_0 and μ_1 , Y would have to be an annulus isotopic to a boundary curve of $\text{supp}(\mu_0, \mu_1)$, which is marked by both μ_1 and μ_0 . But in this case, by definition the annulus would be a component of $\text{supp}(\mu_0, \mu_1)$ so $Y \subseteq \text{supp}(\mu_0, \mu_1)$ after all. In this situation both $\pi_{\mathcal{C}(Y)}(\mu_0 \rfloor \mu_1)$ and $\pi_{\mathcal{C}(Y)}(\mu_1 \rfloor \mu_0)$ are within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(\partial \text{supp}(\mu_0, \mu_1))$, so that $d_{\mathcal{C}(Y)}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0)$ is uniformly bounded. Thus although these terms do not appear in the sum, by raising the threshold we may formally put them into it with only a bounded change to the estimate.

At this stage, the sum is indexed by the set of all $Y \subset S$ such that $Y \pitchfork \mu_0$ and $Y \pitchfork \mu_1$. For such Y , $\pi_{\mathcal{C}(Y)}(\mu_0 \rfloor \mu_1)$ is within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(\mu_0)$, and $\pi_{\mathcal{C}(Y)}(\mu_1 \rfloor \mu_0)$ is within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(\mu_1)$, and so the $d_{\mathcal{C}(Y)}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0)$ is approximated within a uniform additive error by $d_{\mathcal{C}(Y)}(\mu_0, \mu_1)$. By raising the threshold above twice this error we obtain

$$d_{\mathcal{M}(\text{supp}(\mu_0, \mu_1))}(\mu_0 \rfloor \mu_1, \mu_1 \rfloor \mu_0) \approx \sum_{\substack{Y \pitchfork \mu_0 \\ Y \pitchfork \mu_1}} \{ \{ d_{\mathcal{C}(Y)}(\mu_0, \mu_1) \} \}_A \quad (3.4)$$

The fourth summand is, by the quasidistance formula in $\text{open}(\mu_0, \mu_1)$, approximated by

$$d_{\mathcal{M}(\text{open}(\mu_0, \mu_1))}(x_0, x_1) \approx \sum_{Y \subseteq \text{open}(\mu_0, \mu_1)} \{ \{ d_{\mathcal{C}(Y)}(x_0, x_1) \} \}_A \quad (3.5)$$

Now putting these four sums (3.2), (3.3), (3.4), (3.5) together, and recalling that $Y \subseteq \text{open}(\mu_i)$ if and only if $Y \not\pitchfork \mu_i$, it follows that each $Y \subset S$ appears in exactly one of these four sums. Moreover, whenever μ_i appears it can be replaced by x_i with a bounded additive change in the term. Raising the threshold above twice the value of this change, we see that the sum is approximated by the quasidistance formula for $d_{\mathcal{M}(S)}(x_0, x_1)$. \square

3.2. Dimension and nonseparation in the asymptotic cone

We shall apply Proposition 3.1 and the results of [7] to compute the dimension of a product region in the asymptotic cone, and combined with Alexander duality we will obtain nonseparation results.

Recall that the *topological dimension* of a topological space X is the least n such that every open cover of X can be refined to an open cover \mathbf{U} having the property that for any subset of \mathbf{U} whose members have nonempty intersection in X , the cardinality of the subset is $\leq n + 1$. The *locally compact topological dimension* is the least n such that each locally compact subset of X has topological dimension $\leq n$. It follows immediately that the Čech cohomology of each locally compact subset of X vanishes in dimensions above n .

Throughout this paper we usually use the word *dimension* to refer to locally compact topological dimension. Also, given a subset $A \subset X$, we use the term *codimension* to refer to the locally compact dimension of X minus the locally compact dimension of A . Proposition 2.2 is useful in some contexts for computing codimension.

The following theorem is the main result of [7]:

Theorem 3.4. *For each connected, finite type surface F with $\xi(F) \geq 1$, the locally compact topological dimension of $\mathcal{M}_\omega(F)$ is $\xi(F)$.*

Consider a sequence of partial markings (μ_n) . Since an essential subsurface of S can have only finitely many topological types, the type of $\text{open}(\mu_n)$ is ω -a.e. constant, and so the number $r(\text{open}(\mu_n))$ is ω -a.e. constant, a number we denote $r(\text{open}(\bar{\mu}))$.

By combining Theorem 3.4 with Proposition 3.1, in particular the coarse cartesian product formula for $\mathcal{Q}(\mu_n)$ given after the statement of the proposition, plus the fact that finite cartesian products commute with rescaled ultralimits, we immediately obtain:

Lemma 3.5. *If (μ_n) is a sequence of partial markings such that the rescaled ultralimit $\mathcal{Q}_\omega(\bar{\mu})$ of $(\mathcal{Q}(\mu_n))$ is nonempty in $\mathcal{M}_\omega(S)$, then $\mathcal{Q}_\omega(\bar{\mu})$ has locally compact topological dimension $r(\text{open}(\bar{\mu}))$. \square*

Given a topological space X and two subsets A, B , we say that A *separates* B in X if B has nonempty intersection with at least two components of $X - A$; in particular, $B - A$ is disconnected.

Lemma 3.6. *Let (μ_n) be as in Lemma 3.5. If $E \subset \mathcal{M}_\omega(S)$ is a connected oriented manifold of dimension $D \geq r(\text{open}(\bar{\mu})) + 2$ then $\mathcal{Q}_\omega(\bar{\mu})$ does not separate E in X .*

Proof. Since $\mathcal{Q}_\omega(\bar{\mu})$ is closed and E is locally compact, the set $E \cap \mathcal{Q}_\omega(\bar{\mu})$ is locally compact, and therefore by Lemma 3.5 is of topological dimension $\leq r(\text{open}(\bar{\mu})) < D - 1$. By Alexander Duality,

$$\tilde{H}_0(E - \mathcal{Q}_\omega(\bar{\mu})) \approx \check{H}^{D-1}(E \cap \mathcal{Q}_\omega(\bar{\mu}))$$

using reduced homology on the left hand side and Čech cohomology on the right hand side, but the right hand side is trivial. \square

3.3. Cubes and their junctures

Cubes are subsets of $\mathcal{M}(S)$ obtained by choosing an essential subsurface of S whose components all satisfy $\xi \leq 1$, marking the complement of that subsurface, and choosing a geodesic in the marking complex of each component of the subsurface; the cube is parameterized by the product of the chosen geodesics. A special case is a “Dehn twist k -flat” in which S has k components, each one

an annulus, and the chosen geodesic in the marking complex of each annulus is bi-infinite; see Section 3.5 for more details on Dehn twist k -flats. In this language, a “Dehn twist flat” as previously defined is the same thing as a Dehn twist ξ -flat.

Given a connected $V \Subset S$ with $\xi(V) \leq 1$, the marking complex $\mathcal{M}(V)$ is quasi-isometric to a tree which we denote $\mathcal{TM}(V)$. When V is an annulus, $\mathcal{TM}(V)$ is isometric to \mathbf{R} , and the Dehn twist about V acts naturally by translation on $\mathcal{TM}(V)$. In the other two cases, where V is a one-holed torus or four-holed sphere, $\mathcal{TM}(V)$ is isometric to the dual tree of the usual modular diagram for $\mathrm{SL}_2\mathbb{Z}$, on which $\mathrm{MCG}(V)$ acts naturally. Given a geodesic segment $r \subset \mathcal{TM}(V)$ of positive length — finite, half-infinite, or bi-infinite — and given an annulus $U \Subset S$, we say that r is a *twist segment with support U* if one of the following holds: $U \stackrel{i}{\cong} V$ is an annulus; or V is not an annulus, $U \Subset V$, and r is contained in the axis of the Dehn twist about U .

Consider a subset of $\mathcal{M}(S)$ formed as follows. Choose a partial marking μ such that the components W_1, \dots, W_m of $W = \mathrm{open}(\mu)$ satisfy $\xi(W_i) \leq 1$. In each tree $\mathcal{TM}(W_i)$ choose r_i to be a geodesic, finite, half-infinite or bi-infinite (we allow length 0 as well). Let $r = \{r_1, \dots, r_m\}$. The *cube* $C(\mu, W, r)$ is the subset of $\mathcal{Q}(\mu)$ consisting of markings which, in each W_i , restrict to a marking in the geodesic r_i . In other words, under the quasi-isometry

$$\begin{aligned} \mathcal{Q}(\mu) &\approx \mathcal{M}(W_1) \times \cdots \times \mathcal{M}(W_m) \\ &\approx \mathcal{TM}(W_1) \times \cdots \times \mathcal{TM}(W_m) \end{aligned}$$

we have

$$C(\mu, W, r) \approx r_1 \times \cdots \times r_m$$

Junctures of cubes can be described in a reasonably straightforward manner, with careful bookkeeping, in terms of the description of junctures of product sets given in Lemma 3.3. Here are the details.

Lemma 3.7. *The family of cubes has junctures.*

Proof. Given cubes $C(\mu, W, r)$ and $C(\nu, V, s)$, we must construct subcubes which will function as junctures. Denote the components as $W = W_1 \cup \cdots \cup W_m$ and $V = V_1 \cup \cdots \cup V_n$.

First we describe the essential subsurface $\mathrm{open}(\mu, \nu) = W \circledast V$, whose marking complex parameterizes the junctures of $\mathcal{Q}(\mu)$ and $\mathcal{Q}(\nu)$, by Lemma 3.3.

We claim that the components may be reindexed as

$$\begin{aligned} W &= (W_1 \cup \cdots \cup W_k) \cup (W_{k+1} \cup \cdots \cup W_m) \\ V &= (V_1 \cup \cdots \cup V_k) \cup (V_{k+1} \cup \cdots \cup V_n) \end{aligned}$$

where $k \geq 0$, so that the components of $W \circledast V$ are

$$W \circledast V = \underbrace{(W_1 \circledast V_1)}_{U_1} \cup \cdots \cup \underbrace{(W_k \circledast V_k)}_{U_k}$$

and so that for each $i = 1, \dots, k$ one of the following holds: either $W_i \stackrel{i}{\cong} V_i \stackrel{i}{\cong} U_i$; or U_i is an annulus which is essentially contained in W_i and in V_i .

More generally, consider two connected, essential subsurfaces X, Y of S with $\xi(X), \xi(Y) \leq 1$. If $U = X \circledast Y$ is nonempty, it can only be an annulus or all of X and Y . The complement of an annulus in X , if X is not an annulus itself, is either one or two 3-holed spheres. Now if Z is disjoint from Y , we claim that $X \circledast Z$ is empty. For any curve c in $\Gamma(Z) \cap \Gamma(X)$ would have to be essential in X and isotopic to the complement of Y – hence U would be an annulus and c isotopic to its core. This would make c essential in both Y and Z , which is impossible unless Y and Z are isotopic annuli.

In the context of W and V , this implies that the relation $W_i \circledast V_j \neq \emptyset$ is a bijection between a subset of the components of W and a subset of the components of V , and the claim immediately follows.

Now we will construct a quasi-isometric embedding of $\mathcal{Q}(\mu)$ and $\mathcal{Q}(\nu)$ into a product of trees, which will allow us to see their junctures more clearly.

For each $i = 1, \dots, k$ the inclusion $U_i \subset W_i$ induces an embedding $\mathcal{T}\mathcal{M}(U_i) \hookrightarrow \mathcal{T}\mathcal{M}(W_i)$ whose image is a subtree denoted $\tau_i \subset \mathcal{T}\mathcal{M}(W_i)$: either $U_i \stackrel{i}{\cong} W_i$ and $\tau_i = \mathcal{T}\mathcal{M}(W_i)$; or U_i is an annulus and τ_i is the axis in $\mathcal{T}\mathcal{M}(W_i)$ of the Dehn twist about U_i . Similarly, the inclusion $U_i \subset V_i$ induces a quasi-isometric embedding $\mathcal{M}(U_i) \rightarrow \mathcal{T}\mathcal{M}(V_i)$ whose image is a subtree $\sigma_i \subset \mathcal{T}\mathcal{M}(V_i)$.

By composing a coarse inverse of the map $\mathcal{M}(U_i) \rightarrow \tau_i$ with the map $\mathcal{M}(U_i) \rightarrow \sigma_i$, we obtain a quasi-isometry $g_i: \tau_i \rightarrow \sigma_i$. Notice that we may take g_i to be a simplicial isomorphism, as one can verify easily in either of two cases: if $W_i \stackrel{i}{\cong} V_i \stackrel{i}{\cong} U_i$ then these isotopies induce simplicial isomorphisms of marking complexes; and otherwise τ_i and σ_i are the axes in the trees $\mathcal{T}\mathcal{M}(W_i)$ and $\mathcal{T}\mathcal{M}(V_i)$, respectively, of the Dehn twist about U_i , and we can take g_i to be a simplicial isomorphism between these two axes. Let X_i be the tree obtained from the disjoint union of the trees $\mathcal{T}\mathcal{M}(W_i)$ and $\mathcal{T}\mathcal{M}(V_i)$ by gluing τ_i to σ_i isometrically using the map g_i .

Let

$$\Upsilon = \mathbb{R} \times \prod_{i=1}^k X_i \times \prod_{i=k+1}^m \mathcal{T}\mathcal{M}(W_i) \times \prod_{i=k+1}^n \mathcal{T}\mathcal{M}(V_i).$$

This is a product of trees on which we can put the ℓ^1 metric.

Now for $i = k+1, \dots, m$, let $p_i = \pi_{\mathcal{M}(W_i)}(\nu)$ and note that in fact a bounded neighborhood of p_i contains all of $\pi_{\mathcal{M}(W_i)}(\mathcal{Q}(\nu))$. Similarly, for $j = k+1, \dots, n$ let $q_j = \pi_{\mathcal{M}(V_j)}(\mu)$ which approximates $\pi_{\mathcal{M}(V_j)}(\mathcal{Q}(\mu))$.

The product structure of $\mathcal{Q}(\mu)$ (Proposition 3.1) now gives us a quasi-isometric embedding

$$\theta^\mu: \mathcal{Q}(\mu) \rightarrow \Upsilon$$

which is the identity on the $\mathcal{TM}(W_i)$ factors (including those embedded in the X_i), and maps to the constant q_j on each $\mathcal{TM}(V_j)$, $j = k + 1, \dots, n$, and to 0 in the \mathbb{R} factor. Similarly we have

$$\theta^\nu: \mathcal{Q}(\nu) \rightarrow \Upsilon,$$

which is the identity on the $\mathcal{TM}(V_i)$ factors (including those embedded in the X_i), and maps to the constant p_j on each $\mathcal{TM}(W_j)$, $j = k + 1, \dots, m$, and to D in the \mathbb{R} factor, where D is the Hausdorff distance between $E(\mu, \nu)$ and $E(\nu, \mu)$.

Note, by Lemma 3.3, that the images $\theta^\mu(E(\mu, \nu))$ and $\theta^\nu(E(\nu, \mu))$ are parallel products of subtrees, namely

$$\{0\} \times \prod_1^k \sigma_i \times \prod_{k+1}^m \{p_i\} \times \prod_{k+1}^n \{q_i\}$$

and

$$\{D\} \times \prod_1^k \tau_i \times \prod_{k+1}^m \{p_i\} \times \prod_{k+1}^n \{q_i\}$$

recalling that σ_i and τ_i are identified in X_i .

Moreover we note that, by the distance formula (3) in Lemma 3.3 (and its interpretation in terms of projections in (3.2,3.3,3.4,3.5)), θ^μ and θ^ν actually combine to give us a quasi-isometric embedding of the union $\mathcal{Q}(\mu) \cup \mathcal{Q}(\nu)$ into Υ , which we will call θ .

In particular $\theta(E(\mu, \nu))$ and $\theta(E(\nu, \mu))$ are junctures for $\theta(\mathcal{Q}(\mu))$ and $\theta(\mathcal{Q}(\nu))$ in this product of trees. This is a special case of the following easy fact:

Lemma 3.8. *Let $T = T_1 \times \dots \times T_N$ be a product of complete trees with the ℓ^1 metric. Then the family of products of closed subtrees has junctures. Moreover the approximations in the definition of junctures are all exact.*

Proof. For a single tree this is easily checked: Any two subtrees either intersect, in which case the junctures are (two copies of) their common subtree, or are disjoint, in which case the junctures are the unique points of closest approach of each tree to the other. For a product of subtrees in a product of trees, the junctures are the products of junctures in the factors, and the distance formulas in the factors sum to give the desired outcome. \square

Now it is easy to understand how the cubes $C(\mu, W, r)$ and $C(\nu, V, s)$ are situated by considering their θ -images. $\theta(C(\mu, W, r))$ is a product of lines and points in the factors of Υ , with first coordinate 0, and $\theta(C(\nu, V, s))$ is a similar product with first coordinate D . Lemma 3.8 implies that the junctures of

the images are again products of subintervals, and we conclude that the θ -preimages, which are subcubes of the original cubes, are also junctures. \square

The proof of Lemma 3.7 gives some more information about the structure of the junctures of two cubes, which we record here:

Lemma 3.9. *Let $C_1 = C(\mu, W, r)$ and $C_2 = C(\nu, V, s)$. The junctures $C_{ij} = E(C_i, C_j) \subset C_i$ are subcubes of the form $C_{12} = C(\mu, W, r')$ and $C_{21} = C(\nu, V, s')$, where each component of r' or s' is a subinterval or point of the corresponding component of r or s .*

After the renumbering in the proof of Lemma 3.7, the components of r' and s' that are not single points come in pairs r'_i, s'_i such that $U_i = W_i \circlearrowleft V_i \neq \emptyset$, and r'_i, s'_i are images of the same segment of $\mathcal{TM}(U_i)$ under the quasi-isometric embeddings $\mathcal{TM}(U_i) \hookrightarrow \mathcal{TM}(W_i)$ and $\mathcal{TM}(U_i) \hookrightarrow \mathcal{TM}(V_i)$. Furthermore:

- (1) *If U_i is an annulus then r'_i, s'_i are twist segments with support U_i .*
- (2) *If $U_i \stackrel{i}{\cong} V_i \stackrel{i}{\cong} W_i$ then $r'_i = s'_i = r_i \cap s_i$ in the tree $\mathcal{TM}(U_i) = \mathcal{TM}(V_i) = \mathcal{TM}(W_i)$.*

Proof. This is a consequence of the fact that the map θ in the proof of Lemma 3.7 respects the product structures in its domain and range. The image of $C(\mu, W, r)$ in Υ is a product of line segments in the factors X_1, \dots, X_k and $\mathcal{TM}(W_{k+1}), \dots, \mathcal{TM}(W_m)$, and points in the other factors, whereas $C(\nu, V, s)$ maps to a product of line segments in X_1, \dots, X_k and in $\mathcal{TM}(V_{k+1}), \dots, \mathcal{TM}(V_n)$, and points in the rest. Thus, any tree factor in which the juncture factor is a nondegenerate segment is an X_i which corresponds to a pair W_i, V_i that has nontrivial essential intersection U_i , this nondegenerate segment is the intersection of the images of r_i and s_i via the quasi-isometric embeddings $\mathcal{TM}(W_i) \hookrightarrow X_i$ and $\mathcal{TM}(V_i) \hookrightarrow X_i$, the pullbacks of this segment to $\mathcal{TM}(W_i)$ and to $\mathcal{TM}(V_i)$ are the segments r'_i, s'_i respectively, and the pullbacks of these two segments to $\mathcal{TM}(U_i)$ are the same segment; in the special case when U_i is an annulus item (1) is an immediate consequence, and when $U_i \stackrel{i}{\cong} V_i \stackrel{i}{\cong} W_i$ item (2) is an immediate consequence. The pullbacks of the junctures by θ are then subcubes respecting the product structures of the original cubes, and with nondegenerate segments r'_i, s'_i only in the factors $\mathcal{TM}(W_i), \mathcal{TM}(V_i)$ corresponding to the X_i . \square

3.4. Cubes and junctures in the asymptotic cone

From the definition of junctures we can obtain the following statement in the asymptotic cone: Let \mathcal{F} be a family with junctures in \mathcal{M} , let (X_n) and (Y_n) be sequences in \mathcal{F} and let X_ω and Y_ω be their rescaling ultralimits in the cone \mathcal{M}_ω . We find that $E(X, Y)_\omega$ and $E(Y, X)_\omega$ are either disjoint or identical, depending on rate of growth of the Hausdorff distance. Property 2

in the definition of junctures also implies that

$$X_\omega \cap Y_\omega = E(X, Y)_\omega \cap E(Y, X)_\omega,$$

and hence this intersection is either empty or equal to the limit of the junctures.

Now given a sequence of cubes $C^n = C(\mu^n, W^n, r^n)$, which we denote $\overline{C} = C(\overline{\mu}, \overline{W}, \overline{r})$, we can take the cone $C^\omega(\overline{\mu}, \overline{W}, \overline{r})$, which is nonempty provided that the distance from the cubes to the basepoint of $\mathcal{M}(S)$ does not grow too fast. This object has dimension less than or equal to the number of components of W^n for ω -a.e. n . In fact the limit cube is naturally bilipschitz homeomorphic to $r_1^\omega \times \cdots \times r_k^\omega$ where each r_i^ω is an embedded path in the \mathbf{R} -tree $\mathcal{M}_\omega(\overline{W}_i)$ whose length is in $[0, \infty]$, having positive length if and only if the length of the sequence (r_i^n) grows linearly. We will continue calling these objects cubes.

Lemma 3.7 on junctures for cubes implies, using the discussion in the beginning of this section, that the intersection of two cubes in $\mathcal{M}(S)$ is empty or is a cube, possibly a trivial cube, meaning a single point. Moreover this cube is described by data closely related to the original cubes. We will use this in Section 9 to understand the *complex of orthants* in the asymptotic cone.

3.5. Coarse set theory of Dehn twist k -flats

We have already explained in Lemma 3.2 how junctures are examples of coarse intersection. In this section we make a further study of coarse inclusion and coarse equivalence among Dehn twist k -flats in $\mathcal{M}(S)$.

A *Dehn twist k -flat* in $\mathcal{M}(S)$, $0 \leq k \leq \xi(S)$, is a subset of the form $\mathcal{Q}(\mu)$ where μ is a marking such that $\text{base}(\mu)$ is a pants decomposition with exactly k unmarked components. As the terminology suggests, each Dehn twist k -flat is a quasi-isometrically embedded copy of \mathbf{R}^k in $\mathcal{M}(S)$, by Proposition 3.1 (1).

Throughout the paper, the phrase “Dehn twist flat”, when unadorned by a dimension, will by default refer to the top dimensional case, namely a Dehn twist $\xi(S)$ -flat. Sometimes we emphasize this by referring to *maximal Dehn twist flats*.

In a metric space \mathcal{M} , given two subsets $A, B \subset \mathcal{M}$, we say that A is *coarsely included* in B if there exists $r \in [0, \infty)$ such that $A \subset \mathcal{N}_r(B)$. We say that A is *coarsely equivalent* to B if there exists $r \in [0, \infty)$ such that $A \subset \mathcal{N}_r(B)$ and $B \subset \mathcal{N}_r(A)$; the infimum of all such $r \in [0, \infty]$ is equal to the Hausdorff distance between A and B .

Given a $k - 1$ simplex σ in $\mathcal{C}(S)$ and a Dehn twist k -flat $\mathcal{Q}(\mu)$, we say that σ is *represented by* $\mathcal{Q}(\mu)$ if the system of unmarked curves of $\text{base}(\mu)$ is isotopic to σ . In particular, vertices are represented by Dehn twist 1-flats and edges by Dehn twist 2-flats. Note that each Dehn twist k -flat represents a unique $k - 1$ simplex. Conversely, each $k - 1$ simplex is represented by infinitely many Dehn twist k -flats, except in the case $k = \xi$ where each maximal dimension simplex is represented by a unique maximal Dehn twist flat.

Lemma 3.10. *Given simplices $\sigma_i \subset \mathcal{C}(S)$, $i = 0, 1$, and representative Dehn twist k_i -flats $\mathcal{Q}(\mu_i)$ respectively, we have:*

- (1) $\mathcal{Q}(\mu_0)$ is coarsely contained in $\mathcal{Q}(\mu_1)$ if and only if $\sigma_0 \subset \sigma_1$.
- (2) $\mathcal{Q}(\mu_0)$ is coarsely equivalent to $\mathcal{Q}(\mu_1)$ if and only if $\sigma_0 = \sigma_1$.
- (3) Given a simplex $\tau \subset \mathcal{C}(S)$ and a representative Dehn twist l -flat $\mathcal{Q}(\nu)$, $\mathcal{Q}(\nu)$ represents the coarse intersection of $\mathcal{Q}(\mu_0)$ and $\mathcal{Q}(\mu_1)$ if and only if $\tau = \sigma_0 \cap \sigma_1$.

Proof. First note that $\sigma_0 \subset \sigma_1$ if and only if $\mathcal{Q}(\mu_0)$ is equal to $E(\mathcal{Q}(\mu_0), \mathcal{Q}(\mu_1)) = \mathcal{Q}(\mu_0 \rfloor \mu_1)$. The “if” direction of item (1) is then an immediate consequence of Lemma 3.3 (2). Conversely, if $\sigma_0 \not\subset \sigma_1$ then by applying Lemma 3.1 (1) it follows that there exist points of $\mathcal{Q}(\mu_0)$ which are arbitrarily far from $\mathcal{Q}(\mu_0 \rfloor \mu_1)$, and then Lemma 3.3 (3) proves that $\mathcal{Q}(\mu_0)$ is not contained in any finite radius neighborhood of $\mathcal{Q}(\mu_1)$.

Item (2) follows by symmetric applications of item (1).

To prove item (3), first apply Lemma 3.3 to conclude that $\mathcal{Q}(\mu_0 \rfloor \mu_1)$ and $\mathcal{Q}(\mu_1 \rfloor \mu_0)$ are the junctures of $\mathcal{Q}(\mu_0)$ and $\mathcal{Q}(\mu_1)$, then apply Lemma 3.2 to conclude that these two junctures both represent the coarse intersection of $\mathcal{Q}(\mu_0)$ and $\mathcal{Q}(\mu_1)$, and then note that both of these junctures represent the simplex $\sigma_0 \cap \sigma_1$. Now apply item (2) to $\mathcal{Q}(\nu)$ and either of the two junctures. \square

The proof of Lemma 3.10 (3) does not make full use of the uniform control on coarse intersection that is provided by Lemma 3.2. The following lemma make use of this control, which will be needed in Section 10.3.

Lemma 3.11. *For each Dehn twist k -flat $\mathcal{Q}(\mu)$ there exist two maximal Dehn twist flats $\mathcal{Q}(\mu_0)$, $\mathcal{Q}(\mu_1)$ whose junctures $\mathcal{Q}(\mu_0 \rfloor \mu_1)$ and $\mathcal{Q}(\mu_1 \rfloor \mu_0)$ are Dehn twist k -flats at uniform Hausdorff distance from each other and from $\mathcal{Q}(\mu)$, and which uniformly represent the coarse intersection of $\mathcal{Q}(\mu_0)$ and $\mathcal{Q}(\mu_1)$. To be precise:*

$$\begin{aligned} \mathcal{Q}(\mu) \cup \mathcal{Q}(\mu_0 \rfloor \mu_1) \cup \mathcal{Q}(\mu_1 \rfloor \mu_0) &\subset \mathcal{N}_R(\mathcal{Q}(\mu_0)) \cap \mathcal{N}_R(\mathcal{Q}(\mu_1)) \\ &\subset \mathcal{N}_C(\mathcal{Q}(\mu)) \cap \mathcal{N}_C(\mathcal{Q}(\mu_0 \rfloor \mu_1)) \cap \mathcal{N}_C(\mathcal{Q}(\mu_1 \rfloor \mu_0)) \end{aligned}$$

where the constants $C, R \geq 0$ depend only on the topology of S .

Proof of Lemma 3.11. It suffices to prove that the three sets $\mathcal{Q}(\mu)$, $\mathcal{Q}(\mu_0 \rfloor \mu_1)$, $\mathcal{Q}(\mu_1 \rfloor \mu_0)$ are all at uniformly finite Hausdorff dimension from each other, for once this is done we may apply Lemma 3.3 to obtain that $\mathcal{Q}(\mu_0 \rfloor \mu_1)$, $\mathcal{Q}(\mu_1 \rfloor \mu_0)$ are the junctures of $\mathcal{Q}(\mu_0)$, $\mathcal{Q}(\mu_1)$, and then we may apply Lemma 3.2 to obtain the desired uniform control on coarse intersection.

We first prove the lemma in the special case that $k = 0$, so μ is a marking and $\mathcal{Q}(\mu) = \{\mu\}$. Note that a pants decomposition Δ overlaps each component of $\text{base}(\mu)$ if and only if $\text{base}(\mu)$ overlaps each component of Δ , in which case both $\text{base}(\mu) \rfloor \Delta$ and $\Delta \rfloor \text{base}(\mu)$ are markings of S . We shall find such a Δ so

that each of the markings $\text{base}(\mu)\rfloor\Delta$ and $\Delta\rfloor\text{base}(\mu)$ is uniformly close to μ in $\mathcal{M}(S)$.

To find the appropriate Δ , letting $\text{base}(\mu) = \{c_1, \dots, c_\xi\}$, we shall build up $\Delta = \{d_1, \dots, d_\xi\}$ one component at a time. Proceeding by induction, choose the subset $\{d_1, \dots, d_k\}$ so that:

- $\{d_1, \dots, d_k, c_{k+1}, \dots, c_\xi\}$ is a pants decomposition.
- c_k, d_k are connected by an edge in the curve complex of the complexity 1 component of $S - (d_1 \cup \dots \cup d_{k-1} \cup c_{k+1} \cup \dots \cup c_\xi)$ that contains them.
- $\pi_{\mathcal{C}(c_k)}(d_k)$ is a uniform distance from the μ -transversal of c_k .

The last item is possible because each orbit of the action on $\mathcal{C}(c_k)$ of the Dehn twist group $\langle \tau_{c_k} \rangle$ comes uniformly close to each point of $\mathcal{C}(c_k)$. Application of Proposition 3.1 (1) provides a uniform bound to $d(\mu, \text{base}(\mu)\rfloor\Delta)$. Consider the sequence of markings defined by

$$\begin{aligned} \mu_0 &= \text{base}(\mu)\rfloor\Delta \\ &= ((c_1, d_1), \dots, (c_k, d_k)) \\ \mu_i &= ((d_1, c_1), \dots, (d_i, c_i), (c_{i+1}, d_{i+1}), \dots, (c_k, d_k)), \quad i = 1, \dots, k-1 \\ \mu_k &= \Delta\rfloor\text{base}(\mu) \\ &= ((d_1, c_1), \dots, (d_k, c_k)) \end{aligned}$$

The markings μ_i, μ_{i+1} are connected by an edge in $\mathcal{M}(S)$ and so

$$d(\text{base}(\mu)\rfloor\Delta, \Delta\rfloor\text{base}(\mu)) \leq k$$

Now we reduce the general case of the lemma to the special case just proved. Consider any Dehn twist k -flat $\mathcal{Q}(\mu)$ with $k \leq \xi - 1$. Write $\mu = \mu' \cup \mu''$ where μ' consists of the k unmarked curves of $\text{base}(\mu)$ and μ'' consists of the $\xi - k$ marked curves together with their transversals. The subsurface $F = \text{supp}(\mu)$ is the union of the components of $S - \mu'$ that are not 3-holed spheres, and μ'' is a marking of F . Applying the special case we obtain a pants decomposition Δ of F such that $\text{base}(\mu'')\rfloor\Delta$ and $\Delta\rfloor\text{base}(\mu'')$ are markings of F each at uniformly bounded distance from μ'' in $\mathcal{M}(F)$. Let $\mu_0 = \text{base}(\mu) = \mu' \cup \text{base}(\mu'')$ and $\mu_1 = \mu' \cup \Delta$. By construction the sets $\mathcal{Q}(\mu)$, $\mathcal{Q}(\mu_0\rfloor\mu_1)$, $\mathcal{Q}(\mu_1\rfloor\mu_0)$ are all their own junctures among each other, and Lemma 3.3 (2) provides a uniform bound to their Hausdorff distances. \square

4. Consistency theorem

In this section we will derive a coarse characterization of the image of the *curve complex projections map*

$$\Pi: \mathcal{M}(S) \rightarrow \prod_{W \subseteq S} \mathcal{C}(W)$$

defined by $\Pi(\mu) = (\pi_W(\mu))_W$.

Consider the following *Consistency Conditions* on a tuple $(x_W) \in \prod_W \mathcal{C}(W)$, where c_1 and c_2 are a pair of positive numbers:

C1: Whenever $W \pitchfork V$,

$$\min(d_W(x_W, \partial V), d_V(x_V, \partial W)) < c_1.$$

C2: Whenever $V \subseteq W$ and $d_W(x_W, \partial V) > c_2$,

$$d_V(x_V, x_W) < c_1.$$

In C2 recall that $d_V(x_V, x_W)$ is shorthand for $d_V(x_V, \pi_V(x_W))$ and so is only defined if x_W is in the domain of π_V , which it may not be if $d_W(x_W, \partial V)$ is too small. Notice that if C1 and C2 are satisfied with respect to positive constants c_1, c_2 , then they are satisfied with respect to any larger constants.

The conditions C1-C2, with suitable constants, are satisfied by the image of Π , and moreover

Lemma 4.1. *Given K there exist $c_1, c_2 \geq 1$ such that, if $\mu \in \mathcal{M}(S)$ and $(x_W) \in \prod \mathcal{C}(W)$ such that $d_W(x_W, \mu) \leq K$ for all $W \subseteq S$, then (x_W) satisfies C1 and C2 with constants c_1, c_2 .*

Proof. The case $K = 0$, i.e. $(x_W) = \Pi(\mu)$, follows from Behrstock's inequality [3], namely

Lemma 4.2. *There exists m_0 such that for any marking $\mu \in \mathcal{M}(S)$ and subsurfaces $V \pitchfork W$,*

$$\min(d_W(\mu, \partial V), d_V(\mu, \partial W)) < m_0.$$

This gives condition C1. Condition C2, with $c_2 = 1$ and suitable c_1 , follows simply because π_V is determined by intersections, so whenever $V \subseteq W$, $\pi_V \circ \pi_W$ is a bounded distance from π_V when both are defined (Lemma 2.12).

For $K > 0$ we simply observe that (C1-2) are preserved, with suitable change in constants, when all the coordinates of (x_W) are changed a bounded amount. \square

Our main point here is to show that conditions (C1-2) are also *sufficient* for a point to be close to the image of Π , namely:

Theorem 4.3. *Given c_1 and c_2 there exists c_3 such that, if (C1-2) hold with c_1 and c_2 for a point (x_W) , then there exists $\mu \in \mathcal{M}(S)$ such that*

$$d_W(x_W, \mu) < c_3$$

for all $W \subseteq S$.

The proof of this theorem will take up the rest of Section 4.

4.1. Subsurface ordering induced by projections

In order to approach the proof of Theorem 4.3 we will first study more carefully the structure imposed by (C1) and (C2). Recall from [39] that for any two markings of S , there is a natural partial order on the set of component domains of subsurfaces that occur in a hierarchy between those two markings. With this motivation, the structure that we study will involve similar partial orders on the collection of proper, connected, essential subsurfaces (up to isotopy).

Let us fix a tuple (x_W) satisfying (C1-2). Without loss of generality, we will assume $c_1 > \max\{c_2, m_0, B\}$, where m_0 is the constant given by Lemma 4.2 and B is the constant given by Theorem 2.6.

If W, V are proper, connected, essential subsurfaces of S and $k \in \mathbb{N}$, define a relation

$$W \prec_k V$$

to mean that

$$W \pitchfork V \text{ and } d_W(x_W, \partial V) \geq k(c_1 + 4).$$

The role of 4 here is that it is twice the maximal diameter of $\pi_{\mathcal{C}(Y)}(\gamma)$ for a curve system γ — see lemmas 2.10 and 2.11 and the comments thereafter.

We also allow the right hand side of \prec_k to be a marking ρ : define $W \prec_k \rho$ to mean that $d_W(x_W, \rho) \geq k(c_1 + 4)$.

Although \prec_k is not quite an order relation on the set of proper, connected, essential subsurfaces, the family of all \prec_k behaves roughly like a partial order in a way we shall now explore. Let us also define a relation

$$W \ll_k \rho,$$

where ρ is any partial marking, to mean that

$$W \pitchfork \rho \text{ and } d_W(x_W, \rho) \geq k(c_1 + 4).$$

Notice that if ρ is a marking then $W \prec_k \rho$ and $W \ll_k \rho$ are equivalent, since W and ρ always overlap.

We then define $W \ll_k V$ to mean $W \ll_k \partial V$. This is a weaker relation than $W \prec_k V$ because $W \pitchfork \partial V$ allows the possibility that $V \subseteq W$, which cannot happen if $W \pitchfork V$.

Note that if $k \geq p$ then $U \ll_k V \implies U \ll_p V$ and $U \prec_k V \implies U \prec_p V$. Next we point out that property C1 implies the following:

- \prec_k is antisymmetric in the following sense: if $U \prec_k V$ holds, then $V \not\prec_1 U$, and hence $V \not\prec_k U$.

Of course \ll_k is antisymmetric as well, since containment is already antisymmetric. Now we will prove the following lemma, which states that the system of relations is transitive in a certain sense.

Lemma 4.4. *Given an integer $k > 1$ we have:*

- (1) If $U \prec_k V$ and $V \ll_2 W$ then $U \prec_{k-1} W$.
Also, if ρ is a marking, and if $U \prec_k V$ and $V \ll_2 \rho$, then $U \prec_{k-1} \rho$.
- (2) If $U \ll_k V$ and $V \ll_2 W$ then $U \ll_{k-1} W$.
Also, if ρ is a marking, and if $U \ll_k V$ and $V \ll_2 \rho$ then $U \ll_{k-1} \rho$.
- (3) If $U \pitchfork V$ and both $U \ll_k \rho$ and $V \ll_k \rho$ for some partial marking ρ , then U and V are \prec_{k-1} -ordered — that is, either $U \prec_{k-1} V$ or $V \prec_{k-1} U$.

Note that the weak transitivity of parts (1) and (2) tends to “decay” (k decreases) each time it is applied, and hence does not give a partial order. However part (3) can be used to re-strengthen the inequalities under appropriate circumstances.

Proof. Beginning with (1), suppose $U \prec_k V$ and $V \ll_2 W$. From $V \ll_2 W$ we have

$$d_V(x_V, \partial W) \geq 2(c_1 + 4)$$

and from $U \prec_k V$ and property C1 we have

$$d_V(x_V, \partial U) < c_1.$$

By the triangle inequality, together with the fact that $\text{diam} \pi_{\mathcal{C}(V)}(\gamma) \leq 2$ for any disjoint curve system γ (see comments after Lemma 2.11),

$$\begin{aligned} d_V(\partial U, \partial W) &\geq d_V(x_V, \partial W) - d_V(x_V, \partial U) - \text{diam}_V(\partial U) - \text{diam}_V(\partial W) \\ &> 2(c_1 + 4) - c_1 - 4 = c_1 + 4. \end{aligned}$$

In particular $d_V(\partial U, \partial W) > 2$, so $\partial U \pitchfork \partial W$, and so $U \pitchfork W$. Now applying Lemma 4.2 we also get

$$d_U(\partial V, \partial W) < m_0 \leq c_1$$

and hence, using $U \prec_k V$ and the triangle inequality as above,

$$d_U(x_U, \partial W) > (k-1)(c_1 + 4).$$

Hence, $U \prec_{k-1} W$, as desired.

Replacing W (and ∂W) by a marking ρ , the second clause of (1) is proved similarly, noting that the question of whether $U \pitchfork \rho$ is not an issue.

The proof of part (2) is similar. The case not covered by part (1) is when $V \subseteq U$, and we consider two cases, depending on whether $W \subseteq V$ or $W \pitchfork V$.

In the first case suppose that $W \subseteq V$. Since $V \not\cong U$, it follows that ∂V and ∂W together form a curve system in U , and hence $\text{diam}_U(\partial V \cup \partial W) \leq 1$. So by the triangle inequality we have

$$d_U(x_U, \partial W) \geq d_U(x_U, \partial V) - 1 \geq (k-1)(c_1 + 4)$$

and we conclude $U \ll_{k-1} W$.

In the second case suppose that $W \pitchfork V$. Since $V \ll_2 W$ we have

$$d_V(x_V, \partial W) \geq 2(c_1 + 4).$$

Since $V \Subset U$ we know that $\partial W \pitchfork U$. Since $d_U(x_U, \partial V) \geq k(c_1 + 4) > c_2$, by property C2 we have that

$$d_V(x_V, x_U) < c_1$$

and hence

$$d_V(x_U, \partial W) \geq d_V(x_V, \partial W) - d_V(x_V, x_U) - \text{diam}_V(\partial W) > c_1 + 4.$$

But now by Theorem 2.6, this implies that any $\mathcal{C}(U)$ -geodesic $[x_U, \pi_U(\partial W)]$ must pass within distance 1 of ∂V , and we conclude

$$d_U(x_U, \partial W) \geq d_U(x_U, \partial V) - 1 - \text{diam}_U(\partial V) \geq (k-1)(c_1 + 4)$$

and again we have $U \ll_{k-1} W$.

Again replacing W (and ∂W) by a marking ρ , the second clause of (2) is proved similarly, only the second case of the proof being relevant.

Now we prove (3): starting with $U \pitchfork V$ and

$$d_U(x_U, \rho) \geq k(c_1 + 4)$$

and

$$d_V(x_V, \rho) \geq k(c_1 + 4),$$

suppose $U \not\prec_{k-1} V$, so that

$$d_U(x_U, \partial V) < (k-1)(c_1 + 4).$$

Then by the triangle inequality

$$d_U(\rho, \partial V) > c_1$$

and by Lemma 4.2

$$d_V(\partial U, \rho) < m_0.$$

Now by the triangle inequality

$$d_V(x_V, \partial U) > (k-1)(c_1 + 4)$$

so $V \prec_{k-1} U$, and we are done. \square

Let us now define

$$\mathcal{F}_k((x_W), \rho) = \{W \subsetneq S : W \pitchfork \rho \text{ and } d_W(x_W, \rho) \geq k(c_1 + 4)\}.$$

which is the collection of subsurfaces in whose curve complexes x_W and ρ have large distance. If $x_W = \pi_W(x)$ for $x \in \mathcal{M}(S)$, then we use the notation $\mathcal{F}_k(x, \rho)$.

Note that $\mathcal{F}_k((x_W), \rho) = \{W \subsetneq S : W \ll_k \rho\}$, and so we sometimes simplify the notation by simply writing $\mathcal{F}_k(\rho)$. In this vein, when Z is a subsurface, we let $\mathcal{F}_k(Z)$ denote $\mathcal{F}_k(\partial Z)$.

As a corollary of the previous lemma we obtain:

Lemma 4.5. *If $k > 2$ then the relation \prec_{k-1} is a partial order on $\mathcal{F}_k((x_W), \rho)$.*

Proof. All that is needed is to prove that \prec_{k-1} is transitive on $\mathcal{F}_k(\rho)$ — antisymmetry is already established.

Suppose $U, V, W \in \mathcal{F}_k(\rho)$, and $U \prec_{k-1} V$ and $V \prec_{k-1} W$. By Lemma 4.4 part (1), this implies $U \prec_{k-2} W$. In particular $U \pitchfork W$, so by Lemma 4.4 part (3) U and W are \prec_{k-1} -ordered. Antisymmetry together with $U \prec_{k-2} W$ implies that $U \prec_{k-1} W$, as desired. \square

We can also obtain a finiteness statement:

Lemma 4.6. *If $k > 2$ then $\mathcal{F}_k((x_W), \rho)$ is finite.*

Proof. Suppose that $\mathcal{F}_k(\rho)$ is infinite and let $\{Y_i\}$ be an infinite, injective sequence within it. After extracting a subsequence we may assume that $\partial Y_i \rightarrow \lambda$ in $\mathcal{PML}(S)$, the projective measured lamination space of S . Let U be a subsurface filled by a component of λ ; possibly $U = S$. Then ∂Y_i meets U for all sufficiently large i , and $\pi_U(\partial Y_i) \rightarrow \infty$ in $\mathcal{C}(U)$ — that is, $d_U(\partial Y_i, q) \rightarrow \infty$ for any fixed q . This is a consequence of the Kobayashi/Luo argument that $\mathcal{C}(U)$ has infinite diameter, see [38, Proposition 3.6]. Note in the special case that U is an annulus we are obtaining that the twisting of ∂Y_i around U is going to ∞ .

Now, $d_U(x_U, \partial Y_i) \rightarrow \infty$ means for any given p that eventually $U \ll_p Y_i$. However we have $Y_i \ll_k \rho$ by assumption, so

$$U \ll_{p-1} \rho$$

by Lemma 4.4, part (2). However U and ρ are fixed and p is arbitrary, so this is impossible. We conclude that $\mathcal{F}_k(\rho)$ is finite. \square

4.2. Proof of the consistency theorem

As stated at the outset of Section 4.1 we have (x_W) satisfying C1-2 with the same assumptions on c_1 and c_2 . We will construct μ by induction.

Consider $\mathcal{F}_3(x_S)$, which we recall is shorthand for $\mathcal{F}_3((x_W), x_S)$. If $\mathcal{F}_3(x_S) = \emptyset$, let $\mu_0 = x_S$. Otherwise, by Lemmas 4.6 and 4.5, the set $\mathcal{F}_3(x_S)$ is finite and partially ordered by \prec_2 , and so this partial order contains minimal elements. Among these minimal elements, choose one, Y , of maximal complexity $\xi(Y)$, and let $\mu_0 = \partial Y$.

Now consider any $Z \in S$ which overlaps μ_0 . We claim that

$$d_Z(x_Z, \mu_0) < 4(c_1 + 4). \tag{4.1}$$

Suppose otherwise, so $Z \ll_4 \mu_0$. If $\mathcal{F}_3(x_S) = \emptyset$ then $Z \ll_4 x_S$ which implies $Z \in \mathcal{F}_4(x_S) \subseteq \mathcal{F}_3(x_S)$, a contradiction. When $\mathcal{F}_3(x_S) \neq \emptyset$, we would have $Z \ll_4 Y \ll_3 x_S$, and by Lemma 4.4 part (2), $Z \ll_3 x_S$. Hence $Z \in \mathcal{F}_3(x_S)$.

Now since Y was \prec_2 -minimal, we can't have $Z \prec_2 Y$ and we conclude $Y \in Z$. Now Z cannot be \prec_2 -minimal because its complexity is larger than that of Y ,

so there must be $V \in \mathcal{F}_3(x_S)$ with $V \prec_2 Z$. But then Lemma 4.4 part (1) implies $V \prec_1 Y$. In particular $V \pitchfork Y$ and so, arguing as in the proof of Lemma 4.5, we apply Lemma 4.4 part (3) with $\rho = x_S$ to conclude that V and Y are \prec_2 ordered, and since $V \prec_1 Y$ it follows by asymmetry that $V \prec_2 Y$. Again this is a contradiction. We conclude that (4.1) holds.

Now consider the restriction of (x_W) to subsurfaces in $S \setminus \mu_0$. In each component V of $S \setminus \mu_0$, the assumptions on (x_W) still hold, so inductively there is a marking μ_V in $\mathcal{M}(V)$ satisfying

$$d_Z(\mu_V, x_Z) < c_3(V) \tag{4.2}$$

for all $Z \subseteq V$. We append the μ_V to μ_0 to obtain a marking μ' which almost fills the surface except that it has no transversal data on the curves of μ_0 . By (4.1) and (4.2), it satisfies a bound on $d_Z(x_Z, \mu')$ for every $Z \subseteq S$ except the annuli whose cores are components of μ_0 . Let μ be the enlargement of μ' obtained by setting the transversal on each $\gamma \in \mu_0$ to be x_γ . Now we obtain a bound on $d_Z(x_Z, \mu)$ for all Z , so μ is the desired marking and the proof is complete.

5. Σ -hulls

A Σ -hull of a finite set in $\mathcal{M}(S)$ (and then, taking limits, of a finite set in $\mathcal{M}_\omega(S)$) is a substitute for convex hull which is well adapted to the presence in $\mathcal{M}(S)$ of both hyperbolicity and product structure. In “hyperbolic directions” it looks like a hyperbolic convex hull, and in product regions the hull of two points can be a rectangle. In general it is a hybrid of these.

In this section we focus on the “coarse Σ -hull” of a finite set $A \subset \mathcal{M}(S)$, a parameterized family of sets $\Sigma_\epsilon(A)$ which are coarsely well-defined for sufficiently large ϵ (see Lemma 5.4 (3)). Our main goal is Proposition 5.2, in which we show that Σ -hulls admit coarse retractions.

In Section 6 we will apply this to Σ -hulls in the asymptotic cone, showing that they are contractible and vary continuously with their extreme points.

5.1. Hulls in hyperbolic spaces

If $A \subset X$ is a finite subset of a δ -hyperbolic geodesic space X , let $\text{hull}_X(A)$ denote the union of geodesics $[a, a']$ with $a, a' \in A$, the *hyperbolic hull* of A in X . We will need the following properties of this construction, which are easy exercises.

Lemma 5.1. *The sets $\text{hull}_X(A)$ satisfy the following properties, with implicit constants depending only on the hyperbolicity constant of X and the cardinality of A :*

- (1) $\text{hull}_X(A)$ is quasi-convex.

- (2) If $x \in X$ and $y \in \text{hull}_X(A)$ is the point nearest to x , and if $y' \in \text{hull}_X(A)$, then $d(y, y') \prec d(x, y') - d(x, y)$.
- (3) The map $A \mapsto \text{hull}_X(A)$ is coarsely Lipschitz in the Hausdorff metric.

Also, for points $x \in X$ and closed, quasiconvex subsets $B \subset X$:

- (4) The nearest point retraction, which takes the pair (x, B) to a point of B closest to x , is coarsely Lipschitz in both x and B (in terms of Hausdorff distance for B), with implicit constants depending only on the hyperbolicity constant of X and the quasiconvexity constants of B .

□

We will apply these properties to curve complexes of surfaces and their subsurfaces, as follows: if A is a finite subset of $\mathcal{M}(S)$ then we let $\text{hull}_S(A)$ denote $\text{hull}_{\mathcal{C}(S)}(A')$, where A' is the set of curves in the bases of the markings in A . Similarly, if $W \subset S$ we let $\text{hull}_W(A)$ denote $\text{hull}_{\mathcal{C}(W)}(\pi_W(A))$ where $\pi_W: \mathcal{M}(S) \rightarrow \mathcal{C}(W)$ is the usual subsurface projection.

5.2. Σ -hulls and their projections

If A is a finite subset of $\mathcal{M}(S)$, and $\epsilon > 0$, we define

$$\Sigma_\epsilon(A) = \{\mu \in \mathcal{M}(S) : d_W(\mu, \text{hull}_W(A)) \leq \epsilon, \quad \forall W \subseteq S\}.$$

Here W varies over all essential subsurfaces of S (including S) and $\text{hull}_W(A)$ is the hyperbolic hull as defined in §5.1. These sets, one for each $\epsilon \geq 0$, are called the *coarse Σ -hulls* or just the *Σ -hulls* of A in $\mathcal{M}(S)$. Usually we will assume ϵ is large enough so that the conclusions of Propositions 5.2 and 5.4 apply.

It is clear that $A \subset \Sigma_\epsilon(A)$ and that $\Sigma_\epsilon(A) \subset \Sigma_{\epsilon'}(A)$ if $\epsilon \leq \epsilon'$, but a priori not much else. (For the reader familiar with the constructions in Masur-Minsky [39], we note one of our motivations for this definition: there exists ϵ_0 such that, if $\epsilon > \epsilon_0$ then $\Sigma_\epsilon(A)$ contains every hierarchy path between points $a, a' \in A$.)

In order to understand Σ -hulls better we will need a family of coarse retractions.

Proposition 5.2. *There exists $\epsilon_0 \geq 0$ depending only on $\xi(S)$ such that for each $\epsilon \geq \epsilon_0$ the following hold. Given a finite set $A \subset \mathcal{M}(S)$ there exists a map*

$$p_A: \mathcal{M} \rightarrow \Sigma_\epsilon(A)$$

which is a coarse retraction. That is,

- (1) $p_A|_{\Sigma_\epsilon(A)}$ is uniformly close to the identity.
- (2) $p_A(x)$ is coarse-Lipschitz not just in x , but jointly in x and in A (using the Hausdorff metric on A).
- (3) For each $W \subseteq S$, let y_W be a nearest point on $\text{hull}_W(A)$ to $\pi_W(x)$. Then

$$d_W(p_A(x), y_W)$$

is uniformly bounded.

The implicit constants depend only on ϵ , $\xi(S)$, and the cardinality of A .

Proof. The proof will be an application of the Consistency Theorem 4.3. Given $x \in \mathcal{M}(S)$, for each $W \subseteq S$ let $y_W = y_W(x, A)$ be a nearest point to $\pi_W(x)$ on $\text{hull}_W(A)$.

Lemma 5.3. *For any $x \in \mathcal{M}(S)$ and finite $A \subset \mathcal{M}(S)$, the tuple $(y_W(x, A))_W$ satisfies the consistency conditions (C1-2) of §4, with constants c_1, c_2 depending only on the cardinality of A .*

Proof. To prove (C1), let $U \pitchfork V$. First, by Lemma 4.2 we have

$$\min(d_U(x, \partial V), d_V(x, \partial U)) < m_0$$

Consider the case where $d_U(x, \partial V) < m_0$; the other case is similar. Now if $d_U(x, y_U) < 2m_0 + 2$, we are done, because $d_U(y_U, \partial V) < 3m_0 + 4$. We have used here that $\text{diam}_U(x)$ is bounded above by 2 — see comments after Lemma 2.11 — and we shall use the same bound several times below.

If $d_U(x, y_U) \geq 2m_0 + 2$ then $d_U(x, \text{hull}_U(A)) \geq 2m_0 + 2$ since y_U was the nearest point to $\pi_U(x)$, and we conclude by the triangle inequality that

$$d_U(\partial V, \text{hull}_U(A)) \geq m_0. \quad (5.1)$$

Now since $y_V \in \text{hull}_V(A)$, there must exist $a, b \in A$ and $a' \in \pi_V(a)$, $b' \in \pi_V(b)$ such that $y_V \in [a', b']$. Now $d_U(\partial V, a)$ and $d_U(\partial V, b)$ are $\geq m_0$ by (5.1), and it follows by Lemma 4.2 that $d_V(\partial U, a) < m_0$ and $d_V(\partial U, b) < m_0$, and so $d_V(a', b') < 2m_0 + 6$. Since $y_V \in [a', b']$ we conclude that one of $d_V(y_V, a')$, $d_V(y_V, b')$ is $< m_0 + 3$, implying that $d_V(y_V, \partial U) < 2m_0 + 5$, and again we are done.

That is, we have shown that (C1) holds with $c_1 = 3m_0 + 5$.

It remains to prove (C2). Let $V \Subset W$, and suppose that $d_W(y_W, \partial V) > 4$. We will bound $d_V(y_V, y_W)$.

Suppose first that $d_V(y_W, x) \geq m_0$. Then by Theorem 2.6, the $\mathcal{C}(W)$ -geodesic $[y_W, \pi_W(x)]$ must pass through a point t within 1 of ∂V . By the assumption that $d_W(y_W, \partial V) > 4$, it follows that $d_W(t, y_W) > 3$ and hence $d_W(t, x) < d_W(x, y_W) - 3$. Now let γ be a $\mathcal{C}(W)$ -geodesic $[\pi_W(a), \pi_W(b)]$ for $a, b \in A$. If γ were to pass within 1 of ∂V then it would pass within 2 of t , so there would be a point of γ which is within $d_W(x, y_W) - 1$ of $\pi_W(x)$. This contradicts the choice of y_W as a closest point to $\pi_W(x)$. We conclude, by Theorem 2.6 that $\text{diam}_V(\gamma) < m_0$, and hence $\text{diam}_V(A) < m_0$.

Moreover, since y_W itself is on such a geodesic, $d_V(y_W, A) < m_0$. Since $y_V \in \text{hull}_V(A)$ we also have $d_V(y_V, A) < m_0$ and we conclude $d_V(y_V, y_W) < 3m_0$.

Now suppose that $d_V(y_W, x) < m_0$. Let $a, b \in A$ be such that $y_W \in [\pi_W(a), \pi_W(b)]$. Now, by our assumption that $d_W(y_W, \partial V) > 4$, we have

that $\pi_W(\partial V)$ may be within distance 1 of either subsegment $[\pi_W(a), y_W]$ or $[y_W, \pi_W(b)]$, but not both. Suppose the former. Then by Theorem 2.6 we have $d_V(y_W, b) < m_0$. This yields that $d_V(x, A) < 2m_0$, and hence that the closest point y_V to $\pi_V(x)$ is within $3m_0$ of $\pi_V(y_W)$.

Hence we have proved (C1) and (C2) both hold with constants $c_1 = 3m_0 + 5$ and $c_2 = 4$. \square

We turn now to the proof of Proposition 5.2. Using Lemma 5.3, the definition and properties of p_A follow directly from Theorem 4.3: given $x \in \mathcal{M}(S)$ and (y_W) as in part (3), Lemma 5.3 tells us that (y_W) satisfies conditions (C1-2) with uniform constants, and hence by Theorem 4.3 there exists $\mu \in \mathcal{M}(S)$ with $d_W(\mu, y_W) < c_3$ for uniform c_3 and all $W \subseteq S$. We take $\epsilon_0 = c_3$, for any $\epsilon \geq \epsilon_0$ we define $p_A(x) \equiv \mu \in \Sigma_\epsilon(A)$, and clearly (3) holds.

Finally, let us show that the rest of the proposition follows from (3). To see (1), let $x \in \Sigma_\epsilon(A)$, i.e., for all W we have $d_W(x, \text{hull}_W(A)) \leq \epsilon$, and so $d_W(x, y_W) \leq \epsilon$. Now if $\mu = p_A(x)$ we have from (3) that $d_W(\mu, y_W)$ is uniformly bounded, and hence we have a uniform bound on $d_W(x, \mu)$. Corollary 2.9 of the quasidistance formula now gives us a bound on $d_{\mathcal{M}}(x, p_A(x))$.

To prove (2), suppose that we have $d_{\mathcal{M}}(x, x') < b$ and $d_H(A, A') < b$, where d_H is Hausdorff distance in \mathcal{M} . The coarse-Lipschitz property of π_W (Lemma 2.11) implies that for any W we have bounds of the form $d_W(x, x') < b'$, and $d_{H, \mathcal{C}(W)}(A, A') < b'$. The latter implies a Hausdorff distance bound

$$d_{H, \mathcal{C}(W)}(\text{hull}_W(A), \text{hull}_W(A')) \leq b''$$

by Lemma 5.1(3). If y'_W is a nearest point to $\pi_W(x')$ in $\text{hull}_W(A')$ then, by Lemma 5.1(4) we obtain a uniform bound on $d_W(y_W, y'_W)$.

But (3) now implies that $d_W(p_A(x), p_{A'}(x'))$ is uniformly bounded for all W . Again the quasidistance formula gives us a uniform bound on $d_{\mathcal{M}}(p_A(x), p_{A'}(x'))$. \square

As a consequence of Proposition 5.2 we obtain the following facts:

Lemma 5.4. *There exists $\epsilon_0 \geq 0$ depending only on $\xi(S)$ such that for all $\epsilon, \epsilon' \geq \epsilon_0$ and all I there exist K, C , and ϵ'' such that if $A, A' \subset \mathcal{M}(S)$ each have cardinality $\leq I$ then*

- (1) $\text{diam}(\Sigma_\epsilon(A)) \leq K \text{diam}(A) + C$
- (2) *If $A' \subset \Sigma_\epsilon(A)$ then $\Sigma_\epsilon(A') \subset \Sigma_{\epsilon''}(A)$.*
- (3) $d_H(\Sigma_\epsilon(A'), \Sigma_\epsilon(A)) \leq K d_H(A', A) + C$.
- (4) $d_H(\Sigma_\epsilon(A), \Sigma_{\epsilon'}(A)) \leq C$

Proof. Parts (1) and (2) follow from the definition of $\Sigma_\epsilon(A)$ and the quasidistance formula.

To prove (3), note by Proposition 5.2 (2) the maps p_A and $p_{A'}$ differ by at most $K d_H(A', A) + C''$, for some K, C'' (with the appropriate dependence). Since $p_{A'}$ is uniformly close to the identity on $\Sigma_\epsilon(A')$, the restriction of p_A to $\Sigma_\epsilon(A')$ must be within $K d_H(A', A) + C'$ of the identity, for some C' . It follows that $\Sigma_\epsilon(A')$ is within $K d_H(A', A) + C$ of $\Sigma_\epsilon(A)$, for some C . The opposite inclusion is obtained in the same way.

To prove (4), we may assume that $\epsilon \leq \epsilon'$ from which it immediately follows that $\Sigma_\epsilon(A) \subset \Sigma_{\epsilon'}(A)$. Consider the two projection maps $p_{A,\epsilon}: \mathcal{M}(S) \rightarrow \Sigma_\epsilon(A)$, $p_{A,\epsilon'}: \mathcal{M}(S) \rightarrow \Sigma_{\epsilon'}(A)$. Given $\mu \in \Sigma_{\epsilon'}(A)$, by applying Proposition 5.2 (3) we obtain a uniform bound on $d_W(y_W, p_{A,\epsilon'}(\mu))$ and on $d_W(y_W, p_{A,\epsilon}(\mu))$ and so also on $d_W(p_{A,\epsilon'}(\mu), p_{A,\epsilon}(\mu))$, over all essential subsurfaces $W \subset S$. Corollary 2.9 then gives a bound on $d(p_{A,\epsilon'}(\mu), p_{A,\epsilon}(\mu))$. Also, Proposition 5.2 (1) gives a bound on $d(\mu, p_{A,\epsilon'}(\mu))$, and since $p_{A,\epsilon}(\mu) \in \Sigma_\epsilon(A)$ we are done. \square

We shall also have use for the following lemma, where m_0 is the constant in Lemma 4.2.

Lemma 5.5. *There exists $m_1 \geq 0$ depending only on $\xi(S)$ such that if $A \subset \mathcal{M}(S)$ is any subset and $W \subset S$ is any essential subsurface satisfying*

$$\text{diam}_{\mathcal{C}(W)}(A) > m_1,$$

then for all essential subsurfaces $U \subseteq S$ with $U \pitchfork \partial W$ we have

$$d_{\mathcal{C}(U)}(\partial W, \text{hull}_U(A)) \leq m_0$$

Proof. Since $U \pitchfork \partial W$ we either have $W \subseteq U$ and $W \not\stackrel{i}{=} U$, or $W \pitchfork U$. We treat these two cases separately.

First, if $W \subset U$ and $W \not\stackrel{i}{=} U$ then Theorem 2.6 immediately implies that

$$d_{\mathcal{C}(U)}(\partial W, \text{hull}_U(A)) \leq 1$$

as long as $\text{diam}_{\mathcal{C}(W)}(A) > B$.

Now, consider the case that $W \pitchfork U$. As long as $\text{diam}_{\mathcal{C}(W)}(A) > 2m_0 + 2$, it follows that there exists $a \in A$ for which $d_{\mathcal{C}(W)}(a, \partial U) \geq m_0$. Then by Lemma 4.2 it follows that $d_{\mathcal{C}(U)}(a, \partial W) < m_0$. \square

6. Contractibility and homology

In this section we prove contractibility of Σ -hulls in $\mathcal{M}_\omega(S)$ (Lemma 6.2), and use this to develop Σ -compatible chains in the cone. This has applications in Sections 7 and 8, as well as providing another proof of Hamenstädt's theorem on the homological dimension of $\mathcal{M}_\omega(S)$.

If \mathbf{A} is a finite set in \mathcal{M}_ω represented by a sequence (A_n) , the Σ -hull of \mathbf{A} , denoted $\Sigma(\mathbf{A})$, is the ultralimit of the coarse Σ -hulls $\Sigma_\epsilon(A_n)$, where ϵ is a fixed constant chosen sufficiently large so that the the lemmas in Section 5 apply.

Note that changing ϵ does not change $\Sigma(\mathbf{A})$, by Lemma 5.4 (4), nor does changing the representatives, by Lemma 5.4 (3). In fact Lemma 5.4 applied in the limit gives:

Lemma 6.1. *For all $I \geq 0$ there exists a constant K , depending also on $\xi(S)$, such that if $\mathbf{A} \subset \mathcal{M}_\omega$ is a set of cardinality $\leq I$ then:*

- (1) $\text{diam}(\Sigma(\mathbf{A})) \leq K \text{diam}(\mathbf{A})$
- (2) *If $\mathbf{A}' \subset \Sigma(\mathbf{A})$ then $\Sigma(\mathbf{A}') \subset \Sigma(\mathbf{A})$*

The retractions p_{A_n} of Proposition 5.2 ultraconverge to a Lipschitz retraction

$$p_{\mathbf{A}} : \mathcal{M}_\omega \rightarrow \Sigma(\mathbf{A})$$

whose Lipschitz constant depends only on $\xi(S)$ and the cardinality of \mathbf{A} . Moreover, Proposition 5.2 implies that $p_{\mathbf{A}}$ is jointly continuous in its arguments and in the points of \mathbf{A} . With this we can establish:

Lemma 6.2. *$\Sigma(\mathbf{A})$ is contractible.*

Proof. First we note that $\Sigma(\mathbf{A})$ is path-connected: $\mathcal{M}_\omega(S)$ is path-connected since it is the asymptotic cone of a path-metric space. Hence given $\mathbf{a}, \mathbf{b} \in \Sigma(\mathbf{A})$, let $\gamma(t)$ be a path connecting them and note that $p_{\mathbf{A}} \circ \gamma$ is a path in $\Sigma(\mathbf{A})$ connecting them.

Now write $\mathbf{A} = \{\mathbf{a}_0, \mathbf{a}_1, \dots, \mathbf{a}_k\}$, and for $j = 1, \dots, k$ let $\mathbf{a}_j(t)$ be a path in $\Sigma(\mathbf{A})$ from \mathbf{a}_0 to \mathbf{a}_j , where $\mathbf{a}_j(0) = \mathbf{a}_0$ and $\mathbf{a}_j(1) = \mathbf{a}_j$.

Let $\mathbf{A}_t = \{\mathbf{a}_0, \mathbf{a}_1(t), \dots, \mathbf{a}_k(t)\}$ for $t \in [0, 1]$, and let p_t be the retraction from \mathcal{M}_ω to $\Sigma(\mathbf{A}_t)$. p_t varies continuously in t , takes values within $\Sigma(\mathbf{A})$, and we note that p_1 restricted to $\Sigma(\mathbf{A})$ is the identity while p_0 is a constant. Hence $\Sigma(\mathbf{A})$ is contractible. \square

6.1. Σ -compatible chains and homological dimension

In this subsection we use Σ -hulls as a device to control singular chains in $\mathcal{M}_\omega(S)$, in terms of what we call Σ -compatible chains. With these we compute the homological dimension from the result of [7] that the topological (covering) dimension of compact subsets of $\mathcal{M}_\omega(S)$ is bounded by $\xi(S)$. We also recall a local homology theorem of Kleiner-Leeb [35] (Theorem 6.8 below) and its corollary 6.9, which we will use in Sections 7 and 8 to control the support of embedded top-dimensional manifolds in $\mathcal{M}_\omega(S)$ in terms of Σ -compatible chains.

A *polyhedron* is a finite simplicial complex, and a *polyhedral pair* (P, Q) consists of a polyhedron P and a subcomplex Q . A *polyhedral n -chain* in a space X is a continuous map from an n -dimensional polyhedron to X , where the domain is equipped with a specified orientation and coefficient on each n -simplex. A *polyhedral n -cycle* is a polyhedral n -chain such that, in the simplicial chain complex of the domain, the linear combination of the n -simplices has zero

boundary. Polyhedral chains and cycles in X represent singular chains and cycles in the traditional sense, by restricting the map to individual simplices and taking the indicated formal linear combination. Every homology class in X can be represented by a polyhedral cycle.

A continuous map $f: P \rightarrow \mathcal{M}_\omega$ from a polyhedron to \mathcal{M}_ω is Σ -compatible if for each face $\tau \subset P$,

$$f(\tau) \subset \Sigma(f(\tau^{(0)})),$$

where $\tau^{(0)}$ denotes the 0-skeleton of τ . By applying Lemma 5.4 (1) it follows that if $f: P \rightarrow \mathcal{M}_\omega$ is Σ -compatible then for every face $\tau \subset P$,

$$\text{diam}(f(\tau)) \leq \text{diam}(\Sigma(f(\tau^{(0)}))) \leq C \text{diam}(f(\tau^{(0)})), \quad (6.1)$$

where the constant $C = C(\dim \tau)$ depends explicitly on $\dim \tau = \#\text{vertices}(\tau) - 1$, and also depends implicitly on

Lemma 6.3. *Suppose (P, Q) is a finite dimensional polyhedral pair, where the zero skeleton of Q coincides with the zero skeleton of P . Then any Σ -compatible map $f_0: Q \rightarrow \mathcal{M}_\omega$ can be extended to a Σ -compatible map $f: P \rightarrow \mathcal{M}_\omega$.*

Proof. The map f may be constructed by induction on the relative k -skeleton using the contractibility of hulls. \square

Lemma 6.4. *If $\epsilon > 0$ and $f_0: P \rightarrow \mathcal{M}_\omega$ is a map from a finite polyhedron to \mathcal{M}_ω , then there is a map $f_1: P \rightarrow \mathcal{M}_\omega$ such that*

1. f_1 factors through a polyhedron of dimension $\leq \xi(S)$.
2. $d(f_0, f_1) < \epsilon$.

(Here $d(f, g) = \sup_{x \in P} d(f(x), g(x))$).

Proof. Pick $\rho > 0$.

Let $Y := f_0(P) \subset \mathcal{M}_\omega$. Since the topological dimension of Y is $\leq \xi(S)$ by [7], there is an open cover $\mathbf{U} = \{U_i\}_{i \in I}$ of Y such that $P' := \text{Nerve}(\mathbf{U})$ has dimension at most $\xi(S)$, and $\text{diam}(U_i) < \rho$ for all $i \in I$. Let $\{\phi_i: Y \rightarrow [0, 1]\}_{i \in I}$ be a partition of unity subordinate to \mathbf{U} , and $\phi: Y \rightarrow P'$ be the map with barycentric coordinates given by the ϕ_i 's.

Next, for each $i \in I$, pick $x_i \in U_i$, and using Lemma 6.3 construct a Σ -compatible map $\alpha: P' \rightarrow \mathcal{M}_\omega$ with the property that $\alpha(U_i) = x_i$ (recall that the vertex set of $\text{Nerve}(\mathbf{U})$ consists of elements of \mathbf{U}).

Set $f_1 := \alpha \circ \phi \circ f_0 \rightarrow \mathcal{M}_\omega$.

We now estimate $d(f_0, f_1)$.

Pick $x \in P$. If $\phi \circ f_0(x)$ lies in an open face $\tau \subset P'$ whose vertices are U_{i_1}, \dots, U_{i_k} , then $f_0(x) \in U_{i_1} \cap \dots \cap U_{i_k}$, and

$$f_1(x) \in \Sigma(\{x_{i_1}, \dots, x_{i_k}\}).$$

Therefore for a constant C depending only on $\xi(S)$ we have:

$$\begin{aligned} d(f_0(x), f_1(x)) &\leq d(f_0(x), x_{i_1}) + d(x_{i_1}, f_1(x)) \\ &\leq \rho + C \operatorname{diam}(\{x_{i_1}, \dots, x_{i_k}\}) \\ &\leq \rho + 2C\rho. \end{aligned} \tag{6.2}$$

So when $\rho < \frac{\epsilon}{1+2C}$ we will have $d(f_0, f_1) < \epsilon$. \square

Lemma 6.5. *Let P be a finite polyhedron. Given a pair of maps $f_0, f_1: P \rightarrow \mathcal{M}_\omega$, there is a homotopy $\{f_t\}_{t \in [0,1]}$ from f_0 to f_1 whose tracks have diameter $< C d(f_0, f_1)$, where $C = C(\dim P)$.*

Proof. Pick $\rho > 0$. By subdividing P we may assume without loss of generality that for $i \in \{0, 1\}$ and every face τ of P ,

$$\operatorname{diam}(f_i(\tau)) < \rho.$$

Let $P = P_1, P_2, \dots, P_k, \dots$ be a sequence of successive barycentric subdivisions of P , so the mesh size tends to zero. For $i \in \{0, 1\}$, $k \in \mathbb{Z}_+$, let $f_{i,k}: P_k \rightarrow \mathcal{M}_\omega$ be a Σ -compatible map agreeing with f_i on the 0-skeleton of P_k . Since f_i is uniformly continuous, the diameter estimate (6.1) implies that $f_{i,k}$ converges uniformly to f_i as $k \rightarrow \infty$.

We will construct the homotopy from f_0 to f_1 as an infinite concatenation of homotopies

$$f_0 \dots \overset{H_{0,3}}{\sim} f_{0,3} \overset{H_{0,2}}{\sim} f_{0,2} \overset{H_{0,1}}{\sim} f_{0,1} \overset{H}{\sim} f_{1,1} \overset{H_{1,1}}{\sim} f_{1,2} \overset{H_{1,2}}{\sim} f_{1,3} \dots f_1.$$

The homotopy $f_{i,k} \overset{H_{i,k}}{\sim} f_{i,k+1}$ is constructed as follows. Triangulate $P \times [0, 1]$ such that the 0-skeleton lies in $P \times \{0, 1\}$, and the induced triangulation of $P \times \{j\}$ agrees with P_{k+j} , for $j \in \{0, 1\}$. Now apply Lemma 6.3 to get a homotopy from $f_{i,k}$ to $f_{i,k+1}$. The homotopy H is constructed similarly.

Now consider the track of the point $x \in P$ during the homotopy $H_{i,k}$. The point $x \in P$ lies in some open simplex τ_k of P_k . Let $\operatorname{Sd} \tau_k$ be the barycentric subdivision of τ_k , a subcomplex of P_{k+1} . By the uniform continuity of f_i and the fact that $f_{i,k} \rightarrow f_i$ uniformly, it follows that the diameter of $f_{i,k}(\tau_k) \cup f_{i,k+1}(\operatorname{Sd} \tau_k)$ tends to zero as $k \rightarrow \infty$. Now for every $t \in [0, 1]$, the point $(x, t) \in P \times [0, 1]$ lies in a face of the subdivision of $P \times [0, 1]$ used to construct $H_{i,k}$, and this face has vertices in $(\tau_k \times \{0\}) \cup (\operatorname{Sd} \tau \times \{1\})$. By Lemma 6.3, we get $d(H_{i,k}(x, t), f_i(x)) < \delta_k$, where $\delta_k < \rho$ and $\delta_k \rightarrow 0$ as $k \rightarrow \infty$. It follows that the concatenation of

$$H_{i,j}, H_{i,j+1}, \dots$$

has tracks of diameter tending to zero as $j \rightarrow \infty$, yielding a homotopy $f_{i,1} \sim f_i$ whose tracks have diameter $< C_1\rho$.

Similar estimates imply that the tracks of H have diameter $< C_2(d(f_0, f_1) + \rho)$. So if ρ is sufficiently small, we obtain the desired homotopy. \square

We now give some corollaries of Lemmas 6.4 and 6.5.

Our first corollary produces acyclic sets in $\mathcal{M}_\omega(S)$. A set $X \subset \mathcal{M}_\omega(S)$ is Σ -convex if, for any finite set $\mathbf{A} \subset X$ the hull $\Sigma(\mathbf{A})$ is in X as well.

Corollary 6.6. *If $X \subset \mathcal{M}_\omega(S)$ is open and Σ -convex then X is acyclic.*

Notice that, as a consequence, $\mathcal{M}_\omega(S)$ is itself acyclic. But contractibility of $\mathcal{M}_\omega(S)$ was already known, as a consequence of the fact that $\mathcal{MCG}(S)$ is automatic [43], by using a folk theorem which says that the combing lines of an automatic structure induce a contraction in the asymptotic cone.

Proof of Corollary 6.6. We claim that any polyhedral cycle $f: P \rightarrow X$ can be refined and then approximated by a Σ -compatible cycle, and openness allows us to do this within X . To see how, refine P until the mesh size is sufficiently small, apply Lemma 6.3 to the pair (Q, Q^0) where Q is the refined polyhedron and Q^0 is its 0-skeleton, and then use Lemma 6.1 to show that the new map $f': Q \rightarrow \mathcal{M}_\omega(S)$ is sufficiently close to the old to still be in X . Moreover Lemma 6.5 tells us the new and old maps are homotopic with similar control, so that the homotopy may be made to lie in X . The old and new cycles therefore represent the same homology class.

Now let $\mathcal{C}(Q)$ be the cone of Q , the join of Q with a single point p . Extend f' to $Q \cup p$ by taking $f'(p)$ to be in $f'(Q^0)$. Applying Lemma 6.3 extend f' to a Σ -compatible map on $\mathcal{C}(Q)$. By Σ -convexity this polyhedral chain lies in X , so our original cycle bounds in X , and therefore X is acyclic. \square

Our next corollary computes the homological dimension of $\mathcal{M}_\omega(S)$. Note that in Hamenstädt's approach [26] the homological statement comes directly.

Corollary 6.7. *If (U, V) is an open pair in \mathcal{M}_ω , then $H_k(U, V) = \{0\}$ for all $k > \xi(S)$.*

Proof. Pick $[c] \in H_k(U, V)$. Then there is a finite polyhedral pair (P, Q) and a continuous map of pairs $f_0: (P, Q) \rightarrow (U, V)$ such that

$$[c] \in \text{Im}(H_k(P, Q) \xrightarrow{f_{0*}} H_k(U, V)).$$

Pick $\epsilon > 0$. Applying Lemma 6.4, we obtain a continuous map $f_1: P \rightarrow \mathcal{M}_\omega$ with $d(f_0, f_1) < \epsilon$, such that f_1 factors through a polyhedron P' of dimension at most $\xi(S)$. By Lemma 6.5 there is a homotopy $\{f_t\}_{t \in [0,1]}$ whose tracks have diameter $< C\epsilon$, where $C = C(\dim P)$.

If ϵ is sufficiently small, then f_1 will induce a map of pairs $(P, Q) \rightarrow (U, V)$, and the homotopy $\{f_t\}$ will be a homotopy of maps of pairs, so that f_0 and f_1 induce the same map $H_k(P, Q) \rightarrow H_k(U, V)$. But since f_1 factors through a polyhedron P' of dimension $\leq \xi(S)$, by subdividing P' if necessary we can arrange that f_1 factors as $(P, Q) \rightarrow (P', Q') \rightarrow (U, V)$, where (P', Q') is a

polyhedral pair of dimension $\leq \xi(S)$. This implies that $f_{1*} = 0$. Hence $[c] = 0$. \square

We are now in a position to apply the following local homology results of Kleiner-Leeb [35], which we will be using in the proof of Theorems 7.10 and 8.7.

Theorem 6.8. *Let X be a contractible metric space and suppose $H_k(U, V) = 0$ for any open pair $V \subset U \subset X$ and $k > n$. If $M \subset X$ is an embedded n -manifold then*

$$H_n(M, M - p) \rightarrow H_n(X, X - p)$$

is injective for any $p \in M$.

Corollary 6.9. *Let X be a contractible metric space and suppose $H_k(U, V) = 0$ for any open pair (U, V) and $k > n$. Let $M \subset X$ be an oriented compact n -manifold with boundary, and let C be a singular chain in X , such that $\partial C = \partial M$. Then $M \subset C$.*

By Corollary 6.7, we will be able to apply Corollary 6.9 in the setting of $\xi(S)$ -dimensional manifolds in $\mathcal{M}_\omega(S)$.

7. Separation properties

In this section we develop the notion of *jets*, which are local structures in the cone corresponding to sequences of geodesics in subsurface complexes. Projections to a jet serve to control separation properties in the cone. The two main results of the section are Theorem 7.2 and Theorem 7.7, which are concerned with separation properties of *microscopic* and *macroscopic* jets, respectively. Much of the technical work is done in Lemma 7.4. Section 7.4 provides a brief digression, where we deduce information about the tree-graded structure of $\mathcal{M}_\omega(S)$ as an application of microscopic jets.

7.1. Jets

Recall the following definition from [39]. Consider a finite type surface Y . If $\xi(Y) \geq 2$, a *tight geodesic* in $\mathcal{C}(Y)$ is a sequence of simplices $\sigma = (w_0, \dots, w_n)$ such that any selection of vertices $v_i \in w_i$ yields a geodesic in the 1-skeleton of $\mathcal{C}(Y)$, and such that for each $1 \leq i \leq n - 1$, the curve system w_i is the boundary of the subsurface filled by w_{i-1} and w_{i+1} . If $\xi(Y) = 1$, every geodesic in $\mathcal{C}(Y)$ is considered to be tight. If $Y \subset S$ is an essential annulus, then every geodesic in $\mathcal{C}(Y)$ is considered to be tight as long as it satisfies a technical finiteness condition on the endpoints of arcs representing the vertices. It is shown in [39] that any two vertices in $\mathcal{C}(Y)$ can be joined by a tight geodesic, and there are only finitely many possibilities. When the sequence (w_0, \dots, w_n) is understood, we use the shorthand notation $[w_0, w_n]$, and we also refer to $[w_i, w_j] = (w_i, \dots, w_j)$ as a *subsegment* of $[w_0, w_n]$.

Let $a, b \in \mathcal{M}(S)$, let $W \subseteq S$ be a connected essential subsurface, and let g be a tight geodesic in $\mathcal{C}(W)$ from an element of $\pi_{\mathcal{C}(W)}(a)$ to an element of $\pi_{\mathcal{C}(W)}(b)$. If $\sigma = [\alpha, \beta]$ is a subsegment of g , we call (σ, a, b) a *tight triple* supported in W . Let $|\sigma|$ denote the length of σ in $\mathcal{C}(W)$. Although we have suppressed the geodesic g from the notation for a tight triple, when we need to refer to it we shall call it the *ambient geodesic*.

We also associate to the triple (σ, a, b) the following pair of points in $\mathcal{M}(W)$: $\iota(\sigma, a, b) = \alpha \rfloor \pi_{\mathcal{M}(W)}(a)$, the *initial marking* of the triple; and $\tau(\sigma, a, b) = \beta \rfloor \pi_{\mathcal{M}(W)}(b)$, the *terminal marking*. Up to the usual bounded ambiguity one can think of ι as $\alpha \cup \pi_{\mathcal{M}(W_\alpha)}(a)$, where W_α denotes the union of $W \setminus \alpha$ with the annuli whose cores are the curves of α ; and similarly for τ .

We define

$$\|\sigma\|_{(W,a,b)} = \text{dist}_{\mathcal{M}(W)}(\iota(\sigma, a, b), \tau(\sigma, a, b)).$$

Using Theorem 2.8 we can establish the following properties of this notion of size:

Lemma 7.1. *Given $\sigma = [\alpha, \beta], W, a, b$ as above, letting $\Phi(\sigma)$ denote the set of subsurfaces $Y \Subset W$ that do not overlap some simplex of σ , we have*

$$(1) \quad \|\sigma\|_{(W,a,b)} \approx d_{\mathcal{C}(W)}(\alpha, \beta) + \sum_{Y \in \Phi(\sigma)} \left\{ \left\{ d_{\mathcal{C}(Y)}(a, b) \right\} \right\}_A$$

(2) *If σ is written as a concatenation of subintervals $\sigma_1 * \dots * \sigma_k$, then*

$$\|\sigma\|_{(W,a,b)} \approx \sum_i \|\sigma_i\|_{(W,a,b)}$$

where the constant A and the constants of approximation depend only on $\xi(W)$.

Proof. Let $\iota = \iota(\sigma, a, b)$ and $\tau = \tau(\sigma, a, b)$. To prove (1), first recall that Theorem 2.8 gives us, for large enough A and uniform constants of approximation (depending on A), that

$$d_{\mathcal{M}(W)}(\iota, \tau) \approx \sum_{Y \Subset W} \left\{ \left\{ d_{\mathcal{C}(Y)}(\iota, \tau) \right\} \right\}_A$$

On the other hand, Theorem 2.6 gives a constant B such that, if $Y \Subset W$, $Y \not\stackrel{j}{\subseteq} W$, and if Y overlaps every simplex of σ , then $\text{diam}_{\mathcal{C}(Y)}(\sigma) \leq B$; in particular since α, β are in σ , and since α, β are contained in ι and τ respectively, we get

$$d_{\mathcal{C}(Y)}(\iota, \tau) \leq B$$

Thus, if the threshold constant A is raised above B all of these terms drop out of the sum, leaving the $Y = W$ term, and what is almost the summation in (1), indexed by $Y \in \Phi(\sigma)$, except with $d_{\mathcal{C}(Y)}(\iota, \tau)$ in place of $d_{\mathcal{C}(Y)}(a, b)$.

Now consider $Y \Subset W$ which does not overlap some simplex of σ . By tightness of the ambient geodesic g containing σ , the set of simplices in g not overlapping Y is a contiguous sequence of at most 3 simplices. If Y overlaps β then the rest

of g between β and $\pi_{\mathcal{C}(W)}(b)$ consists of simplices overlapping Y and so, since τ contains β , Theorem 2.6 implies

$$d_{\mathcal{C}(Y)}(\tau, b) \leq B.$$

If Y does not overlap β then $Y \subset W_\beta$. By definition the restriction of τ to W_β is $\pi_{\mathcal{M}(W_\beta)}(b)$, and again we have a uniform bound on $d_{\mathcal{C}(Y)}(\tau, b)$. The same logic yields a uniform bound on $d_{\mathcal{C}(Y)}(\iota, a)$. Thus, at the cost of again raising the threshold, we can replace ι, τ in the sum by a, b — thus completing the proof of (1).

To prove (2), we simply apply the approximation of (1) to each σ_i separately and sum, noting that any $Y \subset W$ there are at most 3 (successive) simplices disjoint from it, and hence it can be in at most 4 different $\Phi(\sigma_i)$. This bounds the overcounting by a factor of 4, and gives the estimate. \square

A *jet*, denoted J , is a quadruple $(\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$, where (σ_n, a_n, b_n) are tight triples with σ_n supported in W_n , and we assume that \bar{a} and \bar{b} have ultralimits in $\mathcal{M}_\omega(S)$ (i.e., that they do not go to ∞ faster than linearly). We refer to \bar{W} as the *support surface* of the jet J . The sequence of initial points $\iota_n = \iota(\sigma_n, a_n, b_n)$ defines a point $\iota_\omega(\bar{\sigma}, \bar{a}, \bar{b}) \in \mathcal{M}_\omega(\bar{W})$, which we will call the *basepoint* of the jet and denote $\iota(J)$ or just ι when J is understood. Similarly one obtains $\tau(J)$.

Call a jet *microscopic* if $\|\sigma_n\|_{(W_n, a_n, b_n)}$ grows sublinearly — that is, if

$$\frac{1}{s_n} \|\sigma_n\|_{(W_n, a_n, b_n)} \rightarrow_\omega 0.$$

A jet J is *macroscopic* if it is not microscopic, which occurs if and only if $\iota(J) \neq \tau(J)$. Often we write $\|\sigma_n\|_J$ to denote $\|\sigma_n\|_{(W_n, a_n, b_n)}$.

7.2. Projection and separation properties of microscopic jets

Let $J = (\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$ be a microscopic jet with basepoint $\iota \in \mathcal{M}_\omega(\bar{W})$. As in Section 3 we have product regions

$$\mathcal{Q}(\partial W_n) \cong \mathcal{M}(W_n) \times \mathcal{M}(W_n^c)$$

which give rise in the cone to

$$\mathcal{Q}_\omega(\partial \bar{W}) \cong \mathcal{M}_\omega(\bar{W}) \times \mathcal{M}_\omega(\bar{W}^c).$$

We let $\mathcal{L}_n(J)$ denote the slice $\mathcal{Q}(\iota_n \cup \partial W_n)$, which by Proposition 3.1 can be identified with $\{\iota_n\} \times \mathcal{M}(W_n^c)$. In the cone we get

$$\mathcal{L}_\omega(J) = \mathcal{Q}_\omega(\iota \cup \partial \bar{W}) \cong \{\iota\} \times \mathcal{M}_\omega(\bar{W}^c).$$

Applying Lemma 3.5, the locally compact dimension of $\mathcal{L}_\omega(J)$ equals $\delta(\bar{W}^c)$ which, by applying the codimension formula, Proposition 2.2, equals $\xi(S) - \delta(\bar{W}) = \xi(S) - \xi(\bar{W})$.

Denote by $\pi_{\sigma_n} : \mathcal{M}(S) \rightarrow \sigma_n$ the composition of projection $\mathcal{M}(S) \rightarrow \mathcal{C}(W_n)$ with closest point projection $\mathcal{C}(W_n) \rightarrow \sigma_n$.

Projection equivalence. In terms of the jet J we can define a relation on sequences (x_n) in $\mathcal{M}(S)$ as follows. Say that $(x_n) \sim_{\bar{\sigma}} (x'_n)$ if

$$d_{\mathcal{C}(W_n)}(\pi_{\sigma_n}(x_n), \pi_{\sigma_n}(x'_n))$$

is bounded for ω -a.e. n . It is immediate that this is an equivalence relation on the ultraproduct $\mathcal{M}(\bar{S})$. We will deduce the following stronger result.

Theorem 7.2. *For any microscopic jet J , the relation $\sim_{\bar{\sigma}}$ descends to an equivalence relation on $\mathcal{M}_{\omega}(S) \setminus \mathcal{L}_{\omega}(J)$. Moreover, every equivalence class is open.*

This theorem is a consequence of the following more quantitative statement:

Lemma 7.3. *There exists $C > 0$ such that for any microscopic jet J , if (ξ_n) and (ξ'_n) are sequences in $\mathcal{M}(S)$ representing $\xi, \xi' \in \mathcal{M}_{\omega}(S)$, and if $(\xi_n), (\xi'_n)$ are inequivalent under $\sim_{\bar{\sigma}}$, then*

$$d(\xi, \xi') \geq C d(\xi, \mathcal{L}_{\omega}(J)).$$

We now show how the lemma implies the theorem.

Proof of Theorem 7.2. If $\xi = \xi'$ then Lemma 7.3 implies either $(\xi_n) \sim_{\bar{\sigma}} (\xi'_n)$ or $\xi \in \mathcal{L}_{\omega}(J)$. Hence in the complement of $\mathcal{L}_{\omega}(J)$ the equivalence relation $\sim_{\bar{\sigma}}$ descends to an equivalence relation in the asymptotic cone.

Further, if $\xi \notin \mathcal{L}_{\omega}(J)$ then Lemma 7.3 implies that there is a positive radius neighborhood of ξ consisting of points represented by sequences which are $\sim_{\bar{\sigma}}$ equivalent to ξ . Hence equivalence classes are open. \square

Lemma 7.3 is an immediate consequence of the following stronger statement, which will have other applications in what follows:

Lemma 7.4. *There exist $K, C > 0$ such that for any microscopic jet J , if (ξ_n) and (ξ'_n) are sequences in $\mathcal{M}(S)$ representing points $\xi, \xi' \in \mathcal{M}_{\omega}(S)$, and if for ω -a.e. n we have $d_{\mathcal{C}(W_n)}(\pi_{\sigma_n}(\xi_n), \pi_{\sigma_n}(\xi'_n)) > K$, then*

$$d(\xi, \xi') \geq C d(\xi, \mathcal{L}_{\omega}(J)).$$

Proof. Proposition 3.1 (2) gives us the following estimate on distance to $\mathcal{L}_n(J)$:

$$\begin{aligned} d_{\mathcal{M}(S)}(\xi_n, \mathcal{L}_n(J)) &\approx \sum_{Y \pitchfork (\iota_n \cup \partial W_n)} \{ \{ d_{\mathcal{C}(Y)}(\xi_n, \iota_n \cup \partial W_n) \} \}_A \\ &\approx \sum_{Y \pitchfork W_n} \{ \{ d_{\mathcal{C}(Y)}(\xi_n, \iota_n) \} \}_A + \sum_{Y \pitchfork \partial W_n} \{ \{ d_{\mathcal{C}(Y)}(\xi_n, \partial W_n) \} \}_A \end{aligned}$$

where A is any sufficiently large threshold and the approximation constants depend only on A .

For each Y indexing this sum, we will show an inequality of the form

$$d_{\mathcal{C}(Y)}(\xi'_n, \xi_n) \geq d_{\mathcal{C}(Y)}(\xi_n, \iota_n) - d_{\mathcal{C}(Y)}(\iota_n, \tau_n) - q \quad (7.1)$$

if $Y \Subset W_n$, and of the form

$$d_{\mathcal{C}(Y)}(\xi'_n, \xi_n) \geq d_{\mathcal{C}(Y)}(\xi_n, \partial W_n) - q \quad (7.2)$$

if $Y \pitchfork \partial W_n$, where q is a uniform constant. Since the left hand sides of these inequalities are terms in the quasidistance formula for $d_{\mathcal{M}(S)}(\xi_n, \xi'_n)$, we will obtain (with the usual threshold adjustment)

$$d_{\mathcal{M}(S)}(\xi_n, \xi'_n) \geq p' d_{\mathcal{M}(S)}(\xi_n, \mathcal{L}_n(J)) - p'' d_{\mathcal{M}(W_n)}(\iota_n, \tau_n) - q' \quad (7.3)$$

where p', p'', q' are additional constants. This will be sufficient, since by assumption $\lim_{\omega} d_{\mathcal{M}(S)}(\iota_n, \tau_n)/s_n = \lim_{\omega} \|\sigma_n\|_J/s_n = 0$, and hence the second term disappears in the asymptotic cone. We proceed to establish (7.1) and (7.2).

Let

$$x_n = \pi_{\mathcal{C}(W_n)}(\xi_n), \quad x'_n = \pi_{\mathcal{C}(W_n)}(\xi'_n),$$

and

$$z_n = \pi_{\sigma_n}(\xi_n), \quad z'_n = \pi_{\sigma_n}(\xi'_n).$$

Let $h_n = [x_n, z_n]$ and $h'_n = [x'_n, z'_n]$ be $\mathcal{C}(W_n)$ -geodesic segments. Because z_n is a nearest point to x_n on σ_n (and similarly for z'_n and x'_n), and $\mathcal{C}(W_n)$ is δ -hyperbolic, there is a constant K_δ such that, if $d(z_n, z'_n) > K_\delta$, the union $T_n = \sigma_n \cup h_n \cup h'_n$ can be considered as a finite tree, and the distance function of $\mathcal{C}(W_n)$ restricted to T_n is approximated by the distance function along the tree, up to some additive error δ' .

In the case that $Y \stackrel{i}{\equiv} W_n$, we immediately find that

$$\begin{aligned} d_{\mathcal{C}(W_n)}(x'_n, x_n) &\geq d_{\mathcal{C}(W_n)}(x_n, z_n) - \delta' \\ &\geq d_{\mathcal{C}(W_n)}(x_n, \iota_n) - \text{diam}_{\mathcal{C}(W_n)}(\sigma_n) - \delta'. \end{aligned} \quad (7.4)$$

which is (7.1) in this case.

Consider next the case that $Y \Subset W_n$ and $Y \not\stackrel{i}{\equiv} W_n$. Let B be the bound in Theorem 2.6. Suppose first that ∂Y is disjoint from a radius 1 neighborhood of h_n in $\mathcal{C}(W_n)$. Then $z_n \pitchfork Y$, and $d_{\mathcal{C}(Y)}(x_n, z_n) \leq B$. Moreover, ∂Y can only be disjoint from simplices on one side of z_n in σ_n (not both) since σ_n is a tight geodesic. It follows that $\pi_Y(z_n)$ is within B of either $\pi_Y(\iota_n)$ or $\pi_Y(\tau_n)$, and hence

$$\min\{d_{\mathcal{C}(Y)}(x_n, \iota_n), d_{\mathcal{C}(Y)}(x_n, \tau_n)\} \leq 2B.$$

It follows that

$$d_{\mathcal{C}(Y)}(x'_n, x_n) \geq 0 \geq d_{\mathcal{C}(Y)}(x_n, \iota_n) - d_{\mathcal{C}(Y)}(\iota_n, \tau_n) - 2B - 3,$$

which again gives (7.1).

Now suppose that ∂Y intersects a radius 2 neighborhood of h_n . Assuming $K > \max(K_\delta, 2\delta' + 4)$, and using the remark above about the tree T_n , it

follows that ∂Y is disjoint from a radius 2 neighborhood of h'_n . Hence the same argument as above gives

$$\min\{d_{\mathcal{C}(Y)}(x'_n, \iota_n), d_{\mathcal{C}(Y)}(x'_n, \tau_n)\} \leq 2B.$$

If $d_{\mathcal{C}(Y)}(x'_n, \iota_n) \leq 2B$ then the triangle inequality gives

$$d_{\mathcal{C}(Y)}(x_n, x'_n) \geq d_{\mathcal{C}(Y)}(x_n, \iota_n) - 2B - 3$$

and if $d_{\mathcal{C}(Y)}(x'_n, \tau_n) \leq 2B$ then the triangle inequality gives

$$d_{\mathcal{C}(Y)}(x_n, x'_n) \geq d_{\mathcal{C}(Y)}(x_n, \iota_n) - d_{\mathcal{C}(Y)}(\iota_n, \tau_n) - 2B - 6.$$

Either way this again gives us (7.1).

Now consider the case when $Y \pitchfork \partial W_n$. We may assume $d_{\mathcal{C}(Y)}(\xi_n, \partial W_n) > A$. As in Section 4, for essential subsurfaces U and partial markings γ of S , for each $k > 0$ define the relation $U \ll_k \gamma$ by

$$U \pitchfork \gamma \text{ and } d_{\mathcal{C}(U)}(\xi_n, \gamma) \geq k c_1 + 4$$

where c_1 is a constant which is chosen as follows. As noted in Section 4, the consistency conditions hold for $(\pi_U(\xi_n))$ for any sufficiently large constants c_1, c_2 . Without loss of generality, we will assume $c_1 > \max\{c_2, m_0, B\}$, where m_0 is the constant given by Lemma 4.2 and B is the constant given by Theorem 2.6.

We may assume $A > 4(c_1 + 4)$, and setting $k = \lfloor d_{\mathcal{C}(Y)}(\xi_n, \partial W_n)/(c_1 + 4) \rfloor \geq 4$, we have

$$Y \ll_k W_n.$$

Moreover we have $d_{\mathcal{C}(W_n)}(\xi_n, \xi'_n) > d_{\mathcal{C}(W_n)}(z_n, z'_n) - \delta' > K - \delta'$, so assuming $K - \delta' > 2(c_1 + 4)$ we get

$$W_n \ll_2 \xi'_n.$$

Now Lemma 4.4 implies that

$$Y \ll_{k-1} \xi'_n$$

so in particular

$$d_{\mathcal{C}(Y)}(\xi_n, \xi'_n) \geq (\lfloor d_{\mathcal{C}(Y)}(\xi_n, \partial W_n)/(c_1 + 4) \rfloor - 1)(c_1 + 4) \geq d_{\mathcal{C}(Y)}(\xi_n, \partial W_n) - 2(c_1 + 4). \quad (7.5)$$

This gives us (7.2). \square

Finding microscopic jets. The next lemma constructs microscopic jets having properties that we will utilize in Sections 7.4 and 8.

Lemma 7.5. *Let $a_\omega, b_\omega \in \mathcal{M}_\omega(\bar{S})$ be represented by $\bar{a}, \bar{b} \in \mathcal{M}(\bar{S})$, let (W_n) be a sequence of connected essential subsurfaces, and suppose that $d_{\mathcal{C}(W_n)}(a_n, b_n) \rightarrow_\omega \infty$. Then there exists a microscopic jet $J = (\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$ such that*

$$\bar{a} \not\sim_{\bar{\sigma}} \bar{b}$$

Proof. Let $\mu_n = \pi_{\mathcal{M}(W_n)}(a_n)$ and $\nu_n = \pi_{\mathcal{M}(W_n)}(b_n)$. Let ℓ_n be a tight $\mathcal{C}(W_n)$ geodesic between μ_n and ν_n .

We use a counting argument to produce a sequence of subsegments σ_n of ℓ_n with $|\sigma_n| \rightarrow \infty$ but $\|\sigma_n\|_{(W_n, a_n, b_n)}$ growing sublinearly. Let $f(n)$ be an integer valued function going to $+\infty$ more slowly than $|\ell_n|$, meaning that $f(n) \rightarrow +\infty$ but $|\ell_n|/f(n) \rightarrow +\infty$. Divide ℓ_n into $f(n)$ subsegments, each of length between $(|\ell_n|/f(n)) - 1$ and $(|\ell_n|/f(n)) + 1$. The sum of the $\|\cdot\|_{(W_n, a_n, b_n)}$ -sizes of these subsegments is approximated by $\|\ell_n\|_{(W_n, a_n, b_n)}$ up to bounded multiple by part (2) of Lemma 7.1, and this in turn is bounded by a multiple of $d_{\mathcal{M}(S)}(\mu_n, \nu_n)$ by part (1) of Lemma 7.1, and hence by a multiple of s_n . Therefore there must be a fixed C such that there is, for ω -a.e. n , a subsegment σ_n with $\|\sigma_n\|_{(W_n, a_n, b_n)} \leq C s_n / f(n)$.

Sublinear growth of $\|\sigma_n\|_{(W_n, a_n, b_n)}$ implies that $J = (\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$ is a microscopic jet. By construction, a_n and b_n project to opposite ends of σ_n , and therefore \bar{a}, \bar{b} are inequivalent under $\sim_{\bar{\sigma}}$. \square

7.3. Linear/sublinear decomposition of macroscopic jets

For macroscopic jets, the way in which the linear growth happens turns out to be important.

Consider a macroscopic jet $J = (\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$ with $\sigma_n = [\alpha_n, \beta_n]$, and with initial and terminal markings $\iota_n, \tau_n \in \mathcal{M}(W_n)$. We will say that the jet J has *sudden growth* if there exist simplices y_n and z_n on σ_n such that

- $\|[\alpha_n, y_n]\|_J$ grows sublinearly,
- $\|[y_n, z_n]\|_J$ grows linearly, and
- $d_{\mathcal{C}(W_n)}(y_n, z_n)$ is bounded for ω -a.e. n .

We say that J has *gradual growth* if it does not have sudden growth.

Given $\bar{z} \in \mathcal{M}(\bar{S})$ we say that \bar{z} *escapes linearly along J* if

$$\|[\alpha_n, \pi_{\sigma_n}(z_n)]\|_J$$

has linear growth, otherwise \bar{z} *escapes sublinearly*.

Although linear and sublinear escape are only defined in the ultraproduct $\mathcal{M}(\bar{S})$, the following lemma says that they are defined in the asymptotic cone $\mathcal{M}_\omega(S)$ as long as the jet in question has gradual growth.

Lemma 7.6. *Let J be a macroscopic jet with the gradual growth property. The linear/sublinear escape properties for sequences descend to the ultralimits when these lie in $\mathcal{M}_\omega(S) \setminus \mathcal{L}_\omega(J)$. In other words, we can decompose $\mathcal{M}_\omega(S) \setminus \mathcal{L}_\omega(J)$ as a disjoint union,*

$$\mathcal{M}_\omega(S) \setminus \mathcal{L}_\omega(J) = \Omega_J \cup \Lambda_J$$

so that $z \in \Lambda_J$ implies that any sequence (z_n) representing z escapes linearly along J , and $z \in \Omega_J$ implies any (z_n) escapes sublinearly.

We will establish the following.

Theorem 7.7. *Let J be a macroscopic jet with the gradual growth property. Then Λ_J and Ω_J are both open. Moreover, Λ_J is Σ -convex, and therefore is acyclic.*

Note that Ω_J is not acyclic at all — indeed it is not even connected. It breaks up into uncountably many connected components as an application of Theorem 7.2.

We will establish both the theorem and the lemma as consequences of the following more quantitative fact:

Lemma 7.8. *There exists $C > 0$ such that the following holds for any macroscopic jet J with gradual growth. Suppose that $\bar{\xi}, \bar{\xi}' \in \mathcal{M}(\bar{S})$ represent $\xi, \xi' \in \mathcal{M}_\omega(S)$ and that*

- (1) $\xi \notin \mathcal{L}_\omega(J)$
- (2) $d_{\mathcal{M}_\omega(S)}(\xi, \xi') < C d(\xi, \mathcal{L}_\omega(J))$

Then either ξ_n and ξ'_n both escape linearly along J , or both escape sublinearly.

Proof of Lemma 7.8: Write $J = (\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$, let C be the constant in Lemma 7.4, and suppose, by way of contradiction, that (1) and (2) hold but one of $\bar{\xi}, \bar{\xi}'$ escapes sublinearly and the other escapes linearly. After renaming the one that escapes linearly $\bar{\zeta}$ and the one that escapes sublinearly $\bar{\eta}$, we find that $\pi_{\sigma_n}(\eta_n)$ must precede $\pi_{\sigma_n}(\zeta_n)$ along σ_n for ω -a.e. n , and that $\|[\alpha_n, \pi_{\sigma_n}(\eta_n)]\|_J$ grows sublinearly while $\|[\pi_{\sigma_n}(\eta_n), \pi_{\sigma_n}(\zeta_n)]\|_J$ grows linearly (this uses the additivity property (2) in Lemma 7.1).

Since $\bar{\eta}$ escapes sublinearly, the restricted jet

$$J' = (\bar{\sigma}', \bar{W}, \bar{a}, \bar{b}),$$

defined by letting $\sigma'_n = [\alpha_n, \pi_{\sigma_n}(\eta_n)]$, is microscopic. Moreover, gradual growth implies, as above, that if we enlarge σ'_n in the forward direction by an amount which is bounded for ω -a.e. n then we still obtain a microscopic jet. Thus, we may produce a new microscopic jet extending σ'_n along σ_n by any bounded amount which is larger than the constant, K , needed to apply Lemma 7.4. In this new jet $J'' = (\bar{\sigma}'', \bar{W}, \bar{a}, \bar{b})$ we find that $\pi_{\sigma''_n}(\eta_n) = \pi_{\sigma_n}(\eta_n)$, while $\pi_{\sigma''_n}(\zeta_n)$ equals (up to bounded error) the forward endpoint of σ''_n , and hence

$$d_{\mathcal{C}(W_n)}(\pi_{\sigma''_n}(\xi'_n), \pi_{\sigma''_n}(\xi_n)) > K.$$

Thus Lemma 7.4 implies that $d(\xi, \xi') \geq C d(\xi, \mathcal{L}_\omega(J))$. (Note that $\iota(J) = \iota(J'')$, so $\mathcal{L}_\omega(J) = \mathcal{L}_\omega(J'')$.) This contradicts our hypothesis that $d(\xi, \xi') < C d(\xi, \mathcal{L}_\omega(J))$. Hence it must hold that $\bar{\xi}$ escapes linearly if and only if $\bar{\xi}'$ does. \square

Lemma 7.6 follows immediately from Lemma 7.8 by considering the case $\xi = \xi'$.

Proof of Theorem 7.7. The openness of Λ_J and Ω_J is an easy consequence of Lemma 7.8. It remains to prove that Λ_J is Σ -convex, for we can then apply Corollary 6.6 to conclude that Λ_J is acyclic.

Let $A \subset \Lambda_J$ be finite and let \bar{A} represent it. Then each $\bar{a} \in \bar{A}$ has projections to σ_n which escape linearly. The projection of $\Sigma_\epsilon(A_n)$ to σ_n is, up to bounded error, the projection in $\mathcal{C}(W_n)$ of $\text{hull}_{W_n}(A_n)$ to σ_n , and hyperbolicity implies that this is contained (up to bounded error) in the hull along σ_n of the projections of A_n . Hence any point in the hull has projections that escape linearly, and so is in Λ_J . This proves Σ -convexity. \square

7.4. Classification of pieces

We now record an application of Theorem 7.7, which classifies the maximal subsets of \mathcal{M}_ω which can not be separated by a point. This result will help to motivate the statement of Theorem 8.1.

First, let us recall the notions of *pieces* and *tree-graded* spaces as defined by Druţu–Sapir [12].

A complete geodesic metric space X is called *tree-graded* if there exists a collection of proper closed convex subsets, \mathcal{P} , called *pieces*, which pairwise intersect in at most one point and such that every non-trivial simple geodesic triangle in X is contained in one piece. It is an easy observation that if X contains a point whose removal disconnects it, then X is tree-graded. Further, in any tree-graded space X there exists a unique *finest* way to write X as a union of pieces none of which can be separated by a point [12].

In terms of $\mathcal{C}(S)$ -distance, we now provide a complete criterion for when two points in $\mathcal{M}_\omega(S)$ can be globally separated by a point. In particular, the following result describes the pieces in the finest decomposition of $\mathcal{M}_\omega(S)$ as a tree-graded space. We note that by results of [4] such pieces can not be realized as asymptotic cones of subgroups of $\mathcal{M}(S)$.

Theorem 7.9. *Suppose that $\xi(S) \geq 2$. For any pair of points $\mu, \nu \in \mathcal{M}_\omega(S)$, the following are equivalent:*

- (1) *No point of $\mathcal{M}_\omega(S)$ separates μ from ν .*
- (2) *In any neighborhoods of μ, ν , respectively, there exist points μ', ν' with representative sequences $(\mu'_n), (\nu'_n)$ such that*

$$\lim_{\omega} d_{\mathcal{C}(S)}(\mu'_n, \nu'_n) < \infty$$

Remark. Theorem 7.9 can be thought of as a first step in working out the results of Section 8, which provide conditions under which a finite subset of $\mathcal{M}_\omega(S)$ is separated by a product region. Nevertheless, Theorem 7.9 will not be used in the rest of the paper.

Proof. We begin by showing (2) implies (1). Suppose first that μ and ν have representative sequences (μ_n) and (ν_n) for which $\lim_\omega d_{\mathcal{C}(S)}(\mu_n, \nu_n) < \infty$. Hence there is a fixed $m \geq 0$ such that $d_{\mathcal{C}(S)}(\mu_n, \nu_n) = m$ for ω -a.e. n , and we can let $v_{n,0}, \dots, v_{n,m}$ denote the simplices of a tight geodesic in $\mathcal{C}(S)$ connecting $v_{n,0} \in \text{base}(\mu_n)$ to $v_{n,m} \in \text{base}(\nu_n)$. For a fixed i let $\bar{v}_i = \langle v_{n,i} \rangle$. The regions $\mathcal{Q}(v_{n,i})$ have the structure described in Lemma 3.1, and in particular the cone $\mathcal{Q}_\omega(\bar{v}_i)$ is nontrivial (not a singleton) and connected.

Since $v_{n,i}$ and $v_{n,i+1}$ are disjoint (here we use $\xi(S) \geq 2$), we have $v_{n,i} \cap v_{n,i+1} = v_{n,i} \cup v_{n,i+1}$, so by Lemma 3.3 the intersection $\mathcal{Q}_\omega(\bar{v}_i) \cap \mathcal{Q}_\omega(\bar{v}_{i+1})$ is equal to $\mathcal{Q}_\omega(\bar{v}_i \cup \bar{v}_{i+1})$. This again is not a singleton, and it follows that the union

$$\mathcal{Q}_\omega(\bar{v}_0) \cup \dots \cup \mathcal{Q}_\omega(\bar{v}_m)$$

cannot be disconnected by a point. Since $\mu \in \mathcal{Q}_\omega(\bar{v}_0)$ and $\nu \in \mathcal{Q}_\omega(\bar{v}_m)$, this gives property (1) in this case where $\lim_\omega d_{\mathcal{C}(S)}(\mu_n, \nu_n) < \infty$.

Now for general μ and ν satisfying (2), the above argument implies that μ and ν can be approximated arbitrarily closely by μ' and ν' which cannot be separated by a point. Since maximal subsets without cutpoints are closed [12], this completes the proof that (2) implies (1).

We now establish that (1) implies (2), by proving the contrapositive. Namely suppose that (2) fails to hold for μ and ν , so that there exists $r > 0$ such that whenever $d(\mu, \mu') \leq r$ and $d(\nu, \nu') \leq r$, we have $d_{\mathcal{C}(S)}(\mu'_n, \nu'_n) \rightarrow \infty$ for any representative sequences. We can assume $r < d_{\mathcal{M}_\omega}(\mu, \nu)/2$.

Note that Proposition 5.2 implies that $\Sigma_\epsilon(\mu_n, \nu_n)$ is coarsely connected. In particular, by projecting a continuous path in $\mathcal{M}(S)$ from μ_n to ν_n into $\Sigma_\epsilon(\mu_n, \nu_n)$, we can obtain points $\mu'_n, \nu'_n \in \Sigma_\epsilon(\mu_n, \nu_n)$ such that $d_{\mathcal{M}(S)}(\mu_n, \mu'_n)$ and $d_{\mathcal{M}(S)}(\nu_n, \nu'_n)$ are in the interval $[\frac{1}{2}rs_n, rs_n]$ for all sufficiently large n . Fix such a pair of sequences (μ'_n) and (ν'_n) . It follows that $d_{\mathcal{C}(S)}(\mu'_n, \nu'_n) \rightarrow \infty$, and by Lemma 7.5 there exists a microscopic jet $J = (\bar{\sigma}, \bar{S}, \bar{\mu}', \bar{\nu}')$ such that $\bar{\mu}' \not\sim_{\bar{\sigma}} \bar{\nu}'$.

Since μ'_n and ν'_n are in $\Sigma_\epsilon(\mu_n, \nu_n)$, the segments σ_n must be within a bounded distance of any $\mathcal{C}(S)$ -geodesic between $\text{base}(\mu_n)$ and $\text{base}(\nu_n)$. It follows that $\pi_{\sigma_n}(\mu_n)$ and $\pi_{\sigma_n}(\nu_n)$ are within bounded distance of $\pi_{\sigma_n}(\mu'_n)$ and $\pi_{\sigma_n}(\nu'_n)$, respectively. Hence we also have $\bar{\mu} \not\sim_{\bar{\sigma}} \bar{\nu}$.

Now $\mathcal{L}_\omega(J) = \{\iota(J)\}$ since J is built on the main surface S . We claim that $\mu \neq \iota(J)$ and $\nu \neq \iota(J)$. This follows from two facts about Σ -hulls:

First, for any $a, b \in \mathcal{M}(S)$ we claim that, if $a', b' \in \Sigma_\epsilon(a, b)$, then

$$d(a, \Sigma_\epsilon(a', b')) \gtrsim d(a, \{a', b'\})$$

(with uniform constants). In the projection to each $\mathcal{C}(W)$, the Σ -hulls map to coarse intervals, and the corresponding inequality is simply the fact that if two intervals are nested then the endpoints of the inner one separate its interior

from the endpoints of the outer one. The statement then follows from the quasidistance formula, Theorem 2.8.

Second, if v is a vertex on a tight $\mathcal{C}(S)$ -geodesic g from a vertex of $\text{base}(a)$ to a vertex of $\text{base}(b)$, then $v]a$ is in $\Sigma_\epsilon(a, b)$, for a uniform ϵ . This will follow by showing, for all $W \subseteq S$, that $\pi_{\mathcal{C}(W)}(v]a)$ is uniformly close to $\text{hull}_W(a, b)$. If $W \not\cap v$, the projections of $v]a$ and a are by definition close. If $W \cap v$ and $W \neq S$, then tightness of g implies that either the subsegment from v to a or the one from v to b consists of simplices overlapping W , and so Theorem 2.6 implies that one of $d_W(v, a)$ or $d_W(v, b)$ is uniformly bounded. If $W = S$ then v is already in $\text{hull}_S(a, b)$.

Applying the first fact and the choice of r, μ'_n and ν'_n , we see that

$$d_{\mathcal{M}(S)}(\mu_n, \Sigma_\epsilon(\mu'_n, \nu'_n)) \gtrsim \frac{1}{2}rs_n.$$

From the second fact and the definition of $\iota_n(J)$, we see that $\iota_n(J) \in \Sigma_\epsilon(\mu'_n, \nu'_n)$, for a uniform ϵ . Hence $\boldsymbol{\mu} \neq \boldsymbol{\iota}(J)$. The same applies to $\boldsymbol{\nu}$.

We have shown that $\bar{\boldsymbol{\mu}} \not\sim_{\sigma} \bar{\boldsymbol{\nu}}$, and that $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are different from $\boldsymbol{\iota}(J)$. Hence by Theorem 7.2, $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are separated by the point $\boldsymbol{\iota}(J)$.

□

7.5. Manifolds and jets

As an application of the linear/sublinear decomposition associated to a macroscopic jet with the gradual growth property, together with the homology result Theorem 6.8 and Corollary 6.9, we can obtain:

Theorem 7.10. *Let E be a $\xi(S)$ -dimensional connected manifold in $\mathcal{M}_\omega(S)$ and let J be a macroscopic jet with the gradual growth property. Suppose the supporting subsurface \bar{W} of J has $\xi(\bar{W}) > 1$. Then if $E \cap \mathcal{L}_\omega(J) \neq \emptyset$, we conclude*

$$E \cap \Lambda_J = \emptyset.$$

Proof. Let $\mathbf{q} \in \mathcal{L}_\omega(J) \cap E$. Suppose on the contrary that $E \cap \Lambda_J \neq \emptyset$. Now $\mathcal{L}_\omega(J)$ has codimension at least 2 since $\xi(\bar{W}) > 1$. Applying Theorem 3.6 the set $\mathcal{L}_\omega(J)$ cannot separate E which has dimension $\xi(S)$. However, $\mathcal{L}_\omega(J)$ does separate Λ_J from Ω_J in \mathcal{M}_ω , by Theorem 7.7. We conclude that $E \setminus \mathcal{L}_\omega(J)$ is contained in Λ_J .

Now let B be a ball in E containing \mathbf{q} in its interior. Since $\mathcal{L}_\omega(J) \cap \partial B$ is a compact set of codimension ≥ 1 in ∂B , for any $\epsilon > 0$ there exists a triangulation of ∂B with vertices outside $\mathcal{L}_\omega(J)$ and mesh size ϵ . Using Lemmas 6.1, 6.3, and 6.5, as in the proof of Theorem 7.7, B can be deformed to a Σ -compatible chain C , such that every point moves at most $c\epsilon$ (with c a uniform constant) and the 0-skeleton does not move at all. Since the 0-skeleton is contained in Λ_J , by Theorem 7.7 all of C is contained in Λ_J as well. Let U be the r -chain

giving the homotopy of ∂B to C , i.e. $\partial U = \partial B - C$ and U is supported in a $c\epsilon$ neighborhood of ∂B .

Since C sits within Λ_J , Theorem 7.7 also implies that it bounds an r -chain B' in Λ_J .

Corollary 6.9 now implies, since B is embedded and $\partial B = \partial(B' + U)$, that $B \subset B' + U$. Assuming we have chosen ϵ so that $c\epsilon < d(\mathbf{q}, \partial B)/2$, we find that \mathbf{q} cannot be in U . Hence $\mathbf{q} \in B' \subset \Lambda_J$. This is a contradiction. \square

8. Local finiteness for manifolds

Our main goal in this section is Theorem 8.7, which says that any top-dimensional submanifold of $\mathcal{M}_\omega(S)$ is locally contained in a union of finitely many cubes.

This will be a consequence of Theorem 8.5, in which we will consider the Σ -hull of a finite number of points in a connected top-dimensional manifold in $\mathcal{M}_\omega(S)$, and show that it is always contained in a finite complex made of cubes of the appropriate dimension. In order to do this we will prove Theorem 8.1, which will show that points in the manifold can be represented by sequences of markings whose projections to all but the simplest subsurfaces remain bounded. This in turn will be possible because of the separation theorems established in Section 7.

8.1. Trimming theorem

In Theorem 7.9 we showed that for any $\mu, \nu \in \mathcal{M}_\omega(S)$, if no point in $\mathcal{M}_\omega(S)$ separates μ, ν then — after perturbation — the curve complex distance $d_{\mathcal{C}(S)}(\mu_n, \nu_n)$ is ω -a.e. bounded.

Theorem 7.9 can be regarded as a baby version of Theorem 8.1. Given a top-dimensional manifold E in $\mathcal{M}_\omega(S)$, the Alexander duality argument given in Lemma 3.6 shows that E *cannot* be separated by any product region of codimension ≥ 2 in $\mathcal{M}_\omega(S)$; such product regions are associated to sequences of connected, essential subsurfaces \overline{W} such that $\xi(\overline{W}) > 1$. From this, Theorem 8.1 will conclude that the curve complex diameters of finite subsets of E — after trimming — are ω -a.e. bounded, and this will be true in the curve complexes of *all* \overline{W} such that $\xi(\overline{W}) > 1$.

A manifold in $\mathcal{M}_\omega(S)$ is *top dimensional* if it is of dimension $\xi(S)$, equal to the locally compact dimension of $\mathcal{M}_\omega(S)$.

Theorem 8.1. *Let \overline{A} be a finite set of elements in $\mathcal{M}(\overline{S})$. Suppose A_ω is contained in a connected top-dimensional manifold $E \subset \mathcal{M}_\omega(S)$. There exist ϵ and k_0 , a new set, \overline{A}' , and an onto map $\tau: \overline{A} \rightarrow \overline{A}'$ with the following properties.*

- (1) $\tau(\overline{a})_\omega = a_\omega$ for each $\overline{a} \in \overline{A}$
- (2) $A'_n \subset \Sigma_\epsilon(A_n)$ for ω -a.e. n .

(3) For any \overline{W} with $\xi(\overline{W}) > 1$,

$$\text{diam}_{\mathcal{C}(W_n)}(A'_n) < k_0$$

for ω -a.e. n .

Notation: because \overline{A} is finite we can think of it, up to ultraproduct equivalence, as a sequence of finite sets A_n , and we can think of the trimming map τ as a sequence of maps from A_n to A'_n . With slight abuse of notation we will also use τ to denote these maps, thus writing for example $\tau(a_n) \in \tau(A_n) = A'_n$.

Proof. We will argue by induction on the cardinality of \overline{A} . The case of cardinality 1 is trivial, so let us consider the case that \overline{A} has two points.

The fundamental step of the proof is the following lemma, which “trims” \overline{A} to reduce its projections to a given subsurface sequence \overline{W} :

Lemma 8.2. *Suppose \overline{A} has two elements and A_ω is contained in a connected top-dimensional manifold $E \subset \mathcal{M}_\omega(S)$. Let \overline{W} be represented by a sequence (W_n) of connected, essential subsurfaces with $\xi(\overline{W}) > 1$. There exists a map $\tau: \overline{A} \rightarrow \Sigma_\epsilon(\overline{A})$ such that*

$$(\tau\overline{x})_\omega = x_\omega$$

for each $\overline{x} \in \overline{A}$, and

$$\text{diam}_{\mathcal{C}(W_n)}(\tau(A_n))$$

is bounded, for ω -a.e. n . The constant ϵ depends only on the topological type of S .

Proof. We may assume that $\text{diam}_{\mathcal{C}(W_n)}(A_n) \rightarrow_\omega \infty$, for if not then $\text{diam}_{\mathcal{C}(W_n)}(A_n)$ is ω -a.e. bounded, and we may simply take τ to be inclusion of \overline{A} into $\Sigma_\epsilon(\overline{A})$. All we need is that $\text{diam}_{\mathcal{C}(W_n)}(A_n) \rightarrow_\omega \infty$ is greater than the constant m_1 of Lemma 5.5 for ω -a.e. n , for we may then apply that lemma, concluding that for each connected, essential subsurface U of S ,

$$\text{if } U \cap W_n \text{ for } \omega\text{-a.e. } n \text{ then } d_{\mathcal{C}(U)}(\partial W_n, \text{hull}_U(A_n)) \leq m_0 \text{ for } \omega\text{-a.e. } n \quad (8.1)$$

This fact is used repeatedly below, although only once with explicit details, in the proof of (3) below.

We find the trimming map τ in stages. First let $\tau_1(\overline{x})$, for $\overline{x} \in \overline{A}$, be defined as follows:

$$\tau_1(\overline{x}) = \begin{cases} \overline{x} & x_\omega \notin \mathcal{Q}_\omega(\partial\overline{W}) \\ \pi_{\mathcal{Q}(\partial\overline{W})}(\overline{x}) & x_\omega \in \mathcal{Q}_\omega(\partial\overline{W}). \end{cases} \quad (8.2)$$

The notation $\pi_{\mathcal{Q}(\mu)}: \mathcal{M}(S) \rightarrow \mathcal{Q}(\mu)$ denotes the map $\nu \mapsto \mu \rfloor \nu$ from Section 3.1. In particular the sequence $\pi_{\mathcal{Q}(\partial W_n)}$ gives rise to a map $\pi_{\mathcal{Q}(\partial\overline{W})}: \mathcal{M}(\overline{S}) \rightarrow \mathcal{Q}(\partial\overline{W})$.

We claim that $\tau_1(\overline{x})$ has the properties

$$(1) \tau_1(\overline{x})_\omega = x_\omega,$$

- (2) Either $\tau_1(\bar{x})_\omega \notin \mathcal{Q}_\omega(\partial\bar{W})$, or $\tau_1(x_n) \in \mathcal{Q}(\partial W_n)$ for ω -a.e. n .
 (3) $\tau_1(\bar{x}) \in \Sigma_\epsilon(\bar{A})$ for suitable ϵ .

Property (1) is immediate from the definition and the fact that (by Proposition 3.1 (3)) $\pi_{\mathcal{Q}(\partial W)}(x)$ realizes, within bounded factor, the distance from x to $\mathcal{Q}(\partial W)$. Property (2) similarly follows from the definition.

To see Property (3), note it is obvious when $\tau_1(\bar{x}) = \bar{x}$. Hence, assume that $\tau_1(\bar{x}) = \pi_{\mathcal{Q}(\partial\bar{W})}(\bar{x}) \in \mathcal{Q}(\partial\bar{W})$. By definition of Σ_ϵ , it suffices to bound the distance from $\pi_{\mathcal{C}(U)}(\tau_1(x_n))$ to $\text{hull}_U(A_n)$, for every $U \subseteq S$ and ω -a.e. n . We treat separately the cases $U \pitchfork \partial W_n$ and $U \not\pitchfork \partial W_n$.

If $U \not\pitchfork \partial W_n$ for ω -a.e. n , then $d_{\mathcal{C}(U)}(\tau_1(x_n), x_n)$ is uniformly bounded: this can be seen easily from the definition of the projection $\pi_{\mathcal{Q}(\partial W_n)}$ and the coarse composition properties of subsurface projection maps, Lemma 2.12. Since $x_n \in A_n \subset \Sigma_\epsilon(A_n)$, this gives the desired bound for (3).

If $U \pitchfork \partial W_n$ for ω -a.e. n , then $d_{\mathcal{C}(U)}(\tau_1(x_n), \partial W_n)$ is bounded by Lemma 2.11, and by applying (8.1) we obtain (3).

Now for notational simplicity let us assume that we have replaced \bar{A} by $\tau_1(\bar{A})$. For each $\bar{x} \in \bar{A}$ properties (1) and (3) become trivial, and property (2) holds with the consequence that either $x_\omega \notin \mathcal{Q}_\omega(\partial\bar{W})$, or $\bar{x} \in \mathcal{Q}(\partial\bar{W})$. Recalling that \bar{A} has two elements, write $\bar{A} = \{\bar{a}, \bar{b}\}$, and the discussion separates into two cases: neither of a_ω, b_ω is in $\mathcal{Q}_\omega(\partial\bar{W})$; or one of \bar{a}, \bar{b} is in $\mathcal{Q}(\partial\bar{W})$.

Case 1: Neither a_ω nor b_ω lies in $\mathcal{Q}_\omega(\partial\bar{W})$.

We claim that already $d_{\mathcal{C}(W_n)}(a_n, b_n)$ is bounded for ω -a.e. n , and hence there is nothing left to do in this case. Suppose otherwise, that $d_{\mathcal{C}(W_n)}(a_n, b_n) \rightarrow_\omega \infty$. Then Lemma 7.5 yields a microscopic jet $J = (\bar{\sigma}, \bar{W}, \bar{a}, \bar{b})$, built from $\mathcal{C}(W_n)$ -geodesics σ_n for which $(a_n) \not\sim_{\bar{\sigma}} (b_n)$. Moreover, by our assumption that $a_\omega, b_\omega \notin \mathcal{Q}_\omega(\partial\bar{W})$, we know that neither of them is contained in $\mathcal{L}_\omega(J) \subset \mathcal{Q}_\omega(\partial\bar{W})$. It follows from Lemma 7.2 that $\mathcal{L}_\omega(J)$ separates a_ω from b_ω .

However, $\mathcal{L}_\omega(J)$ is homeomorphic to $\mathcal{M}_\omega(\bar{W}^c)$, which by Proposition 2.2 has codimension at least 2 since $\xi(\bar{W}) > 1$, so by Lemma 3.6 it cannot separate E . This contradiction implies that in fact $d_{\mathcal{C}(W_n)}(a_n, b_n)$ is bounded ω -a.s.

Case 2: At least one of \bar{a}, \bar{b} , say \bar{a} , lies in $\mathcal{Q}(\partial\bar{W})$.

Now we consider the projections of a_ω and b_ω to the factor $\mathcal{M}_\omega(\bar{W})$ of $\mathcal{Q}_\omega(\partial\bar{W}) \approx \mathcal{M}_\omega(\bar{W}) \times \mathcal{M}_\omega(\bar{W}^c)$. Using the coarse product structure $\mathcal{Q}(\partial W_n) \approx$

$\mathcal{M}(\overline{W}_n) \times \mathcal{M}(\overline{W}_n^c)$, let

$$\begin{aligned} a_n &\approx \left(\pi_{\mathcal{M}(\overline{W}_n)}(a_n), \pi_{\mathcal{M}(\overline{W}_n^c)}(a_n) \right) \\ &= (\alpha_n, \beta_n) \\ \pi_{\mathcal{Q}(\partial W_n)}(b_n) &\approx \left(\pi_{\mathcal{M}(\overline{W}_n)}(b_n), \pi_{\mathcal{M}(\overline{W}_n^c)}(b_n) \right) \\ &= (\gamma_n, \delta_n) \end{aligned}$$

Case 2a: $\alpha_\omega = \gamma_\omega$.

In this case, we can simply adjust \bar{a} so that the projections, before rescaling, are a bounded distance apart: replace a_n by

$$\tau_2(a_n) = (\gamma_n, \beta_n).$$

We need to check, as before, that $\tau_2(a_n) \in \Sigma_\epsilon(A_n)$ for a fixed ϵ . This is again done by considering projections to all connected, essential subsurfaces $U \subset S$, treating separately the cases $U \Subset W_n$, $U \Subset W_n^c$, or $U \pitchfork \partial W_n$, for ω -a.e. n ; in the latter case (8.1) is again applied, and remaining details are left to the reader. Similarly $\tau_2(\bar{a})_\omega = a_\omega$, and of course $d_{\mathcal{C}(W_n)}(\tau_2(a_n), b_n)$ is now bounded ω -a.s.

Note that this argument works whether or not $\bar{b} \in \mathcal{Q}(\partial \overline{W})$. If it is, then the roles of \bar{a} and \bar{b} can be reversed.

Case 2b: $\alpha_\omega \neq \gamma_\omega$.

In this case, consider the jet $J = (\bar{\sigma}, \overline{W}, \bar{a}, \bar{b})$, where $\sigma_n = [x_n, y_n]$ with $x_n \in \pi_{\mathcal{C}(W_n)}(a_n)$ and $y_n \in \pi_{\mathcal{C}(W_n)}(b_n)$. Notice that the initial marking of J is $\bar{\iota}(J) = \bar{\alpha}$, because $\iota_n(J) = \iota(\sigma_n, a_n, b_n) = x_n \rfloor \pi_{\mathcal{M}(W_n)}(a_n) = x_n \rfloor \alpha_n = \alpha_n$ where the last equation follows since x_n is a vertex of α_n . Similarly the terminal marking of J is $\bar{\tau}(J) = \bar{\gamma}$. Since $\alpha_\omega \neq \gamma_\omega$ it follows that the jet J is macroscopic.

In the arguments to follow we abbreviate the notation for the jet norm $\|\sigma_n\|_J = \|\sigma_n\|_{(W_n, a_n, b_n)}$ to $\|\cdot\|$, and similarly for jets built on subsegments of σ_n , because in all cases the jet norm is equal to $\|\cdot\|_{(W_n, a_n, b_n)}$ which is independent of how the subsegments were chosen.

Since a_ω and b_ω are contained in a connected top-dimensional manifold E , we can apply Theorem 7.10 to conclude that J cannot have the gradual growth property, for $\mathcal{L}_\omega(J) = \mathcal{Q}_\omega(\bar{\alpha} \cup \partial \overline{W})$ contains a_ω , and if J had gradual growth then b_ω would be in Λ_J , but then Theorem 7.10 would forbid b_ω from being in E .

Since J has sudden growth, the following must occur: in σ_n there must be points p_n and q_n such that $\|[x_n, p_n]\|$ grows sublinearly and $\|[p_n, q_n]\|$ grows linearly, while $d_{\mathcal{C}(W_n)}(p_n, q_n)$ stays bounded.

Let $\tau_3(a_n)$ be the marking obtained by projecting a_n to $\mathcal{Q}(p_n \cup \partial W_n)$. Since a_n is already in $\mathcal{Q}(\partial W_n)$ it follows that $\tau_3(a_n)$ is coarsely equal to the projection

of a_n to $\mathcal{Q}(p_n)$. We have

$$\begin{aligned} d_{\mathcal{M}(S)}(a_n, \tau_3(a_n)) &\approx d_{\mathcal{M}(S)}(a_n, \mathcal{Q}(p_n)) \\ &\approx d_{\mathcal{M}(W_n)}(\pi_{\mathcal{M}(W_n)}(a_n), \mathcal{Q}_{W_n}(p_n)) \end{aligned}$$

the last equation following from Proposition 3.1, where $\mathcal{Q}_{W_n}(p_n)$ denotes the product region in $\mathcal{M}(W_n)$ associated to the partial marking p_n . This product region contains the projection to $\mathcal{M}(W_n)$ of $p_n \rfloor \gamma_n$ and so

$$\begin{aligned} d_{\mathcal{M}(W_n)}(\pi_{\mathcal{M}(W_n)}(a_n), \mathcal{Q}_{W_n}(p_n)) &\prec d_{\mathcal{M}(W_n)}(\pi_{\mathcal{M}(W_n)}(a_n), p_n \rfloor \gamma_n) \\ &\approx d_{\mathcal{M}(W_n)}(x_n \rfloor \alpha_n, p_n \rfloor \gamma_n) \\ &\approx \|[x_n, p_n]\| \end{aligned}$$

The quantity $d_{\mathcal{M}(S)}(a_n, \tau_3(a_n))$ therefore grows sublinearly, and so $\tau_3(\bar{a})_\omega = a_\omega$. As before we can show that $\tau_3(\bar{a}) \in \Sigma_\epsilon(\bar{A})$.

In the case that $\bar{b} \in \mathcal{Q}(\partial\bar{W})$, by reversing the direction of J so that its initial marking is $\bar{\gamma}$, the same argument as above shows that the reversed jet has sudden growth, and we obtain a path sequence $[v_n, u_n]$ in σ_n such that $\|[v_n, u_n]\|$ grows linearly and $\|[u_n, y_n]\|$ grows sublinearly while $d_{\mathcal{C}(W_n)}(v_n, u_n)$ stays bounded ω -a.e. As with \bar{a} , in this case we define $\tau_3(b_n) = \pi_{\mathcal{Q}(u_n \cup \partial W_n)}(b_n)$, and so $\tau_3(\bar{b})_\omega = b_\omega$. In the case that $b_\omega \notin \mathcal{Q}_\omega(\partial\bar{W})$, then we simply let $\tau_3(\bar{b}) = \bar{b}$, and let $u_n = v_n = y_n$. In σ_n we have a sequence $x_n \cdots p_n \cdots q_n \cdots v_n \cdots u_n \cdots y_n$, where we can ensure that q_n precedes v_n by shortening the segments $[p_n, q_n]$ and $[v_n, u_n]$ if necessary, maintaining the property that $\|[p_n, q_n]\|$ and $\|[v_n, u_n]\|$ each grow linearly.

We claim now that $d_{\mathcal{C}(W_n)}(\tau_3(a_n), \tau_3(b_n)) \approx d_{\mathcal{C}(W_n)}(p_n, u_n)$ is bounded ω -a.s. For if it were not then neither would $d_{\mathcal{C}(W_n)}(q_n, v_n)$ be bounded ω -a.s., and so as in case (1) we could extract a microscopic jet $J' = (\bar{\sigma}', \bar{W}, \bar{a}, \bar{b})$ with σ'_n a subsegment of $[q_n, v_n]$. Neither of the points a_ω and b_ω can be in $\mathcal{L}_\omega(J')$. For a_ω , this follows from the fact that $\|[p_n, q_n]\|$ grows linearly and hence insulates a_n from σ'_n — that is, by Lemma 7.1 (1) and the quasidistance formula we obtain, term-by-term, a linearly growing lower bound for $d_{\mathcal{M}(W_n)}(\alpha_n, \iota_n(J'))$. For b_ω this is the same argument if $\bar{b} \in \mathcal{Q}(\partial\bar{W})$, and if not it is even easier for b_ω is not even in $\mathcal{Q}_\omega(\partial\bar{W})$.

Hence, Lemma 7.2 would imply that $\mathcal{L}_\omega(J')$ separates E , and applying Lemma 3.6 this would again contradict the assumption that $\xi(\bar{W}) > 1$.

We conclude that, in case (2), we can find $\tau_3(\bar{A})$ such that $\text{diam}_{\mathcal{C}(W_n)}(\tau_3(\bar{A}))$ is bounded ω -a.s. This concludes the proof of Lemma 8.2, where τ is the composition of the appropriate τ_i . \square

Hierarchies of geodesics. Before we can continue the proof of Theorem 8.1 we must recall a few of the details of the construction of hierarchies of tight geodesics from [39]. A hierarchy $H = H(a, b)$ is associated to any $a, b \in \mathcal{M}(S)$,

and is a certain collection of tight geodesics in curve complexes of connected, essential subsurfaces of S . The subsurface whose complex contains a geodesic h is called its *domain* $D(h)$. The properties relevant to us are the following:

Theorem 8.3. *Let $a, b \in \mathcal{M}(S)$ and $H(a, b)$ a hierarchy between them.*

- (1) *There is a unique main geodesic g_H with $D(g_H) = S$, whose endpoints lie on $\text{base}(a)$ and $\text{base}(b)$.*
- (2) *For any geodesic $h \in H$ other than g_H , there exists another geodesic $k \in H$ such that, for some simplex v in k , $D(h)$ is either a component of $D(k) \setminus v$, or an annulus whose core is a component of v . We say that $D(h)$ is a component domain of k .*
- (3) *A subsurface in S can occur as the domain of at most one geodesic in H .*
- (4) *For each $h \in H$, the endpoints of h are within uniformly bounded distance of $\pi_{D(h)}(a)$ and $\pi_{D(h)}(b)$.*
- (5) *For each connected, essential subsurface $W \subset S$, if $d_{\mathcal{C}(W)}(a, b) > m_0$, then there exists $h \in H(a, b)$ with $D(h) = W$.*

Define the complexity of a geodesic $g \in H(a, b)$ to be $\xi(D(g))$. The following counting argument allows us, under the appropriate circumstances, to bound the cardinality of the set of geodesics with a given lower bound on complexity.

Lemma 8.4. *For all $t \geq 1$, all $a, b \in \mathcal{M}(S)$, and all $k \geq 1$, if $d_{\mathcal{C}(W)}(a, b) \leq k$ for all subsurfaces W with $\xi(W) \geq t$, then the hierarchy $H(a, b)$ contains at most $O(k^{\xi(S)-t})$ geodesics of complexity $\xi = t - 1$.*

Proof. The proof is by induction, using the properties listed in Theorem 8.3. Every subsurface of complexity $\xi = s$ in $H(a, b)$ appears as a component domain in some geodesic of complexity $> s$. Hence the number of $\xi = s$ geodesics is bounded by the number of $\xi > s$ geodesics times the length bound on those geodesics. \square

We now turn to the proof of Theorem 8.1 in the case that \overline{A} has two elements. We first apply Lemma 8.2 with $\overline{W} = \overline{S}$. Thus we obtain $\tau(\overline{A})$, such that $\text{diam}_{\mathcal{C}(S)}(\tau(A_n))$ is ω -a.s. bounded. Again for notational convenience we replace \overline{A} by $\tau(\overline{A})$ and continue.

Writing $\overline{A} = \{\overline{a}, \overline{b}\}$ as before, we consider hierarchies $H_n = H(a_n, b_n)$. By property (5), for any \overline{W} with $\text{diam}_{\mathcal{C}(W_n)}(A_n) \rightarrow_{\omega} \infty$, W_n must be a domain in H_n for ω -a.e. n . The main geodesics g_{H_n} have bounded length for ω -a.e. n , by property (1) and the bound on $\text{diam}_{\mathcal{C}(S)}(A_n)$. Applying Lemma 8.4 with $t = \xi(S)$ we obtain an ω -a.s. bound on the number of geodesics of complexity $\xi(S) - 1$ in H_n , because the domain of such a geodesic must be a component domain of the main geodesic g_{H_n} . We have therefore bounded how many \overline{W} exist with $\xi(\overline{W}) = \xi(S) - 1$ and with W_n a domain in H_n for ω -a.e. n — we

use here the general fact that the ultraproduct of a sequence of sets X_n of finite cardinality $\leq k$ has cardinality $\leq k$. For each such \overline{W} successively, use Lemma 8.2 again to find $\tau(\overline{A})$ such that $\text{diam}_{\mathcal{C}(W_n)}(\tau(\overline{A})_n)$ is bounded, and again replace \overline{A} by $\tau(\overline{A})$ and continue.

Every time we apply Lemma 8.2, we maintain the boundedness that we had for $\text{diam}_{\mathcal{C}(\overline{U})}$ for any previous \overline{U} . This is because $\tau(\overline{A})$ always lies in $\Sigma_\epsilon(\overline{A})$, so in the projections to $\mathcal{C}(U_n)$, it follows that $\pi_{\mathcal{C}(U_n)}(\tau(A_n))$ lies uniformly near the hull of $\pi_{\mathcal{C}(U_n)}(A_n)$ which is bounded. Hence after finitely many steps we have diameter bounds for all \overline{W} with $\xi(\overline{W}) = \xi(S) - 1$.

This procedure repeats $\xi(S)$ times. At the k^{th} step we have bounds on the lengths of all geodesics of complexity $\geq \xi(S) - k + 1$ that occur in the hierarchy, and by applying Lemma 8.4 we bound the *number* of geodesics of complexity $\geq \xi(S) - k$. A finite number of applications of Lemma 8.2 renders bounded the projections to those surfaces without spoiling the previous ones.

The procedure ends when all projections to surfaces of $\xi > 1$ are bounded. The final set, which we might denote $\tau^N(\overline{A})$ (for some N which grows with $\xi(S)$ and the bounds at each level), lies in $\Sigma_{\epsilon'}(\overline{A})$ (where ϵ' depends on ϵ and N), and each $\tau^N(\overline{x})$ defines the same point in the cone as \overline{x} .

This concludes the proof of Theorem 8.1 when \overline{A} has two elements. We are now ready for the inductive step, where we write \overline{A} as $\{\overline{a}\} \cup \overline{B}$, and we assume that there is already a bound on $\text{diam}_{\mathcal{C}(W_n)}(B_n)$ for ω -a.e. n , whenever $\xi(\overline{W}) > 1$.

We wish to prove an analogue of Lemma 8.2, and there is a similar breakup into cases. Let \overline{W} be such that $\xi(\overline{W}) > 1$ and $\text{diam}_{\mathcal{C}(W_n)}(A_n) \rightarrow_\omega \infty$. First we note as in the proof of Lemma 8.2 that we may assume (after a first trimming operation τ_1) that each element $\overline{x} \in \overline{A}$ either satisfies

$$x_\omega \notin \mathcal{Q}_\omega(\partial\overline{W})$$

or satisfies

$$x_n \in \mathcal{Q}(\partial W_n)$$

for ω -a.e. n (or as we wrote above, $\overline{x} \in \mathcal{Q}(\partial\overline{W})$).

Case 1a': Suppose that there is at least one element $\overline{b} \in \overline{B}$ with $b_\omega \notin \mathcal{Q}_\omega(\partial\overline{W})$, and that $a_\omega \notin \mathcal{Q}_\omega(\partial\overline{W})$. Then the same argument as Case 1 of Lemma 8.2 shows that $\text{diam}_{\mathcal{C}(W_n)}(a_n, b_n)$ is bounded. Since $\text{diam}_{\mathcal{C}(W_n)}(B_n)$ was already bounded, this gives us the desired bound for A_n .

Case 1b': Suppose that there is at least one element $\overline{b} \in \overline{B}$ with $b_\omega \notin \mathcal{Q}_\omega(\partial\overline{W})$, but that $\overline{a} \in \mathcal{Q}(\partial\overline{W})$. Depending on whether or not $\pi_{\mathcal{M}_\omega(\overline{W})}(a_\omega) = \pi_{\mathcal{M}_\omega(\overline{W})}(b_\omega)$, we can apply the argument of Cases 2a and 2b of Lemma 8.2, concluding that a_n can be replaced by $\tau_2(a_n)$, for which $d_{\mathcal{C}(W_n)}(\tau_2(a_n), b_n)$ is bounded. Again since $\text{diam}_{\mathcal{C}(W_n)}(B_n)$ is assumed bounded we are done.

Case 2a': Suppose that $\bar{b} \in \mathcal{Q}(\partial\bar{W})$ for each $\bar{b} \in \bar{B}$, and suppose that also $\pi_{\mathcal{M}_\omega(\bar{W})}(B_\omega)$ is a single point.

In this case, choose one element $\bar{b}_0 \in \bar{B}$. Now apply the argument of Case 2a and 2b in Lemma 8.2 to \bar{b}_0 and \bar{a} . Note that here \bar{b}_0 plays the role that \bar{a} played in 2a and 2b, whereas \bar{a} itself may or may not be in $\mathcal{Q}(\partial\bar{W})$. This step produces $\tau_3(\bar{b}_0)$ which possibly modifies the $\mathcal{M}(\bar{W})$ component of \bar{b}_0 (and similarly for \bar{a}), so that afterwards their $\mathcal{C}(W_n)$ distance is ω -a.s. bounded. Define τ_3 on the remaining elements of \bar{B} by making their $\mathcal{M}(\bar{W})$ components equal to that of $\tau_3(\bar{b}_0)$. This is a sublinear change which as before produces points in $\Sigma_\epsilon(B_n)$. We now have the desired bound on $\text{diam}_{\mathcal{C}(W_n)}(\tau_3(A_n))$.

Case 2b': Again suppose that $\bar{b} \in \mathcal{Q}(\partial\bar{W})$ for each $\bar{b} \in \bar{B}$, but now suppose that $\pi_{\mathcal{M}_\omega(\bar{W})}(B_\omega)$ contains at least 2 distinct points. Let $\bar{b}_1, \bar{b}_2 \in \bar{B}$ have distinct projections to $\mathcal{M}_\omega(\bar{W})$.

If $\bar{a} \in \mathcal{Q}(\partial\bar{W})$, and $\pi_{\mathcal{M}_\omega(\bar{W})}(a_\omega) = \pi_{\mathcal{M}_\omega(\bar{W})}(b_{i\omega})$ for $i = 1$ or $i = 2$, then as in Case 2a of Lemma 8.2, we can replace the $\mathcal{M}(\bar{W})$ component of \bar{a} to agree with that of \bar{b}_1 or \bar{b}_2 , respectively, and are done.

If $\bar{a} \in \mathcal{Q}(\partial\bar{W})$ but $\pi_{\mathcal{M}_\omega(\bar{W})}(a_\omega)$ is different from both $\pi_{\mathcal{M}_\omega(\bar{W})}(b_{1\omega})$ and $\pi_{\mathcal{M}_\omega(\bar{W})}(b_{2\omega})$, or if $a_\omega \notin \mathcal{Q}_\omega(\partial\bar{W})$, then we work with \bar{b}_1 and \bar{a} as follows.

If $a_\omega \notin \mathcal{Q}_\omega(\partial\bar{W})$ then let $\tau_4(\bar{a}) = \bar{a}$. If $\bar{a} \in \mathcal{Q}(\partial\bar{W})$, we argue as in Case 2b of Lemma 8.2, first to show that a jet from \bar{a} to \bar{b}_1 cannot have gradual growth, and then to modify \bar{a} : along the geodesic from $x_n \in \pi_{\mathcal{C}(W_n)}(a_n)$ to $y_n \in \pi_{\mathcal{C}(W_n)}(b_{1n})$, we find p_n and q_n such that $\| [x_n, p_n] \|$ grows sublinearly, $\| [p_n, q_n] \|$ grows linearly, and $d_{\mathcal{C}(W_n)}(p_n, q_n)$ is ω -a.s. bounded. We then let $\tau_4(\bar{a}) = (p_n \cup \partial W_n] a_n$.

Unlike Case 2b of Lemma 8.2, we do not attempt to modify \bar{b}_1 . Now if $d_{\mathcal{C}(W_n)}(\tau_4(a_n), b_{1n})$ is still unbounded, we find a microscopic jet J' built from subgeodesics σ_n of $[q_n, y_n]$, so that $a_\omega \notin \mathcal{L}_\omega(J')$ by the same argument at Case 2b. The points a_n and b_{1n} project to opposite sides of σ_n so $\bar{a} \not\sim_{\bar{\sigma}} \bar{b}_1$. Hence if $b_{1\omega} \notin \mathcal{L}_\omega(J')$, then we are done, because $\mathcal{L}_\omega(J')$ then separates a_ω from $b_{1\omega}$ and hence separates E , which is a contradiction. But if $b_{1\omega} \in \mathcal{L}_\omega(J')$ then we must have $b_{2\omega} \notin \mathcal{L}_\omega(J')$, because $(b_1)_\omega$ and $(b_2)_\omega$ have distinct images in $\mathcal{M}_\omega(\bar{W})$. Since $d_{\mathcal{C}(W_n)}(b_{1n}, b_{2n})$ is ω -a.s. bounded, we also have $\bar{a} \not\sim_{\bar{\sigma}} \bar{b}_2$, and hence $\mathcal{L}_\omega(J')$ separates $b_{2\omega}$ from a_ω , and we still have a contradiction.

We conclude that $d_{\mathcal{C}(W_n)}(\tau_4(a_n), b_{1n})$ is ω -a.s. bounded, which is what we wanted to show.

This gives the analogue of Lemma 8.2 for $\bar{A} = \bar{B} \cup \{\bar{a}\}$. Now we finish the proof as we did before: we repeatedly apply this result, bounding first the lengths of the main geodesics in hierarchies between elements of \bar{A} , and then

inducting downward to bound the lengths of geodesics of lower complexities, until only domains of complexity 1 are left with unbounded diameters. \square

8.2. Finitely many cubes

As a consequence of Theorem 8.1, we will show that the Σ -hull of a finite number of points in a connected top-dimensional manifold is composed of finitely many cubes (in the sense of Section 3.1). From this we'll get the statement on finitely many orthants in a neighborhood of a point.

Theorem 8.5. *If A is a finite subset of a connected top-dimensional manifold E in $\mathcal{M}_\omega(S)$, then $\Sigma(A)$ is contained in a finite union of cubes.*

The first step towards establishing this theorem is the following, here we are using \mathcal{N}_r to denote a radius r neighborhood.

Lemma 8.6. *For each integer N , there exists a constant k_2 , such that for any finite set $A \subset \mathcal{M}(S)$ with $\#A = N$ and $\text{diam}_{\mathcal{C}(S)}(A_n) < k_0$ the following holds for each $a \in A$: the set $\Sigma_\epsilon(A)$ is contained in the union of sets*

$$G'(A, \mathcal{U}, a) = \mathcal{N}_{k_2}(G(A, \mathcal{U}, a))$$

where \mathcal{U} varies over all sets of the form $\mathcal{U}(\mu, a)$ for $\mu \in \Sigma_\epsilon(A)$.

Proof. Fix ϵ large enough for Proposition 5.2 (on retractions of Σ -hulls) to apply. For later use fix another constant $k_0 > \max\{3(m_0 + 4), 2m_0 + \epsilon, B - \epsilon\}$, where m_0 is the constant of Lemma 4.2 and B is the constant of Theorem 2.6.

Fix $a \in A$, and now consider any $\mu \in \Sigma_\epsilon(A)$. Following Section 4, we use μ and a to define a partial order among certain subsurfaces of S . As noted in Section 4, the projections $(\pi_W(\mu))$ satisfy the consistency conditions with any $c_1 \geq m_0$. Now define \ll_k and \prec_k as in section 4, that is, $V \ll_k W$ if and only if $V \pitchfork \partial W$ and

$$d_{\mathcal{C}(V)}(\mu, \partial W) > k(c_1 + 4), \tag{8.3}$$

whereas $V \prec_k W$ if and only if $V \ll_k W$ and $V \pitchfork W$. We choose c_1 so that $k_0 + \epsilon = 3(c_1 + 4)$. In particular $\mathcal{F}_3(\mu, a) = \{W : W \ll_3 a\}$ is the set

$$\{W \subsetneq S : d_{\mathcal{C}(W)}(\mu, a) > k_0 + \epsilon\}.$$

Lemma 4.5 now tells us that \prec_2 is a partial order on $\mathcal{F}_3(\mu, a)$. Moreover, by Lemma 4.4, if $V, W \in \mathcal{F}_3(\mu, a)$ and if $V \pitchfork W$ then V, W are \prec_2 -ordered.

The set $\mathcal{F}_3(\mu, a)$ is finite — using the quasidistance formula for example, or Lemma 4.6 — so we can let $\mathcal{V} = \mathcal{V}(\mu, a)$ be the set of \prec_2 -minimal elements. Any two elements of \mathcal{V} are disjoint or nested in S , so let $\mathcal{U} = \mathcal{U}(\mu, a)$ be the subset of \mathcal{V} consisting of elements maximal with respect to containment in S . Hence \mathcal{U} enumerates the components of an essential subsurface of S , which we abuse notation by also calling \mathcal{U} . Recall that $\mathcal{Q}(\partial\mathcal{U})$ has a natural product structure

$\mathcal{M}(\mathcal{U}^c) \times \mathcal{M}(\mathcal{U})$. We claim that μ is within uniformly bounded distance of a subset of $\mathcal{Q}(\partial\mathcal{U})$ of the form

$$G(A, \mathcal{U}, a) = \{\pi_{\mathcal{M}(\mathcal{U}^c)}(a)\} \times \prod_{U \in |\mathcal{U}|} \Sigma_{\epsilon'}(\pi_{\mathcal{M}(U)}(A)) \quad (8.4)$$

where $|\mathcal{U}|$ is the set of components of \mathcal{U} , and $\Sigma_{\epsilon'}$ is defined within $\mathcal{M}(U)$ just as it was in $\mathcal{M}(S)$. The constant ϵ' depends only on ϵ and $\xi(S)$.

To prove this, we first bound $d_{\mathcal{M}(S)}(\mu, \mathcal{Q}(\partial\mathcal{U}))$. By Proposition 3.1, we just need to establish a bound on $d_{\mathcal{C}(W)}(\mu, \partial\mathcal{U})$ for all W that overlap $\partial\mathcal{U}$. By hypothesis, $\text{diam}_{\mathcal{C}(S)}(A_n) < k_0$ which implies by Theorem 2.6 a uniform bound on $d_{\mathcal{C}(S)}(\mu, \partial\mathcal{U})$; hence we now assume $W \subsetneq S$. Suppose that $d_{\mathcal{C}(W)}(\mu, \partial\mathcal{U}) > 4(c_1 + 4)$. In particular $W \ll_4 U$ for some $U \in \mathcal{U}$ such that $W \pitchfork \partial\mathcal{U}$. Since $U \ll_3 a$, by Lemma 4.4 (2) we have $W \ll_3 a$, so that $W \in \mathcal{F}_3(\mu, a)$.

If $W \pitchfork U$ then $W \prec_3 U$ and in particular $W \prec_2 U$, contradicting the minimality of U . Hence W must contain U . However, by choice of \mathcal{U} this means W cannot be \prec_2 -minimal, so there exists $Z \in \mathcal{F}_3(\mu, a)$ such that $Z \prec_2 W$. By Lemma 4.4 (1), $Z \prec_2 W \ll_3 U$ implies that $Z \prec_1 U$. But in particular this means $Z \pitchfork U$ so they are \prec_2 -ordered. $U \prec_2 Z$ would contradict $Z \prec_1 U$, so we must have $Z \prec_2 U$, but this contradicts again the minimality of U .

We conclude that for all W such that $W \pitchfork \partial\mathcal{U}$, $d_{\mathcal{C}(W)}(\mu, \partial\mathcal{U}) \leq 4(c_1 + 4)$, and this gives a bound of the form

$$d_{\mathcal{M}(S)}(\mu, \mathcal{Q}(\partial\mathcal{U})) \leq k_1$$

for some k_1 depending on c_1 (and hence on m_0 and k_0).

Next we claim that $\pi_{\mathcal{M}(\mathcal{U}^c)}(\mu)$ is uniformly close to $\pi_{\mathcal{M}(\mathcal{U}^c)}(a)$. For this, by the quasidistance formula we need to bound $d_{\mathcal{C}(W)}(\mu, a)$ for all $W \subset \mathcal{U}^c$. Suppose that $d_{\mathcal{C}(W)}(\mu, a) > 3(c_1 + 4)$, and so $W \in \mathcal{F}_3(\mu, a)$. Since W is disjoint from all components of \mathcal{U} and hence of \mathcal{V} , it is not \prec_2 -ordered with or isotopic to any of them. W cannot be \prec_2 -minimal as then it would have to be one of \mathcal{V} . Hence there is some $W' \prec_2 W$ which is \prec_2 -minimal — but then W' is in \mathcal{V} , and again we have a contradiction.

Finally we consider $\pi_{\mathcal{M}(\mathcal{U})}(\mu)$. Since $\mu \in \Sigma_{\epsilon}(A)$, for each connected subsurface $W \subseteq \mathcal{U}$ we have $\pi_W(\mu) \in \mathcal{N}_{\epsilon}(\text{hull}_W(A))$, where $\mathcal{N}_{\epsilon}(\cdot)$ denotes the radius ϵ neighborhood (this involves an abuse of distance notation, as explained under the heading ‘‘Subsurface projections’’ in Section 2.1.1). But $\text{hull}_W(A)$ is within uniformly bounded distance of $\text{hull}_W(\pi_{\mathcal{M}(\mathcal{U})}(A))$ by the coarse composition property of projections (Lemma 2.12). Hence $\pi_{\mathcal{M}(\mathcal{U})}(\mu) \in \Sigma_{\epsilon'}(\pi_{\mathcal{M}(\mathcal{U})}(A))$ for some ϵ' depending on ϵ and $\xi(S)$.

This establishes that μ is within uniform distance of the set $G(A, \mathcal{U}, a)$ described in (8.4), proving the Lemma. \square

Proof of Theorem 8.5. From Theorem 8.1 we may assume that \mathbf{A} is represented by (A_n) such that, for ω -a.e. n , $\text{diam}_{\mathcal{C}(W)}(A_n)$ is bounded by some fixed k_0 whenever $\xi(W) > 1$. Let us consider an arbitrary $A \subset \mathcal{M}(S)$, of fixed cardinality $\#A = \#\mathbf{A}$, satisfying this condition.

Let k_2 denote the constant given by Lemma 8.6 for $\#A$. Now we would like to bound the number of \mathcal{U} that can occur in the output of Lemma 8.6.

If W occurs as a component of $\mathcal{U}(\mu, a)$ for some μ , then $d_{\mathcal{C}(W)}(\mu, a) > k_0 + \epsilon$. Since $\pi_W(\mu) \in \mathcal{N}_\epsilon(\text{hull}_W(A))$,

$$\text{diam}_{\mathcal{C}(W)}(A) > k_0.$$

By our assumptions about A , this means $\xi(W) \leq 1$. Now define

$$\begin{aligned} \mathcal{S}_1 &= \{U \subsetneq S : \xi(U) = 1 \text{ and } \text{diam}_U(A) > k_0\}, \\ \mathcal{S}_0 &= \{U \subsetneq S : \xi(U) = 0 \text{ and } \text{diam}_U(A) > k_0\}. \end{aligned}$$

By Theorem 8.3 (5), every element in $\mathcal{S}_0 \cup \mathcal{S}_1$ must be the domain of some geodesic in $H(a, b)$ for some $a, b \in A$. Hence the counting argument, Lemma 8.4, directly gives a bound on the cardinality of \mathcal{S}_1 .

There is no uniform bound for the cardinality of \mathcal{S}_0 , but we can control the number of annuli $U \in \mathcal{S}_0$ which are components of $\mathcal{U}(\mu, a)$ for some $\mu \in \Sigma_\epsilon(A)$. The main reason for this is the following statement:

- (*) For each connected, essential $W \subset S$, each essential annulus $U \Subset W$, each $a \in A$, and each $\mu \in \Sigma_\epsilon(A)$, if $\text{diam}_{\mathcal{C}(W)}(A) > k_0$, and if U is a component of $\mathcal{U}(\mu, a)$, then $d_{\mathcal{C}(W)}(\partial U, A) \leq k_3$ for k_3 depending on k_0 and $\#A$.

To prove (*), define $\mathcal{F}_3(\mu, a)$ using μ as before, and consider two cases depending on whether $W \in \mathcal{F}_3(\mu, a)$.

If $W \notin \mathcal{F}_3(\mu, a)$ then $d_{\mathcal{C}(W)}(\mu, a) \leq k_0 + \epsilon$. Since $d_{\mathcal{C}(U)}(\mu, a) > k_0 + \epsilon > B$, any $\mathcal{C}(W)$ -geodesic from $\pi_W(\mu)$ to $\pi_W(a)$ must pass within distance 1 in $\mathcal{C}(W)$ of ∂U , by Theorem 2.6. It follows that $d_{\mathcal{C}(W)}(\partial U, a) \leq d_{\mathcal{C}(W)}(\mu, a) + 1 \leq k_0 + \epsilon + 1$.

If $W \in \mathcal{F}_3(\mu, a)$ then, since $U \Subset W$, the surface W cannot be a \prec_2 -minimal element of $\mathcal{F}_3(\mu, a)$, because then W would have been included in $\mathcal{U}(\mu, a)$ instead of U . Hence there is some element $Y \in \mathcal{V}(\mu, a)$ such that $Y \prec_2 W$.

We claim that $d_{\mathcal{C}(W)}(\partial Y, b)$ is bounded for some $b \in A$. The argument is similar to the partial-order arguments in Section 4. Since $\mu \in \Sigma_\epsilon(A)$, we have $\pi_Y(\mu) \in \mathcal{N}_\epsilon(\text{hull}_Y(A))$. Also, since $d_{\mathcal{C}(Y)}(\mu, a) > k_0 + \epsilon$ there must be $b \in A$ such that $d_{\mathcal{C}(Y)}(a, b) \geq k_0$. Now $Y \prec_2 W$ implies that $d_{\mathcal{C}(Y)}(\mu, \partial W) \geq 2(c_1 + 4) > m_0$ so that $d_{\mathcal{C}(W)}(\mu, \partial Y) < m_0$ by Lemma 4.2. Further, since $d_{\mathcal{C}(W)}(a, \mu) \geq k_0$, we have $d_{\mathcal{C}(W)}(\partial Y, a) > k_0 - m_0 - 2 > m_0$. Again by Lemma 4.2, we have $d_{\mathcal{C}(Y)}(\partial W, a) < m_0$. Now since $d_{\mathcal{C}(Y)}(a, b) \geq k_0$ we have $d_{\mathcal{C}(Y)}(\partial W, b) > k_0 - m_0 - 2 > m_0$, so applying Lemma 4.2 one more time we get $d_{\mathcal{C}(W)}(\partial Y, b) < m_0$.

Since Y and U are disjoint, we conclude $d_{\mathcal{C}(W)}(\partial U, b) < m_0 + 1$. This finishes the proof of (*).

Now we can control the number of elements in \mathcal{S}_0 which occur as components of $\mathcal{U}(\mu, a)$ for $\mu \in \Sigma_\epsilon(A)$. Given such a U , there exists $b \in A$ such that $d_{\mathcal{C}(U)}(a, b) \geq (k_0 - \epsilon)/2 > m_0$, and so by Theorem 8.3 there exists a geodesic $h \in H(a, b)$ with domain $D(h) = W$ such that $\xi(W) \geq 1$ and U is a component domain of h . Noting that either h has length $\leq k_0$ or $\pi_{\mathcal{C}(W)}(A)$ has diameter $> k_0$, by applying (*) it follows that there exists $c \in A$ such that

$$d_{\mathcal{C}(W)}(\partial U, c) \leq k_4 = \max\{k_0, k_3\}$$

This restricts ∂U , for each c , to a segment of length at most $2k_4$ in h . Now since the number of hierarchies involved is controlled in terms of $\#A$, and the number of $\xi \geq 1$ surfaces appearing is controlled in terms of $\#A$ and k_0 by Lemma 8.4, this gives us a bound on the total number of components of the $\mathcal{U}(\mu, a)$ as μ varies over $\Sigma_\epsilon(A)$.

We now apply this result to the sets A_n in the sequence \bar{A} . Each one is covered by the uniformly bounded number of sets $G'(A, \mathcal{U}, a)$. Taking rescalings, we obtain in the asymptotic cone the statement that $\Sigma(\mathbf{A})$ is contained in a finite union of asymptotic cones of sequences $G(A_n, \mathcal{U}_n, a_n)$, which by (8.4) must be sets of the form

$$\{x_\omega\} \times \prod_{\bar{U} \in |\bar{\mathcal{U}}|} T_{\bar{U}}$$

in $\mathcal{Q}_\omega(\partial \bar{\mathcal{U}})$, where $x_\omega \in \mathcal{M}_\omega(\bar{\mathcal{V}}^c)$, each component \bar{U} of $\bar{\mathcal{U}}$ has $\xi \leq 1$, and each $T_{\bar{U}}$ is the convex hull of a finite set in the \mathbb{R} -tree $\mathcal{M}_\omega(\bar{U})$. Hence each $T_{\bar{U}}$ is a finite tree, so after breaking each tree into a finite union of segments, we obtain the desired finite union of cubes. \square

8.3. Local finiteness

The main application of Theorem 8.5 is the following:

Theorem 8.7. *If $E \subset \mathcal{M}_\omega(S)$ is a connected top-dimensional manifold, then any compact subset of E is contained in a finite union of cubes.*

Proof. It suffices to show that a ball $B \subset E$ is contained in finitely many cubes.

Let $B \subset \text{int}(B')$ where B' is a larger ball. Triangulate $\partial B'$ with simplices of diameter smaller than r , where r will be chosen shortly. Let $f_0: B' \rightarrow E$ be the identity, and let $f_1: \partial B' \rightarrow \mathcal{M}_\omega(S)$ be a Σ -compatible map with respect to the triangulation, which agrees with f_0 on the 0-skeleton; the existence of f_1 follows from Lemma 6.3. By Lemma 6.1 we have $d(f_0, f_1) < Cr$, and by Lemma 6.5 there exists a homotopy $h: \partial B' \times [0, 1] \rightarrow \mathcal{M}_\omega(S)$ with track diameters at most $C'r$, for uniform constants C, C' .

Choose r small enough that $C'r < \frac{1}{2}d(B, \partial B')$. Then we find that the image of h is disjoint from B .

Extend the triangulation of $\partial B'$ to one of B' without adding any vertices. Then using Lemma 6.3 again, f_1 can be extended to a Σ -compatible map $F: B' \rightarrow \mathcal{M}_\omega(S)$ with respect to this triangulation. Let K be the chain which is the sum of F and h – then we note that $\partial K = \partial B'$. By Corollary 6.9, we conclude that $B' \subset K$. Since B is disjoint from h , we have

$$B \subset F.$$

Now F is contained in the Σ -hulls of a finite collection of finite subsets of E . By Theorem 8.5, it must therefore be contained in a finite union of cubes. \square

9. Germs and orthants

In this section, we study the local structure of the set of top-dimensional manifolds passing through a point $\mathbf{x} \in \mathcal{M}_\omega(S)$, by considering the *germs* of such manifolds, and using the Local Finiteness Theorem 8.7 to relate this to the *complex of orthants* through \mathbf{x} .

We then apply this to a study of *Dehn twist flats* in the asymptotic cone $\mathcal{M}_\omega(S)$: given a sequence $\bar{\nu}$ of pants decompositions of S , we obtain a sequence $Q(\bar{\nu})$ of Dehn twist flats in $\mathcal{M}(S)$, and passing to the rescaled ultralimit we obtain $Q_\omega(\bar{\nu})$ which, if nonempty, is by definition a Dehn twist flat in $\mathcal{M}_\omega(S)$. By Proposition 3.1, $Q_\omega(\bar{\nu})$ is a bilipschitz embedded copy of $\mathbf{R}^{\xi(S)}$, although we shall not make use of this property.

The main result is Corollary 9.8, which states that germs of Dehn twist flats passing through \mathbf{x} admit a purely topological characterization. This will be applied in Section 10 in the proof that Dehn twist flats in $\mathcal{M}_\omega(S)$ are preserved by homeomorphisms, and Dehn twist flats in $\mathcal{M}(S)$ are coarsely preserved by quasi-isometries.

9.1. Poset of Germs

Consider the set of closed subsets of \mathcal{M}_ω containing \mathbf{x} , modulo the equivalence $C \sim C'$ if there exists an open neighborhood U of \mathbf{x} such that $C \cap U = C' \cap U$. The equivalence classes are called *germs through \mathbf{x}* , and we let $\gamma(C)$ denote the germ of C through \mathbf{x} . Note that finite intersection and union yield well-defined operations on the set of germs, and the subset relation is well-defined as well. Let \mathcal{G} denote the poset (partially ordered set) of germs at \mathbf{x} ; this is a *lattice*, meaning that least upper bounds and greatest lower bounds exist for all pairs $C, C' \in \mathcal{G}$, namely $C \cup C'$ and $C \cap C'$.

A property P of germs at points of $\mathcal{M}_\omega(S)$ is *topologically characterizable* if any local homeomorphism of $\mathcal{M}_\omega(S)$ taking \mathbf{x} to \mathbf{y} takes germs at \mathbf{x} satisfying P to germs at \mathbf{y} satisfying P . For example, germs of manifolds are topologically characterizable, as is the dimension function on such germs.

9.2. Structure of orthants

In this section we fix $\mathbf{x} \in \mathcal{M}_\omega(S)$ and study the set of germs of cubes in $\mathcal{M}_\omega(S)$ for which \mathbf{x} is a *corner*. These germs will be called *orthants at \mathbf{x}* . The goal in this section is Lemma 9.1 which shows that the poset of nontrivial orthants at \mathbf{x} has the structure of a simplicial flag complex. For this purpose we need to study the relations of equality, subset and intersection of orthants. The main complication that arises is that an orthant has many different representations by sequences of cubes in $\mathcal{M}(S)$, so these relations are not trivial to detect.

Recall from Section 3.3 that if $V \Subset S$ satisfies $\xi(V) \leq 1$ then $\mathcal{TM}(V)$ denotes a particular tree quasi-isometric to the marking complex $\mathcal{M}(V)$, which in the annulus case is a line.

A *cube with distinguished corner* is a cube $C = C(\mu, W, r)$ for which each geodesic $r_i \subset \mathcal{TM}(W_i)$ has a distinguished endpoint $r_i(0)$. The corner of C is, by definition, the marking $\kappa(C) = \{\mu\} \times \prod r_i(0)$, where the right side is interpreted as usual within $\mathcal{Q}(\partial W) \cong \mathcal{M}(W^c) \times \prod \mathcal{M}(W_i)$. Given a sequence $C(\bar{\mu}, \bar{W}, \bar{r})$ of cubes distinguished corners, we obtain in $\mathcal{M}_\omega(S)$ a cube $C^\omega = C^\omega(\bar{\mu}, \bar{W}, \bar{r})$ with corner κ^ω .

We recall a few features of the notation for cubes. First, up to ultraproduct equivalence, the sequence $\bar{W} = (W^n)$ of essential subsurfaces can be identified with a finite set of sequences $\bar{W}_1 = (W_1^n), \dots, \bar{W}_k = (W_k^n)$ of connected essential subsurfaces such that (for ω -a.e. n) the components of W^n are W_1^n, \dots, W_k^n . Second, the sequence \bar{r} can be identified with a finite set of sequences $\bar{r}_1 = (r_1^n), \dots, \bar{r}_k = (r_k^n)$ such that r_i^n is a geodesic segment or ray in the tree $\mathcal{TM}(W_k^n)$ with initial point $r_i^n(0)$. Third, recall from Section 3.4 that the dimension of the asymptotic cube $C^\omega(\bar{\mu}, \bar{W}, \bar{r})$ is equal to the number of components \bar{r}_i of \bar{r} such that the length of the limiting segment $r_i \in \mathcal{M}_\omega(\bar{W})$ is positive — equivalently the lengths $l(r_i^n)$ grow linearly. Finally, given two essential subsurface sequences \bar{W}, \bar{V} , up to reindexing we may assume that for each i, j , if $W_i^n \stackrel{i}{\cong} V_j^n$ for ω -a.e. n then $i = j$; reindexing in this manner is implicit, for example, in the statement of Lemma 9.1.

Define an *orthant at \mathbf{x}* to be the germ $O = \gamma(C^\omega)$ of an asymptotic cube C^ω with distinguished corner $\kappa^\omega = \mathbf{x}$. A *k -orthant* is the germ of an asymptotic cube of dimension k . If an orthant O can be expressed as $O = \gamma(C^\omega(\bar{\mu}, \bar{W}, \bar{r}))$ where all components of \bar{W} are annuli, then we say that O is a *Dehn twist orthant*.

Note that for any orthant $O = \gamma(C^\omega(\bar{\mu}, \bar{W}, \bar{r}))$ at \mathbf{x} , the asymptotic partial marking $\boldsymbol{\mu} \in \mathcal{M}_\omega(\bar{W}^c)$ is determined by the subsurface sequence \bar{W} and the corner \mathbf{x} , being the projection of \mathbf{x} to $\mathcal{M}_\omega(\bar{W}^c)$.

For example, the germ of every Dehn twist flat is a union of $2^{\xi(S)}$ Dehn twist $\xi(S)$ -orthants. To be precise, consider a Dehn twist flat $Q_\omega(\bar{v})$ through

$\mathbf{x} = x_\omega \in \mathcal{M}_\omega(S)$, where $\text{base}(x_i) = \nu_i$. Let \overline{W} be the sequence of annulus neighborhoods of the pants decomposition sequence \overline{v} . The surface \overline{W}^c is empty, so we take $\overline{\mu}$ to be the empty partial marking. The projection of \overline{x} to the line $\mathcal{TM}(\overline{W}_i)$ has two directions denoted $\overline{r}_i^1, \overline{r}_i^2$ (using the usual identification between an ultraproduct of finite sets and a finite subset of an ultraproduct). We may therefore express the germ at \mathbf{x} of $Q_\omega(\overline{v})$ as the union of the Dehn twist orthants $\gamma(C^\omega(\overline{\mu}, \overline{W}, \overline{r}_i^{\hat{j}}))$, where the multi-index \hat{j} varies over $\{1, 2\}^{\xi(S)}$.

First we describe a way to normalize representations of orthants by requiring that Dehn twists be transparently represented. Given a cube sequence $C(\overline{\mu}, \overline{W}, \overline{r})$ and a component \overline{W}_i we say that r_i^ω is a *twist direction* in the \mathbf{R} -tree $\mathcal{M}_\omega(\overline{W}_i)$ if r_i^ω has positive length and there exists a sequence of linearly growing twist segments \overline{s}_i in $\mathcal{TM}(\overline{W}_i)$ such that r_i^ω and s_i^ω have the same germ in $\mathcal{M}_\omega(\overline{W}_i)$; equivalently, one can truncate sublinearly growing initial segments of r_i^ω so that what is left has linearly growing initial subsegments that are twist segments. A cube sequence $C(\overline{\mu}, \overline{W}, \overline{r})$ is said to be *twist normalized* if for each i the segment r_i^ω has positive length, and r_i^ω is a twist direction if and only if \overline{W}_i is an annulus. The dimension of a twist normalized asymptotic cube $C^\omega(\overline{\mu}, \overline{W}, \overline{r})$ equals the number of components of \overline{W} , which is therefore well-defined independent of the choice of a twist normalization. Lemma 9.1 (1) shows more, namely that \overline{W} itself is well defined up to ultraproduct equivalence.

We claim that every orthant O can be represented by a twist normalized cube sequence. To see why, consider an arbitrary representation $O = \gamma(C^\omega(\overline{\mu}, \overline{W}, \overline{r}))$. First, for each i such that r_i^ω has zero length, we can extend $\overline{\mu}$ by the partial marking $\overline{r}_i(0)$, and then we can drop the components \overline{W}_i and \overline{r}_i from the notation, obtaining a new cube sequence representing O . Next, for each i such that r_i^ω is not a twist direction, it is already true that \overline{W}_i is not an annulus sequence. Finally, for each i such that r_i^ω is a twist direction, if \overline{W}_i is not already an annulus sequence then we can replace \overline{W}_i by an annulus sequence $\overline{V}_i \subseteq \overline{W}_i$, and we can replace \overline{r}_i by a linearly growing segment sequence \overline{s}_i in $\mathcal{TM}(\overline{V}_i)$, such that the image of s_i^ω under the embedding $\mathcal{M}_\omega(\overline{V}_i) \hookrightarrow \mathcal{M}_\omega(\overline{W}_i)$ has the same germ as r_i^ω . The result of these replacements is a twist normalized cube sequence still representing the orthant O .

Define \mathcal{O} , the *orthant complex* at \mathbf{x} , to be the poset of all nontrivial orthants at \mathbf{x} — all orthants except for the singleton $\{\mathbf{x}\}$ — with respect to the subset relation. We shall use junctures to show that \mathcal{O} has the structure of a simplicial flag complex. Recall that a simplicial complex is a *flag complex* if, whenever a subgraph of the 1-skeleton is isomorphic to the 1-skeleton of a simplex, it is equal to the 1-skeleton of a simplex in the given complex.

Lemma 9.1. *For any orthants O_1, O_2 , and for any twist normalized representations $O_1 = \gamma(C^\omega(\overline{\mu}, \overline{V}, \overline{r}))$ and $O_2 = \gamma(C^\omega(\overline{v}, \overline{W}, \overline{s}))$, the following hold:*

- (1) $O_1 = O_2$ if and only if for ω -a.e. n the following hold: $W^n \stackrel{i}{\cong} V^n$, and for each i the segments r_i^ω, s_i^ω have the same germ in the \mathbf{R} -tree $\mathcal{M}_\omega(\overline{W}_i) = \mathcal{M}_\omega(\overline{V}_i)$.
- (2) $O_1 \subset O_2$ if and only if for ω -a.e. n the following hold: each component W_i^n of W^n is isotopic to a component V_i^n of V^n and the segments r_i^ω, s_i^ω have the same germ in the \mathbf{R} -tree $\mathcal{M}_\omega(\overline{W}_i) = \mathcal{M}_\omega(\overline{V}_i)$. In this case we say that O is a face of O' .
- (3) $O_1 \cap O_2$ is the maximal common face of O_1 and O_2 .

Moreover, the poset \mathcal{O} is isomorphic to the poset of simplices of a flag complex \mathcal{K} , having one simplex of dimension $k - 1$ for each k -orthant.

Remarks. Since equality, subset, and intersection are well-defined set theoretic operations, it follows from items (1–3) that the face relation and the “maximal common face” are well-defined independent of the choice of twist normalized representations, which is not at all clear a priori.

Proof. While items (1) and (3) formally follow from item (2), the proofs of (1–3) will all follow by studying $O_1 \cap O_2$ using junctures.

Let $C_1^\omega = C^\omega(\overline{\mu}, \overline{V}, \overline{r})$ and $C_2^\omega = C^\omega(\overline{\nu}, \overline{W}, \overline{s})$. To understand $C_1^\omega \cap C_2^\omega$, recall from Section 3.4 that this intersection is either empty or equal to the common ultralimit of the junctures of the approximating cubes. Since both cubes contain \mathbf{x} , the empty case cannot occur, and we are left to study the junctures.

Lemmas 3.7 and 3.9 show that the junctures of the approximating cubes $C_1^n = C(\mu^n, V^n, r^n)$ and $C_2^n = C(\nu^n, W^n, s^n)$ are themselves subcubes $C_{12}^n \subset C_1^n$ and $C_{21}^n \subset C_2^n$, which have the form

$$C_{12}^n = C(\mu'^n, V^n, r'^n),$$

where r'^n denotes a collection consisting of a subinterval (or point) on each segment of the collection r^n , and similarly

$$C_{21}^n = C(\nu^n, W^n, s'^n),$$

where s'^n is a collection of subintervals or points of s^n . These lemmas also produce an indexing of V^n and W^n and a $k \geq 0$ so that for ω -a.e. n we have: $V_i^n \cap W_j^n \neq \emptyset$ if and only if $1 \leq i = j \leq k$ in which case we set $U_i^n = V_i^n \cap W_i^n$; this occurs only if $r_i'^n, s_i'^n$ have the same length (positive or zero), and all other lengths in r'^n and s'^n are zero. Following our usual ultraproduct convention we can say that $\overline{V}_i \cap \overline{W}_j$ is nonempty if and only if $i = j \in \{1, \dots, k\}$, in which case it is isotopic to \overline{U}_i .

We can and do parametrize each $r_i'^n$ in such a way that $r_i'^n(0)$ is the point nearest $r_i^n(0)$, and so the corner $\kappa(C_{12}^n) = \{\mu'^n\} \times \prod r_i'^n(0)$ is the nearest corner to $\kappa(C_1^n)$. Because the limiting cube C_{12}^ω contains \mathbf{x} , it must be that the

ultralimit $\kappa^\omega(\overline{C}_{12})$ of $\kappa(C_{12}^n)$ equals \mathbf{x} . It follows that for each $i = 1, \dots, k$ we have $r_i^{\prime\omega}(0) = r_i^\omega(0)$, or equivalently the subsegment of r_i^n from $r_i^n(0)$ to $r_i^{\prime n}(0)$ grows sublinearly. Similar comments applied to s_i^n yield a corner $\kappa(C_{21}^n) = \{\nu^n\} \times \prod s_i^n(0)$ such that $\kappa^\omega(\overline{C}_{21}) = \mathbf{x}$. A note of caution is that this does *not* mean that the initial segments of r_i^n and s_i^n overlap — again, there can be large but sublinearly growing initial segments between $r_i^n(0)$ and $r_i^{\prime n}(0)$, and similarly for s_i^n and $s_i^{\prime n}$. Note also from Lemma 3.9 that for each $i = 1, \dots, k$ we have $\text{Length}(r_i^{\prime n}) \approx \text{Length}(s_i^{\prime n})$, and so the sequences $\text{Length}(\overline{r}_i^{\prime n}), \text{Length}(\overline{s}_i^{\prime n})$ both grow linearly or both grow sublinearly.

Hence we conclude that $O_1 \cap O_2$ can be identified with the face of O_1 associated to those components \overline{W}_i of \overline{W} where $\text{Length}(\overline{r}_i^{\prime n})$ grows linearly, by replacing sublinearly growing segments $\overline{r}_i^{\prime n}$ with the basepoints of \overline{r}_i , and replacing linearly growing segments $\overline{r}_i^{\prime n}$ with initial segments of \overline{r}_i that contain them. The resulting sequence of faces has an ultralimit that coincides with $C_1^\omega \cap C_2^\omega$ in a neighborhood of \mathbf{x} , and hence its germ is equal to $O_1 \cap O_2$.

Furthermore, we claim that for each i , $\text{Length}(\overline{r}_i^{\prime n})$ grows linearly if and only if $i \in \{1, \dots, k\}$, $U_i^n \stackrel{i}{\cong} W_i^n \stackrel{i}{\cong} V_i^n$ for ω -a.e. n , and r_i^ω, s_i^ω have the same positive length germ in $\mathcal{M}_\omega(\overline{W}_i) = \mathcal{M}_\omega(\overline{V}_i)$.

Once this claim is proved, it follows that the intersection $O_1 \cap O_2$ can be described as the face of O_1 associated to those components \overline{W}_i such that \overline{W}_i is isotopic to a component \overline{V}_i of \overline{V} and r_i^ω, s_i^ω have the same positive length germ in $\mathcal{M}_\omega(\overline{W}_i) = \mathcal{M}_\omega(\overline{V}_i)$, and $O_1 \cap O_2$ is similarly described as a face of O_2 . Items (1), (2) and (3) are all immediate consequences of this description.

For the “if” direction of the claim, suppose that $\text{Length}(\overline{r}_i^{\prime n})$ grows sublinearly, that $i \in \{1, \dots, k\}$, that $U_i^n \stackrel{i}{\cong} W_i^n \stackrel{i}{\cong} V_i^n$ for ω -a.e. n — so $\text{Length}(\overline{s}_i^{\prime n})$ also grows sublinearly — and that $\text{Length}(r_i^\omega), \text{Length}(s_i^\omega)$ are both nonzero, so $\text{Length}(\overline{r}_i), \text{Length}(\overline{s}_i)$ both grow linearly. By Lemma 3.9 (2), for ω -a.e. n we have $r_i^{\prime n} = s_i^{\prime n} = r_i^n \cap s_i^n$ in the tree $\mathcal{TM}(V_i^n) = \mathcal{TM}(W_i^n)$. By truncating sublinearly growing initial subsegments of r_i^n and of s_i^n , namely the smallest initial segments containing $r_i^{\prime n}$ and $s_i^{\prime n}$, respectively, we obtain linearly growing segments $r_i^{\prime\prime n}, s_i^{\prime\prime n}$ such that $r_i^\omega = r_i^{\prime\prime\omega}$ and $s_i^\omega = s_i^{\prime\prime\omega}$, and such that $r_i^{\prime\prime n}$ and $s_i^{\prime\prime n}$ have disjoint interiors, so $r_i^{\prime\prime\omega}, s_i^{\prime\prime\omega}$ have distinct germs in $\mathcal{M}_\omega(\overline{W}_i) = \mathcal{M}_\omega(\overline{V}_i)$.

For the “only if” direction, suppose that $\text{Length}(\overline{r}_i^{\prime n})$ does grow linearly, implying that $i \in \{1, \dots, k\}$ and that $\text{Length}(\overline{s}_i^{\prime n})$ also grows linearly. If U_i^n is an annulus then, by Lemma 3.9 (1), each of $r_i^{\prime n}$ and $s_i^{\prime n}$ is a twist segment supported by U_i^n , and by Dehn twist normalization it follows that $U_i^n \stackrel{i}{\cong} W_i^n \stackrel{i}{\cong} V_i^n$. If U_i^n is not an annulus then we also have $U_i^n \stackrel{i}{\cong} W_i^n \stackrel{i}{\cong} V_i^n$, because all these surfaces have $\xi = 1$ and U_i^n is essentially contained in each. Applying Lemma 3.9 (2), it follows that $r_i^{\prime\omega}, s_i^{\prime\omega}$ have the same germ in $\mathcal{M}_\omega(\overline{W}_i) = \mathcal{M}_\omega(\overline{V}_i)$. But $r_i^\omega, r_i^{\prime\omega}$ have the same germ, and $s_i^\omega, s_i^{\prime\omega}$ have the same germ, so r_i^ω, s_i^ω have the same germ. This completes the proof of the claim.

We have shown that the intersection of two orthants is an orthant which is equal to a common face of the two. Also, the poset of nontrivial faces of a k -orthant, meaning all faces except the singleton $\{\mathbf{x}\}$, is isomorphic to the lattice of nonempty sets of components of a k -component surface, which is isomorphic to the poset of faces of a $k - 1$ simplex. Having excluded the unique 0-orthant $\{\mathbf{x}\}$ from \mathcal{O} , this completes the proof that \mathcal{O} has the structure of a simplicial complex \mathcal{K} with one simplex of dimension $k - 1$ for each k -orthant.

Now we show that \mathcal{K} is a flag complex.

Let O_1, \dots, O_k be distinct 1-orthants which represent vertices of a complete graph in \mathcal{K} , and so for each i, j there is a 2-orthant O_{ij} whose faces are O_i and O_j . Choose twist normalized representatives $O_i = \gamma(C^\omega(\bar{\mu}_i, \bar{W}_i, \bar{r}_i))$ where \bar{W}_i has a single component. Choose twist normalized representatives $O_{ij} = \gamma(C^\omega(\bar{\nu}_{ij}, \bar{V}_{ij}, \bar{s}_{ij}))$ where \bar{V}_{ij} has two components \bar{V}_i and \bar{V}_j . By item (2) of the lemma, for ω -a.e. n and each i, j the surfaces W_i^n, W_j^n are isotopic to distinct components of V_{ij}^n , and we may choose the notation so that $W_i^n \stackrel{i}{\cong} V_i^n$, $W_j^n \stackrel{j}{\cong} V_j^n$. It follows that W_i^n and W_j^n are disjoint and nonisotopic for n in a set I_{ij} of full ω -measure. The intersection $\cap I_{ij}$ over all (i, j) still has full ω -measure, so we conclude that W_1^n, \dots, W_k^n are pairwise disjoint and non-isotopic for ω -a.e. n , and we obtain an essential subsurface $\bar{W} = \bar{W}_1 \cup \dots \cup \bar{W}_k$. Let $\bar{\sigma}$ be a marking sequence on \bar{W}^c defined as the projection $\pi_{\mathcal{M}(\bar{W}^c)}(\bar{x})$, and let $\bar{r} = (\bar{r}_1, \dots, \bar{r}_k)$. Then we obtain an orthant

$$O = \gamma(C^\omega(\bar{\sigma}, \bar{W}, \bar{r})).$$

We need to check that the corner of O , namely the limit of $\bar{\kappa} = \{\bar{\sigma}\} \times \prod \bar{r}_i(0)$, equals \mathbf{x} . But this is a consequence of the quasidistance formula for $d(\kappa^n, x^n)$, in which we separate the terms $\{d_Z(\kappa^n, x^n)\}_A$ according to whether $Z \Subset (W^n)^c$, $Z \Subset W^n$, or $Z \cap \partial W^n$. The first type of term adds up to an estimate of $d_{\mathcal{M}(W^n)^c}(\sigma^n, x^n)$, which by definition of $\bar{\sigma}$ is bounded. The second type adds up to estimate the finite sum $\sum_i d_{\mathcal{M}(W_i^n)}(r_i^n(0), x^n)$, each of whose terms grows sublinearly since the corner of each 1-orthant O_i is \mathbf{x} . The third type is estimated, termwise, by $\{d_Z(\partial W^n, x^n)\}_A$, which sum up to estimate $d(x^n, \mathcal{Q}(\partial W^n))$, by Lemma 3.1. This again grows sublinearly since $\mathbf{x} \in \cap_i \mathcal{Q}_\omega(\partial \bar{W}_i) = \mathcal{Q}_\omega(\partial \bar{W})$. We conclude that $d(\kappa^n, x^n)$ grows sublinearly, so $\kappa^\omega = x^\omega$.

This tells us that $O \in \mathcal{O}$. It is clear by construction and item (2) of the lemma that O_i are the vertices of O . This completes the proof. \square

The following observation will be used later when we give a topological characterization of Dehn twist flats.

Corollary 9.2. *The germ of every Dehn twist ξ -orthant is contained in a unique Dehn twist flat germ.*

Proof. Consider two Dehn twist flat germs $Q_\omega(\bar{\nu}), Q_\omega(\bar{\nu}')$ whose intersection contains a common Dehn twist ξ -orthant O . Let \bar{W}, \bar{W}' be the sequences of annulus neighborhoods of the pants decomposition sequences $\bar{\nu}, \bar{\nu}'$. We obtain two twist normalized expressions $O = \gamma(C^\omega(\bar{\mu}, \bar{W}, \bar{r})) = \gamma(C^\omega(\bar{\mu}', \bar{W}', \bar{r}'))$. By Lemma 9.1 (1) the sequences \bar{W}, \bar{W}' are ω -equivalent, which implies that $Q_\omega(\bar{\nu}) = Q_\omega(\bar{\nu}')$. \square

9.3. Applying local finiteness

We continue to fix the base point $\mathbf{x} \in \mathcal{M}_\omega(S)$. Consider the subset of \mathcal{G} consisting of germs at \mathbf{x} of submanifolds of \mathcal{M}_ω of dimension $\xi = \xi(S)$. This subset generates a sublattice $\mathcal{F} \subset \mathcal{G}$ by taking finite unions and intersections. Since germs of manifolds are topologically characterizable, it follows that germs in \mathcal{F} are topologically characterizable. Our goal in this section is to produce finer topological properties in \mathcal{F} , in order to yield a topological characterization of germs of Dehn twist flats given in Corollary 9.8.

Let $\widehat{\mathcal{O}} \subset \mathcal{G}$ be the sublattice generated from the orthant complex \mathcal{O} at \mathbf{x} by taking finite unions and intersections. By Lemma 9.1, $\widehat{\mathcal{O}}$ is isomorphic to the lattice of finite subcomplexes of the simplicial complex \mathcal{K} .

The manifold local finiteness theorem, Theorem 8.7, will imply:

Lemma 9.3. $\mathcal{F} \subset \widehat{\mathcal{O}}$.

Proof. Let M be a manifold of dimension $\xi = \xi(S)$ passing through \mathbf{x} . Theorem 8.7 states that there is a neighborhood U of \mathbf{x} such that $M \cap U$ is contained in a finite union of cubes. After subdivision and possibly replacing U by a smaller open set containing \mathbf{x} , we obtain a finite collection of cubes C_1, \dots, C_I , each having \mathbf{x} as a corner, whose union contains $M \cap U$. Applying local compactness of M , choose an open set $V \subset U$ containing \mathbf{x} such that the closure of $M \cap V$ is a compact subset of $M \cap U$.

Suppose $M \cap V$ has nontrivial intersection with the interior of a cube C_i of dimension ξ . For $j \neq i \in \{1, \dots, I\}$ the intersection $C_i \cap C_j$ is contained in the boundary of C_j , by Lemma 9.1, and so $\text{int}(C_i)$ is disjoint from the closed set C_j . It follows that $M \cap \text{int}(C_i)$ is open in M . Combined with invariance of domain it follows that $\text{int}(C_i) \cap V \subset M$, and so $C_i \cap V \subset M$.

Suppose next that $M \cap V$ has nontrivial intersection with a cube C_j of dimension $< \xi$. Again by invariance of domain, $M \cap V$ must meet the interior of some cube C_i of dimension ξ having C_j as a face, and so $C_j \cap V \subset C_i \cap V \subset M$.

We conclude that any germ of a manifold is *equal to* a finite union of orthants, and hence $\mathcal{F} \subset \widehat{\mathcal{O}}$. \square

To clarify the structure of \mathcal{F} , we introduce some more objects.

Consider a top dimensional orthant O and a twist normalized representation $O = \gamma(C^\omega(\bar{\mu}, \bar{W}, \bar{r}))$. Let the components of \bar{W} be $\bar{W}_1, \dots, \bar{W}_\xi$. In each \bar{W}_i we have a ray \mathbf{r}_i in the associated \mathbf{R} -tree $\mathcal{M}_\omega(\bar{W}_i)$. Actually we only need to consider the *germ* of a ray, but we will still denote it \mathbf{r}_i . A component \bar{W}_i is called a *boundary annulus* if for ω -a.e. n the surface W_i^n is an annulus homotopic to the boundary of another component W_j^n , necessarily of complexity 1. Let $b(O) = b(\bar{W})$ denote the number of boundary annuli, which is a well-defined function of O by Lemma 9.1 (1). Note that $b(O) = 0$ if and only if all components of \bar{W} are annuli, if and only if O is a Dehn twist orthant, if and only if each \bar{r}_i is a twist direction — the first “if and only if” is a consequence of top dimensionality, the second is a matter of definition, and the last is a consequence of twist normalization.

Fix $j \in \{1, \dots, \xi\}$, and consider the j^{th} codimension 1 face of O , obtained by restricting the ray \mathbf{r}_j to its initial point. If $O' = \gamma(C^\omega(\bar{\mu}', \bar{W}', \bar{r}'))$ is another twist normalized top dimensional orthant meeting O along the j^{th} codimension 1 face, then by applying Lemma 9.1 (3) it follows that O' is obtained from O in one of the following ways: only the component \bar{W}_j is changed; or \bar{W}' and \bar{W} are equivalent and only the ray germ \mathbf{r}_j is changed (all of these changes are up to ω -equivalence). If \bar{W}_j is of complexity 1 then there are infinitely many different choices for O' , for example there are infinitely many different ray germs \mathbf{r}'_j to choose from in the \mathbf{R} -tree $\mathcal{M}_\omega(\bar{W}_j)$. If \bar{W}_j is a nonboundary annulus then there are also infinitely many different choices for O' , for instance there are infinitely many annuli \bar{W}'_j that can replace \bar{W}_j . However, if \bar{W}_j is a boundary annulus then the only change we can make is to replace \mathbf{r}_j by the unique opposite ray germ $-\mathbf{r}_j$ in the line $\mathcal{M}_\omega(\bar{W}_j)$.

We conclude that, along each of the $b = b(O)$ codimension-1 faces of O associated to boundary annuli, there is a unique orthant adjacent to O . It follows that any manifold germ M containing O must contain all of these unique neighboring orthants. Furthermore each of these orthants still has the same defining surface \bar{W} and the same set of b boundary annuli, and for all the corresponding faces the unique neighboring orthants must be included. We conclude that all 2^b orthants obtained in this way must be contained in the germ M . We call this set a *lune*, and refer to the number b as its *rank*. We note that it is naturally identified with a Euclidean spherical lune $\mathbb{R}^b \times (\mathbb{R}^+)^{\xi-b} \cap \mathbb{S}^{\xi-1}$, with its subdivision into Euclidean orthants (i.e., spherical simplices).

Lemma 9.4. *Lunes are precisely the minimal ξ -dimensional elements of the lattice \mathcal{F} .*

Proof. Consider L a lune of rank b . We first show that $L \in \mathcal{F}$, and that L is a minimal ξ -dimensional element of \mathcal{F} . As discussed above, L is a union of 2^b orthants, and without loss of generality the associated (germs of) rays for these

orthants are of the form

$$\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_b^{j_b}, \mathbf{r}_{b+1}^1, \dots, \mathbf{r}_\xi^1,$$

where $\overline{W}_1, \dots, \overline{W}_b$ are the boundary annuli in a decomposition \overline{W} , and for each $i \in \{1, \dots, b\}$ we have $j_i \in \{1, 2\}$ and the ray germs $\mathbf{r}_i^1, \mathbf{r}_i^2$ are opposite pairs.

For each $i \in \{b+1, \dots, \xi\}$ we define the following objects. First, choose \mathbf{r}_i^2 to be a ray in $\mathcal{M}_\omega(\overline{W}_i)$ with germ distinct from \mathbf{r}_i^1 ; if $\xi(\overline{W}_i) = 1$ make sure that \mathbf{r}_i^2 is not a twist direction. Next, let \overline{V}_i be \overline{W}_i if $\xi(\overline{W}_i) = 1$, and if $\xi(\overline{W}_i) = 0$ let \overline{V}_i be the unique sequence (up to the usual ultraproduct equivalence) of $\xi = 1$ subsurfaces containing \overline{W}_i and disjoint from all the other \overline{W}_k . We can interpret $\mathbf{r}_i^1, \mathbf{r}_i^2$ as ray germs in $\mathcal{M}_\omega(\overline{V}_i)$ via the natural embedding $\mathcal{M}_\omega(\overline{W}_i) \rightarrow \mathcal{M}_\omega(\overline{V}_i)$. Denote $\mathbf{s}_i^1 = \mathbf{r}_i^1$. Finally, choose \mathbf{s}_i^2 to be a ray germ in $\mathcal{M}_\omega(\overline{V}_i)$ which shares its basepoint with $\mathbf{r}_i^1, \mathbf{r}_i^2$ but is distinct from both. Let $\overline{W}[i]$ be the subsurface sequence obtained from \overline{W} by replacing \overline{W}_i by \overline{V}_i .

Now for each $i \in \{b+1, \dots, \xi\}$ and each tuple $\hat{j} = (j_1, \dots, j_\xi) \in \{1, 2\}^\xi$, consider the orthant $O[i](\hat{j})$ formed from $\overline{W}[i]$ and the ray germs

$$\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_b^{j_b}, \dots, \mathbf{r}_{i-1}^{j_{i-1}}, \mathbf{s}_i^{j_i}, \mathbf{r}_{i+1}^{j_{i+1}} \dots$$

in other words, we use $\mathbf{r}_k^{j_k}$ for all k except $k = i$, where we use $\mathbf{s}_i^{j_i}$. Let

$$M[i] = \bigcup_{\hat{j}} O[i](\hat{j}).$$

This is a manifold germ, and our lune L is the intersection

$$L = M[b+1] \cap \dots \cap M[\xi].$$

This shows that L is in \mathcal{F} . Since any ξ -dimensional element C of \mathcal{F} contained in L must contain a top dimensional orthant $O \subset L$, it follows from the paragraph before Lemma 9.4 that C contains L , and so L is minimal. Each lune is therefore a minimal ξ -dimensional elements of \mathcal{F} .

Now let C be any minimal ξ -dimensional element of \mathcal{F} . Then C must contain a ξ -dimensional orthant O , by Lemma 9.3. By the discussion above, C must contain the lune L determined by O , and by the minimality of C , we have $C = L$. Each minimal ξ -dimensional element of \mathcal{F} is therefore a lune. \square

We will also make use of the following:

Lemma 9.5. *For each lune L of rank b there exists a manifold germ containing L which is a union of $2^{\xi-b}$ distinct lunes of rank b , no two of which have an orthant in common.*

Proof. We borrow the notation of Lemma 9.4, but in this case the construction is somewhat easier. For each tuple $\hat{j} = (j_{b+1}, \dots, j_\xi) \in \{1, 2\}^{\xi-b}$ let $L(\hat{j})$ be the lune of rank b which is the union of the orthants associated to ray germs

$$\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_b^{j_b}, \mathbf{r}_{b+1}^{j_{b+1}}, \dots, \mathbf{r}_\xi^{j_\xi}$$

where (j_1, \dots, j_b) varies freely in $\{1, 2\}^b$. The union of these lunes, over all $\hat{j} \in \{1, 2\}^{\xi-b}$, is a manifold germ.

Finally, note that distinct lunes have no orthants in common, by the paragraph before Lemma 9.4. \square

Since $\widehat{\mathcal{O}}$ is isomorphic to the poset of finite subcomplexes of the $(\xi - 1)$ -dimensional simplicial complex \mathcal{K} of Lemma 9.1, each element $C \in \widehat{\mathcal{O}}$ determines a simplicial $(\xi - 1)$ -chain with \mathbb{Z}_2 -coefficients in \mathcal{K} , namely the formal sum of the simplices corresponding to the top dimensional orthants appearing in C . In what follows we will conflate the chain with C when convenient. Given two chains $\alpha, \beta \in C_{\xi-1}(\mathcal{K})$, we say that α is part of β if $\beta = \alpha + \alpha'$ where the chains α and α' have no simplices in common.

We let \mathcal{L} denote the collection of lunes, which by Lemma 9.4 are topologically characterized. Our next goal is a characterization of the rank of lunes as a function on \mathcal{L} :

Lemma 9.6. *The rank is the unique function $f: \mathcal{L} \rightarrow \{0, \dots, \xi\}$ with the following property:*

- *If C is a lune and if $f(C) \leq b \in \{0, \dots, \xi\}$, then $f(C) = b$ if and only if C is part of a nonzero cycle*

$$\sum_{i=1}^{2^{\xi-b}} C_i,$$

such that each C_i is a lune satisfying $f(C_i) \leq b$.

To prove this, we will need a lemma about flag complexes:

Lemma 9.7. *Every nontrivial reduced \mathbb{Z}_2 n -cycle in an n -dimensional flag complex has cardinality at least 2^{n+1} .*

Proof. The lemma obviously holds for 0-dimensional flag complexes, since the support of a nontrivial reduced 0-cycle must contain at least two vertices.

Assume inductively that $n = \dim X > 0$, and that the lemma holds for flag complexes of dimension $< n$. We first observe that the link of any vertex is an $(n - 1)$ -dimensional flag complex, and hence by the induction assumption, the lemma holds for links.

Let M be a \mathbb{Z}_2 n -cycle in X . Consider two adjacent n -simplices σ_1, σ_2 meeting at a codimension 1 face τ . Let v_i be the vertex of σ_i complementary to τ . The link of v_i in M is a \mathbb{Z}_2 $(n - 1)$ -cycle, hence by the assumption has cardinality at least 2^n . The lemma would follow if we show that the stars of v_1 and v_2 do not have common n -simplices.

Suppose there is such a simplex. Then v_1 and v_2 must be joined by an edge e . Now the abstract join $\tau * e$ is an $(n + 1)$ -simplex all of whose edges are in

the complex X . Since X is a flag complex, it must contain an $(n + 1)$ -simplex, but this contradicts $\dim X = n$. \square

Proof of Lemma 9.6. We will refer to the property stated in Lemma 9.6 as *Property S*.

We first show that the rank function has Property S. If C is a lune of rank b then, by Lemma 9.5, C is part of a cycle consisting of $2^{\xi-b}$ lunes of rank b . Conversely, suppose $b \in \{0, \dots, \xi\}$, C is a lune of rank $\leq b$, and C is part of a nonzero cycle $\sum_{i=1}^{2^{\xi-b}} C_i$ where each C_i is a lune of rank $\leq b$, so C_i is composed of $2^{\text{rank}(C_i)} \leq 2^b$ orthants, and the entire cycle is composed of $\leq 2^\xi$ orthants. Nonzero cycles require at least 2^ξ orthants by Lemma 9.7, which implies that $\text{rank}(C_i) = b$ for all i . Since C is part of $\sum_i C_i$, the intersection $C \cap C_j$ must contain a top dimensional orthant for some j . The minimality property for lunes, Lemma 9.4, implies that $C = C_j$ and hence $\text{rank}(C) = \text{rank}(C_j) = b$. This shows that rank has Property S.

Now suppose $f : \mathcal{L} \rightarrow \{0, \dots, \xi\}$ has Property S, but is not equal to rank. Let b be the maximum of the integers $\bar{b} \in \{0, \dots, \xi\}$ such that $f^{-1}(\bar{b}) \neq \text{rank}^{-1}(\bar{b})$.

Suppose C is a lune of rank b . Then C belongs to a cycle $\sum_{i=1}^{2^{\xi-b}} C_i$ where $\text{rank}(C_i) = b$. By choice of b we have $f(C_i) \leq b$. Hence by Property S, we get $f(C) = b$. Thus $\text{rank}^{-1}(b) \subset f^{-1}(b)$.

Now suppose $C \in f^{-1}(b)$. Then C belongs to a cycle $\sum_{i=1}^{2^{\xi-b}} C_i$ where $f(C_i) \leq b$. By the choice of b , we have $\text{rank}(C_i) \leq b$ for all i , and by Lemma 9.7 we get $\text{rank}(C_i) = b$ for all i . We conclude as above, using Lemma 9.4, that $C = C_j$ for some j , and hence $\text{rank}(C) = b$. Thus $f^{-1}(b) \subset \text{rank}^{-1}(b)$. This contradicts the choice of b . \square

All the pieces are now in place for the main result of this section:

Corollary 9.8. *There is a topological characterization of Dehn twist orthant germs and Dehn twist flat germs in $\mathcal{M}_\omega(S)$.*

Proof. Membership in the lattice \mathcal{F} is topologically characterized, as is the dimension function on \mathcal{F} . Therefore Lemmas 9.4 and 9.6 give topological characterizations of lunes and lune rank. Dehn twist ξ -orthants are the lunes of rank 0, so these are also topologically characterizable. Dehn twist orthants of arbitrary dimension are topologically characterized as intersections of sets of Dehn twist ξ -orthants.

It remains to topologically characterize Dehn twist flat germs — those configurations of 2^ξ orthants associated to a Dehn-twist flat through \mathbf{x} . This boils down to finding a topological characterization of those antipodal pairs of Dehn twist 1-orthants (also known as Dehn twist vertices) that occur in Dehn twist flat germs.

Consider a twist normalized Dehn twist vertex $O = \gamma(C^\omega(\bar{\mu}, \bar{W}, \bar{r}))$, so W^n is connected and is an annulus for ω -a.e. n . The *antipodal point* of O is defined to be $-O = \gamma(C^\omega(\bar{\mu}, \bar{W}, -\bar{r}))$, where \bar{r} and $-\bar{r}$ have the same initial point but opposite directions in the line $\mathcal{TM}(\bar{W})$. Note that the antipodal point of O is well-defined, for suppose that $O = \gamma(C^\omega(\bar{\mu}', \bar{W}', \bar{r}'))$ is another twist normalized expression. By Lemma 9.1 (1), for ω -a.e. n we have $W^n \stackrel{i}{\cong} W'^n$ and \mathbf{r}, \mathbf{r}' have the same direction in the line $\mathcal{M}_\omega(\bar{W}) = \mathcal{M}_\omega(\bar{W}')$. Moreover, the asymptotic partial markings $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{M}_\omega(\bar{W}^c)$ are each equal to the projection of \mathbf{x} . It follows that the two asymptotic cubes $C^\omega(\bar{\mu}, \bar{W}, -\bar{r}), C^\omega(\bar{\mu}', \bar{W}', -\bar{r}')$ are equal in $\mathcal{M}_\omega(S)$ and so the two expressions for the antipodal point $\gamma(C^\omega(\bar{\mu}, \bar{W}, -\bar{r}))$ and $\gamma(C^\omega(\bar{\mu}', \bar{W}', -\bar{r}'))$ define the same Dehn twist vertex $-O$.

This shows that the relation of “antipodal point” gives a decomposition of the set of Dehn twist vertices into pairs. We now show that antipodal pairs are topologically characterizable.

Consider a lune of positive rank $b > 0$. The corresponding decomposition \bar{W} contains b boundary annuli associated to which are b antipodal pairs of Dehn twist vertices that span a sphere of dimension $b - 1$ in the orthant complex, which is subdivided in the standard way into 2^b simplices.

These lune boundary spheres are topologically characterizable: they are precisely the $(b - 1)$ -dimensional spheres which may be obtained as the intersection of two lunes of rank b . Moreover the simplicial decomposition of such a sphere is topologically characterizable, since the simplices are Dehn-twist orthants. Hence the pairs of antipodal vertices of lune boundary spheres are topologically characterizable in terms of this simplicial structure.

Now any Dehn twist vertex of \mathcal{O} can be placed into such a lune boundary sphere, simply by extending its defining annulus to a decomposition where it is a boundary annulus. Thus antipodal pairs of Dehn twist vertices may be characterized topologically as those vertex pairs which may be embedded as antipodal vertices in a triangulated lune boundary sphere.

To summarize, germs of Dehn twist flats are topologically characterized as those flag subcomplexes $C \subset \mathcal{O}$ whose vertex set decomposes into a union of ξ antipodal pairs of (twist normalized) Dehn twist vertices

$$\pm O_i = \gamma(C^\omega(\bar{\mu}_i, \bar{W}_i, \pm \bar{r}_i)), \quad i = 1, \dots, \xi$$

such that any choice of one vertex from each pair spans a ξ -orthant: C is the union of those 2^ξ orthants; by Lemma 9.1 (2) applied to any such orthant the annuli $\bar{W}_1, \dots, \bar{W}_\xi$ form a sequence of pants decompositions \bar{W} ; and the Dehn twist flat germ associated to \bar{W} is C . \square

9.4. Characterizing Dehn twist flats

We conclude Section 9 with the following observation, which gives a local topological characterization of Dehn twist flats.

Lemma 9.9. *Suppose $E \subset \mathcal{M}_\omega(S)$ is a connected top-dimensional manifold, and that for every $\mathbf{x} \in E$, the germ of E at \mathbf{x} is the germ of a Dehn twist flat. Then E is contained in a Dehn twist flat. If in addition E is a closed subset of $\mathcal{M}_\omega(S)$, then E is a Dehn twist flat.*

Proof. Pick any Dehn twist flat $E' \subset \mathcal{M}_\omega$ such that the interior in E of $E \cap E'$ is a nonempty set $U \subset E$. Suppose $\mathbf{x} \in E$ lies in the closure of U in \mathcal{M}_ω . Since E' is a closed subset of $\mathcal{M}_\omega(S)$, we have $\mathbf{x} \in E \cap E'$. By the definition of U , the germ of $E \cap E'$ at \mathbf{x} has dimension ξ , so by Corollary 9.2 it follows that E, E' have the same germ at \mathbf{x} , and we conclude that $\mathbf{x} \in U$. Thus U is an open and closed subset of E . Since E is connected, we have $E = U \subset E \cap E' \subset E'$.

If E is a closed subset of $\mathcal{M}_\omega(S)$, then $E \cap E'$ is open and closed in E' , and hence $E \cap E' = E'$. \square

10. Finishing the proofs

We are now ready to prove our main theorems on quasi-isometric classification and rigidity. The proof will follow the general sketch from the introduction. Let us first state the theorems in their complete form:

Theorem 1.1 (Classification of Quasi-Isometries) *Suppose that $\xi(S) \geq 2$.*

If $S \neq S_{1,2}$ then quasi-isometries of $\mathcal{MCG}(S)$ are uniformly close to isometries induced by left-multiplication.

That is, given $K, \delta > 0$ there exists $D > 0$ such that, if $f: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(S)$ is a (K, δ) -quasi-isometry then there exists $g \in \mathcal{MCG}(S)$ such that

$$d(f(x), L_g(x)) \leq D \quad \text{for all } x \in \mathcal{MCG}(S)$$

where L_g is left-multiplication by g .

If $S = S_{1,2}$ then the same result holds if we replace L_g by a quasi-isometry of $\mathcal{MCG}(S_{1,2})$ induced by an element $g \in \mathcal{MCG}(S_{0,5})$ via the standard index 5 embedding $\mathcal{MCG}(S_{1,2})/Z(\mathcal{MCG}(S_{1,2})) \hookrightarrow \mathcal{MCG}(S_{0,5})$.

Here $Z(G)$ denotes the center of G , which is trivial for all mapping class groups with $\xi(S) > 2$; furthermore, in the cases where it is non-trivial, it is finite.

Theorem 1.2 (Quasi-Isometric Rigidity) *If Γ is a finitely generated group quasi-isometric to $\mathcal{MCG}(S)$, then there exists a finite-index subgroup $\Gamma' < \Gamma$ and a homomorphism*

$$\Gamma' \rightarrow \mathcal{MCG}(S)/Z(\mathcal{MCG}(S))$$

with finite kernel and finite index image.

In the cases of mapping class groups of complexity less than 2, quasi-isometric rigidity is either trivial or well known. On the other hand, the analogue of Theorem 1.1 fails in the case of complexity one, as there are quasi-isometries of the free group which are not a bounded distance from isometries.

An immediate consequence of Theorem 1.1 is the following characterization of the quasi-isometry group $\text{QI}(\mathcal{MCG}(S))$, i.e. the group of quasi-isometries of $\mathcal{MCG}(S)$ modulo those that are finite distance from the identity.

Corollary 10.1. *If S has complexity at least 2 then the natural homomorphism*

$$\mathcal{MCG}(S)/Z(\mathcal{MCG}(S)) \rightarrow \text{QI}(\mathcal{MCG}(S))$$

is an isomorphism except when $S = S_{1,2}$, in which case it is an isomorphism to a subgroup of index 5.

The proof of this corollary is embedded in the proof of Theorem 1.2 in §10.4 below.

10.1. Preservation of asymptotic Dehn twist flats

We begin the proof by showing that the local topological characterization of Dehn twist flats in the asymptotic cone, given in Lemma 9.9, implies a global characterization:

Theorem 10.2. *If $\xi(S) \geq 2$, any homeomorphism $f: \mathcal{M}_\omega(S) \rightarrow \mathcal{M}_\omega(S)$ permutes the Dehn twist flats in $\mathcal{M}_\omega(S)$.*

Proof. By Corollary 9.8, any homeomorphism must preserve the set of Dehn twist flat germs in $\mathcal{M}_\omega(S)$ (with arbitrary basepoints) since they are topologically characterized. It follows that, at every point in the image $f(E)$ of a Dehn twist flat E , its germ is equal to the germ of a Dehn twist flat. Lemma 9.9 therefore implies that $f(E)$ is itself a Dehn twist flat. \square

10.2. Coarse preservation of Dehn twist flats

We next descend to the group itself, where we show that a quasi-isometry of $\mathcal{M}(S)$ coarsely preserves Dehn twist flats, in a uniform sense.

Theorem 10.3. *If $\xi(S) \geq 2$, then given $K \geq 1$ and $C \geq 0$ there exists A such that, if $f: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(S)$ is a (K, C) -quasi-isometry and E is a Dehn twist flat in $\mathcal{MCG}(S)$ then there exists a unique Dehn twist flat E' such that the Hausdorff distance between $f(E)$ and E' is at most A .*

Proof. Uniqueness of E' follows from the fact that distinct Dehn twist flats have infinite Hausdorff distance, an immediate consequence of Proposition 3.1 and Lemma 3.10.

The existence proof is essentially an argument by contradiction, using Theorem 10.2. If there is no uniform control of the Hausdorff distance between Dehn twist flats and their quasi-isometric images, while on the other hand in every limiting situation the Dehn twist flats are preserved in the asymptotic cone, then in a sequence of counterexamples we can carefully select basepoints and scales to get configurations in which the image of a Dehn twist flat is simultaneously very close to two distinct Dehn twist flats. This contradicts the fact that distinct Dehn twist flats look different at all scales (Lemma 10.5, which is a consequence of Lemma 3.3).

We will find the following variation of the Hausdorff metric useful. Given subsets A, B of a metric space X and a point $p \in X$, for each $r > 0$ we define

$$D_{r,p}(A, B) = \inf\{s \geq 0 \mid A \cap \mathcal{N}_r(p) \subset \mathcal{N}_s(B) \text{ and } B \cap \mathcal{N}_r(p) \subset \mathcal{N}_s(A)\}$$

Notice that, if $\mathcal{N}_r(p) \subset \mathcal{N}_{r'}(p')$, then $D_{r,p} \leq D_{r',p'}$. This is not quite a distance function — it fails the triangle inequality — but it does give a useful criterion for equality of ultralimits.

In the following lemma we consider a sequence (X_i, p_i) of based metric spaces with ultralimit X_ω . No rescaling is assumed here; in applications below, X_i will be the rescaled $\mathcal{MCG}(S)$.

Lemma 10.4. *Given $(A_i), (B_i)$ two sequences of closed subsets, $A_\omega = B_\omega$ if and only if for each basepoint (q_i) and for each $r \geq 0$ the ultralimit of $D_{r,q_i}(A_i, B_i)$ equals zero.*

Proof. Suppose that for some (q_i) and some $r > 0$ we have $D_{r,q_i}(A_i, B_i) \rightarrow_\omega \epsilon \in (0, \infty]$. Choose $\eta \in (0, \epsilon)$. It follows that, for ω -a.e. i , one of the following is true:

- (1): $A_i \cap \mathcal{N}_r(q_i) \not\subset \mathcal{N}_\eta(B_i)$
- (2): $B_i \cap \mathcal{N}_r(q_i) \not\subset \mathcal{N}_\eta(A_i)$

Furthermore, either (1) is true for ω -a.e. i , or (2) is true for ω -a.e. i ; let us assume the former. From this we conclude that there is a sequence

$$x_i \in A_i \cap \mathcal{N}_r(q_i) \setminus \mathcal{N}_\eta(B_i)$$

for which $x_\omega \in A_\omega$ but the distance between x_ω and B_ω is at least η . Hence $A_\omega \neq B_\omega$.

Suppose next that $\lim_\omega D_{r,q_i}(A_i, B_i) = 0$ for all (q_i) and all $r \geq 0$. To prove that $A_\omega \subset B_\omega$, consider $x_\omega \in A_\omega$ represented by a sequence (x_i) at bounded distance from (q_i) , so there exists some $r \geq 0$ such that $x_i \in A_i \cap \mathcal{N}_r(q_i)$ ω -almost surely. For any integer $k > 0$ it follows that $x_i \in \mathcal{N}_{1/k}(B_i)$ ω -almost surely, so we may choose a sequence $y_i^k \in B_i$ such that $d_i(x_i, y_i^k) < 1/k$ ω -almost surely, and therefore $y_\omega^k \in B_\omega$ and $d(x_\omega, y_\omega^k) \leq 1/k$. The sequence y_ω^k therefore converges to x_ω , but this sequence is in the closed set B_ω , proving that $x_\omega \in B_\omega$. A symmetric argument proves that $B_\omega \subset A_\omega$. \square

We turn now to the proof of Theorem 10.3. Suppose that the theorem is false. Then we may fix $K \geq 1$, $C \geq 0$ so that the following is true: for any $A \geq 0$ there is a (K, C) -quasi-isometry $f: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(S)$, and a Dehn twist flat \mathcal{F} , such that for any Dehn twist flat \mathcal{F}' , the Hausdorff distance between $f(\mathcal{F})$ and \mathcal{F}' is greater than A .

From this symmetric statement we make the further asymmetric conclusion that for each $s > 0$ there is a (K, C) -quasi-isometry $f: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(S)$ and a Dehn twist flat F such that for all Dehn twist flats F' we have

$$f(F) \not\subset \mathcal{N}_s(F')$$

For if not, then there exists $s > 0$ such that for all (K, C) -quasi-isometries f and all Dehn twist flats F there exists a Dehn twist flat F' such that $f(F) \subset \mathcal{N}_s(F')$. The closest point projection π from the (K, C) -quasiflat $f(F)$ to the Dehn twist flat F' moves points a distance at most s and can therefore be regarded as a (K', C') -quasi-isometry from \mathbf{R}^n to \mathbf{R}^n for constants K', C' that depend only on K, C, s . Now any (K', C') -quasi-isometric embedding from \mathbf{R}^n into \mathbf{R}^n is a (K', C'') -quasi-isometry where C'' depends only on K', C' , and n ; see, for example, [34]. It follows that π is uniformly onto, meaning that there exists a constant B depending only on K', C' such that F' is in the B neighborhood of $\pi(f(F))$, and so $F' \subset \mathcal{N}_{s+B}(f(F))$. This shows that $f(F)$ and F' have Hausdorff distance at most $s + B$, which is a contradiction for $A > s + B$.

Fix a sequence s_i diverging to $+\infty$, a sequence of (K, C) -quasi-isometries $f_i: \mathcal{MCG}(S) \rightarrow \mathcal{MCG}(S)$, and a sequence of Dehn twist flats F_i , such that for all i and all Dehn twist flats F' we have

$$f_i(F_i) \not\subset \mathcal{N}_{s_i}(F') \tag{10.1}$$

Since there are finitely many $\mathcal{MCG}(S)$ -orbits of Dehn twist flats, by pre-composing with elements of $\mathcal{MCG}(S)$ and extracting a subsequence we may assume that the F_i take a constant value F . Fix a base point $p_0 \in F$. By post-composing with elements of $\mathcal{MCG}(S)$ we may assume that $f_i(p_0) = p_0$, and in particular $p_0 \in f_i(F)$, for all i . We may therefore pass to the asymptotic cone with base point p_0 and scaling sequence s_i producing a bi-Lipschitz homeomorphism $f_\omega: \mathcal{M}_\omega(S) \rightarrow \mathcal{M}_\omega(S)$ and a Dehn twist flat $F_\omega = \lim_\omega(F)$, the asymptotic cone of F . Applying Theorem 10.2 we obtain a Dehn twist flat $F'_\omega = \lim_\omega F'_i$ such that $f_\omega(F_\omega) = F'_\omega$. By applying Lemma 10.4 it follows that, fixing any $R > 0$,

$$\frac{1}{s_i} D_{Rs_i, p_0}(f_i(F), F'_i) \rightarrow_\omega 0. \tag{10.2}$$

On the other hand, (10.1) implies that there is a point $q_i \in f_i(F) - \mathcal{N}_{s_i}(F'_i)$, and so for any $r > 0$ the following statement is always true:

$$D_{r, q_i}(f_i(F), F'_i) > s_i \tag{10.3}$$

In order to get a contradiction out of (10.2) and (10.3) we shall reapply Theorem 10.2 to a properly chosen sequence of intermediate basepoints, near which $f_i(F)$ is still close to F'_i , but sufficiently far that another Dehn twist flat, F''_i ,

is also close to it. The contradiction will then come from the following rigidity property of Dehn twist flats, which expresses that any Dehn twist flat spreads away from any other one in a linear fashion:

Lemma 10.5. *There exists $\epsilon_1 \in (0, 1)$ such that for any sufficiently large r , any $x \in \mathcal{MCG}(S)$, and any Dehn twist flats F_1, F_2 , if $F_1 \cap \mathcal{N}_{\frac{r}{2}}(x) \neq \emptyset$ and if*

$$F_1 \cap \mathcal{N}_r(x) \subset \mathcal{N}_{\epsilon_1 r}(F_2)$$

then $F_1 = F_2$.

Proof. Express F_i as $\mathcal{Q}(P_i)$ for a pants decomposition P_i ($i = 1, 2$). Lemma 3.3 implies that the junctures of $\mathcal{Q}(P_1)$ and $\mathcal{Q}(P_2)$ are $E_1 = \mathcal{Q}(P_1|P_2)$ and $E_2 = \mathcal{Q}(P_2|P_1)$. Note that $P_1|P_2$ is a marking with base equal to P_1 , and a transversal for each component of P_1 that is not a curve of P_2 . Hence, assuming $F_1 \neq F_2$, E_1 and E_2 must be subflats of strictly smaller dimension. The pair (F_i, E_i) is therefore uniformly quasi-isometric to the pair $(\mathbf{R}^\xi, \mathbf{R}^k)$ for some $k < \xi = \xi(S)$.

If F_1 meets $\mathcal{N}_{\frac{r}{2}}(x)$ then pick $y \in F_1$ such that $d(x, y) \leq \frac{r}{2}$. Lemma 3.3 implies that for any $z \in F_1$ the distance $d(z, F_2)$ is bounded below (up to uniform coarse-Lipschitz error) by the distance $d(z, E_1)$. Now use the elementary fact that for any $y \in \mathbf{R}^\xi$ and any $R > 0$ the Euclidean R -ball around y contains a point z that is not contained in the Euclidean $\frac{R}{2}$ -neighborhood of \mathbf{R}^k . After adjusting for the multiplicative errors, and setting r large enough to overcome the additive errors, we find that for suitable $\epsilon'_1, \epsilon_1 > 0$ we have

$$F_1 \cap \mathcal{N}_{\frac{r}{2}}(y) \not\subset \mathcal{N}_{\epsilon'_1 r}(E_1)$$

and so

$$F_1 \cap \mathcal{N}_{\frac{r}{2}}(y) \not\subset \mathcal{N}_{\epsilon_1 r}(F_2).$$

Since $\mathcal{N}_{\frac{r}{2}}(y)$ is contained in $\mathcal{N}_r(x)$, this is what we wanted to prove. \square

Let $\epsilon_2 = \min\{\frac{1}{R}, \frac{\epsilon_1}{2}\}$. Applying (10.2) and (10.3), ω -almost surely the following two statements are true:

$$\begin{aligned} D_{Rs_i, p_0}(f_i(F), F'_i) &\leq \epsilon_2 Rs_i \\ D_{Rs_i, q_i}(f_i(F), F'_i) &> s_i \geq \epsilon_2 Rs_i \end{aligned}$$

For each i , consider a sequence $(p_{i,k})$ starting at p_0 and ending with q_i , with step size $d(p_{i,k}, p_{i,k+1}) \leq 1$. There must be some j such that, labeling $x_i = p_{i,j}$ and $x'_i = p_{i,j+1}$,

$$D_{Rs_i, x_i}(f_i(F), F'_i) \leq \epsilon_2 Rs_i \tag{10.4}$$

but such that

$$D_{Rs_i, x'_i}(f_i(F), F'_i) \geq \epsilon_2 Rs_i \tag{10.5}$$

Now assuming $\epsilon_2 R s_i > 1$ (which is true for large enough i), we have $\mathcal{N}_{R s_i}(x'_i) \subset \mathcal{N}_{R s_i(1+\epsilon_2)}(x_i)$ and hence (10.5) implies

$$D_{R s_i(1+\epsilon_2), x_i}(f_i(F), F'_i) \geq \epsilon_2 R s_i \quad (10.6)$$

Now we apply Theorem 10.2 again, this time using (x_i) as the basepoints, and we conclude via Lemma 10.4 that there exists a sequence (F''_i) of Dehn twist flats such that, for ω -almost every i ,

$$D_{R s_i(1+\epsilon_2), x_i}(f_i(F), F''_i) < \epsilon_2 R s_i \quad (10.7)$$

and in particular $F''_i \neq F'_i$ for ω -a.e. i , by (10.6). Now (10.7) implies in particular that

$$f_i(F) \cap B_{R s_i(1+\epsilon_2)}(x_i) \subset \mathcal{N}_{\epsilon_2 R s_i}(F''_i).$$

Moreover by (10.4) we have

$$F'_i \cap \mathcal{N}_{R s_i}(x_i) \subset \mathcal{N}_{\epsilon_2 R s_i}(f_i(F))$$

and moreover (by triangle inequality)

$$F'_i \cap \mathcal{N}_{R s_i}(x_i) \subset \mathcal{N}_{\epsilon_2 R s_i}(f_i(F) \cap \mathcal{N}_{R s_i(1+\epsilon_2)}(x_i)).$$

Putting this together we see

$$F'_i \cap \mathcal{N}_{R s_i}(x_i) \subset \mathcal{N}_{2\epsilon_2 R s_i}(F''_i). \quad (10.8)$$

Now since $x_i \in f_i(F_i)$ we note that $\mathcal{N}_{\epsilon_2 R s_i}(x_i)$, which is contained in $\mathcal{N}_{\frac{R s_i}{2}}(x_i)$, intersects F'_i nontrivially. Now (10.8) implies, by Lemma 10.5 (noting $2\epsilon_2 \leq \epsilon_1$), that $F'_i = F''_i$, a contradiction. \square

10.3. Quasi-isometry classification

We are now ready to prove Theorem 1.1, the classification of quasi-isometries of $\mathcal{MCG}(S) \approx \mathcal{M}(S)$. The first step is to show that a (K, C) -quasi-isometry $f: \mathcal{M}(S) \rightarrow \mathcal{M}(S)$ induces an automorphism of the curve complex $\phi: \mathcal{C}(S) \rightarrow \mathcal{C}(S)$.

First we define ϕ on the vertex set $\mathcal{C}_0(S)$. Given a vertex $c \in \mathcal{C}_0(S)$, choose a Dehn twist 1-flat $\mathcal{Q}(\mu)$ representing c , as defined in Section 3.5. We claim that there exists a Dehn twist 1-flat $\mathcal{Q}(\nu)$ at finite Hausdorff distance from $f(\mathcal{Q}(\mu))$. From this claim, choose a vertex $d \in \mathcal{C}_0(S)$ represented by $\mathcal{Q}(\nu)$, define $\phi(c) = d$, and apply Lemma 3.10 (2) to conclude that $\phi(c)$ is well-defined. To prove the claim, first apply Lemma 3.11 to obtain two maximal Dehn twist flats $\mathcal{Q}(\mu_0)$, $\mathcal{Q}(\mu_1)$ whose coarse intersection is represented by $\mathcal{Q}(\mu)$. Then apply Theorem 10.3 to obtain Dehn twist flats $\mathcal{Q}(\mu'_0)$, $\mathcal{Q}(\mu'_1)$ at finite Hausdorff distance from $f(\mathcal{Q}(\mu))$, $f(\mathcal{Q}(\mu_1))$ respectively. The coarse intersection of $\mathcal{Q}(\mu'_0)$ and $\mathcal{Q}(\mu'_1)$ is represented by $f(\mathcal{Q}(\mu))$, and also by $\mathcal{Q}(\mu'_0] \mu'_1)$ according to Lemma 3.11. It follows that $\mathcal{Q}(\mu'_0] \mu'_1)$ is quasi-isometric to $f(\mathcal{Q}(\mu))$ which is quasi-isometric to \mathbf{R} , and so $\mathcal{Q}(\mu'_0] \mu'_1)$ is a Dehn twist 1-flat.

The proof of the claim in the previous paragraph also shows for any $k = 0, \dots, \xi$ that the image under f of any Dehn twist k -flat is coarsely equivalent

to a Dehn twist k -flat. The dimension k only plays a role in the final sentence of the paragraph, in which \mathbf{R} is replaced by \mathbf{R}^k .

To prove that ϕ is a bijection of $\mathcal{C}_0(S)$, apply the same process to a coarse inverse $\bar{f}: \mathcal{M}(S) \rightarrow \mathcal{M}(S)$ to obtain a map $\bar{\phi}: \mathcal{C}_0(S) \rightarrow \mathcal{C}_0(S)$. The curve $c'' = \bar{\phi}(c')$ is represented by a Dehn twist 1-flat $Q(\mu'')$ at finite Hausdorff distance from $\bar{f}(Q(\mu'))$. The latter is at finite Hausdorff distance from $\bar{f}(f(Q(\mu)))$, which is at finite Hausdorff distance from $Q(\mu)$ because f and \bar{f} are coarse inverses. The Dehn twist 1-flat $Q(\mu)$ therefore represents both c and c'' , and it follows that $c = c''$ by applying Lemma 3.10 (2). This shows that $\bar{\phi}\phi$ is the identity, and a similar proof shows that $\phi\bar{\phi}$ is the identity.

We next show that two vertices $c_0, c_1 \in \mathcal{C}(S)$ are endpoints of an edge of $\mathcal{C}(S)$ if and only if $\phi(c_0), \phi(c_1)$ are endpoints of an edge of $\mathcal{C}(S)$. We need only prove the “only if” direction, the converse following by the same argument applied to a coarse inverse for f . Assuming c_0, c_1 are endpoints of an edge in $\mathcal{C}(S)$, this edge is represented by some Dehn twist 2-flat $\mathcal{Q}(\mu)$. The image $f(\mathcal{Q}(\mu))$ is coarsely equivalent to some Dehn twist 2-flat $\mathcal{Q}(\mu')$ which represents an edge with endpoints $d_0, d_1 \in \mathcal{C}(S)$. We must show that $\{\phi(c_0), \phi(c_1)\} = \{d_0, d_1\}$. Let $\mathcal{Q}(\mu_0), \mathcal{Q}(\mu_1)$ be Dehn twist 1-flats representing c_0, c_1 , and let $\mathcal{Q}(\mu'_0), \mathcal{Q}(\mu'_1)$ be Dehn twist 1-flats representing $\phi(c_0), \phi(c_1)$. By Lemma 3.10 (1) each of $\mathcal{Q}(\mu_0), \mathcal{Q}(\mu_1)$ is coarsely contained in $\mathcal{Q}(\mu)$, and so each of $\mathcal{Q}(\mu'_0), \mathcal{Q}(\mu'_1)$ is coarsely contained in $\mathcal{Q}(\mu')$, and by Lemma 3.10 (1) again it follows that $\phi(c_0), \phi(c_1) \in \{d_0, d_1\}$. Since $c_0 \neq c_1$ it follows that $\mathcal{Q}(\mu_0), \mathcal{Q}(\mu_1)$ are not coarsely equivalent, so $\mathcal{Q}(\mu'_0), \mathcal{Q}(\mu'_1)$ are not coarsely equivalent, so $\phi(c_0) \neq \phi(c_1)$, from which it follows that $\{\phi(c_0), \phi(c_1)\} = \{d_0, d_1\}$.

This finishes the description of the automorphism $\phi: \mathcal{C}(S) \rightarrow \mathcal{C}(S)$ induced by f .

Now suppose that $S \neq S_{1,2}$, returning to the case of $S_{1,2}$ in the end. The theorem of Ivanov, Korkmaz, and Luo shows that ϕ is induced by some mapping class $\Phi \in \mathcal{MCG}(S)$ [33, 36, 37]. We must show that the quantity

$$d(f, \Phi) = \inf_{x \in \mathcal{M}(S)} d(f(x), \Phi(x))$$

is bounded by a constant depending only on K, C , and the topology of S .

Consider a Dehn twist flat $Q(\mu)$. Theorem 10.3 tells us that the Hausdorff distance between $f(Q(\mu))$ and some Dehn twist flat $Q(\mu')$ is bounded uniformly in terms of K, C , and the topology of S . We claim that $\mu' = \Phi(\mu)$. To prove this let $\mu = \{c_1, \dots, c_\xi\}$, and let $c'_i = \Phi(c_i)$, and so $\Phi(\mu) = \{c'_1, \dots, c'_\xi\}$. Let $Q(\mu_i)$ be Dehn twist 1-flats representing $c_i, i = 1, \dots, \xi$, and so $Q(\mu_i)$ is coarsely contained in $Q(\mu)$. The image $f(Q(\mu_i))$ has finite Hausdorff distance from some Dehn twist 1-flat $Q(\mu'_i)$ representing c'_i , and so $Q(\mu'_i)$ is coarsely contained in $Q(\mu')$ as well as in $Q(\Phi(\mu))$. It follows that $c'_i \in \mu'$, and so $\Phi(\mu) \subset \mu'$. But $\mu', \Phi(\mu)$ are sets of the same cardinality ξ , and so $\Phi(\mu) = \mu'$, proving the claim.

We now find a uniform bound to $d(f(x), \Phi(x))$ for each $x \in \mathcal{M}(S)$. From Lemma 3.11 we obtain two Dehn twist flats $\mathcal{Q}(\mu_0), \mathcal{Q}(\mu_1)$ whose coarse intersection is represented uniformly by the point x , with constants depending only on the topology of S . Applying Φ , we obtain pants decompositions $\mu'_i = \Phi(\mu_i)$ and the coarse intersection of the corresponding Dehn twist flats $\mathcal{Q}(\mu'_0), \mathcal{Q}(\mu'_1)$ is represented uniformly by the point $\Phi(x)$ in terms of the topology of S only. Letting $\nu'_0 = \mu'_0 \rfloor \mu'_1$ and $\nu'_1 = \mu'_1 \rfloor \mu'_0$, the junctures $\mathcal{Q}(\nu'_0), \mathcal{Q}(\nu'_1)$ also represent the coarse intersection, and so must be single points ν'_0, ν'_1 respectively. Choose points $y_i \in \mathcal{Q}(\mu'_i)$ whose distance to $\Phi(x)$ is $\leq C$ depending only on the topology of S , and so $d(y_0, y_1) \leq 2C$. By Lemma 3.3 (3) we obtain a bound on $d(y_0, \nu'_0)$, and so also a bound on $d(\Phi(x), \nu'_0)$, depending only on the topology on S .

On the other hand, applying the map f , and using the bound on the Hausdorff distance between $f(\mathcal{Q}(\mu_i))$ and $\mathcal{Q}(\mu'_i)$, the point $f(x)$ also uniformly represents the coarse intersection of $\mathcal{Q}(\mu'_0)$ and $\mathcal{Q}(\mu'_1)$, the bounds now being in terms of K, C , and the topology of S . The same argument as in the last paragraph, using Lemma 3.3 (3), now produces a bound on $d(f(x), \nu'_0)$ depending only on K, C , and the topology of S .

Combining the last two paragraphs, we obtain the desired bound on $d(f(x), \Phi(x))$, completing the proof when $S \neq S_{1,2}$.

In the case of $S = S_{1,2}$, we still get an automorphism $\phi: \mathcal{C}(S_{1,2}) \rightarrow \mathcal{C}(S_{1,2})$ as before, but it may no longer be induced by a mapping class of $S_{1,2}$. This is a finite-index problem caused by the hyperelliptic involution $\tau: S_{1,2} \rightarrow S_{1,2}$ which we now recall. See Luo [37].

The map τ interchanges the two punctures of $S_{1,2}$. The quotient of $S_{1,2}$ by τ , minus the four branch points, is $S_{0,5}$. Since τ is central in $\mathcal{MCG}(S_{1,2})$, every element descends to $S_{0,5}$ and we get a map $\beta: \mathcal{MCG}(S_{1,2}) \rightarrow \mathcal{MCG}(S_{0,5})$ whose kernel is the center $\langle \tau \rangle \approx \mathbb{Z}/2\mathbb{Z}$. The image has index 5, because an element of $\mathcal{MCG}(S_{0,5})$ lifts if and only if it preserves the puncture which is the image of the two punctures of $S_{1,2}$. Hence β is a quasi-isometry, and we let $\beta': \mathcal{MCG}(S_{0,5}) \rightarrow \mathcal{MCG}(S_{1,2})$ be a quasi-inverse.

Now any quasi-isometry $f: \mathcal{MCG}(S_{1,2}) \rightarrow \mathcal{MCG}(S_{1,2})$ gives rise to a quasi-isometry $f' = \beta \circ f \circ \beta'$, and Theorem 1.1 applied to $S_{0,5}$ gives an element $g \in \mathcal{MCG}(S_{0,5})$ such that $d(f', L_g)$ is bounded. If g is in the index 5 image of $\mathcal{MCG}(S_{1,2})$ then a preimage $h \in \mathcal{MCG}(S_{1,2})$ works for h , i.e., $d(f, L_h)$ is bounded. If not, then at least we can produce the ‘‘almost-geometric’’ quasi-isometry $L = \beta' \circ L_g \circ \beta$, and obtain a bound on $d(f, L)$. This completes the proof.

10.4. Quasi-isometric rigidity

We conclude with the proof of Theorem 1.2, quasi-isometric rigidity of $\mathcal{MCG}(S)$. The argument here is well-known, cf. [51].

Let $G = \mathcal{MCG}(S)$. We may assume $\xi(S) \geq 2$ as the finite and virtually free cases are already known. Left-multiplication gives a homomorphism $\lambda: G \rightarrow \text{QI}(G)$, where $\text{QI}(G)$ is the group of quasi-isometries of G modulo the bounded-displacement subgroup. The kernel of λ is the center $Z = Z(G)$ (in general, $\ker \lambda$ consists of those elements whose centralizer has finite index in G . For $\mathcal{MCG}(S)$ it is easy to show that the center are the only such elements).

Now supposing $S \neq S_{1,2}$, Theorem 1.1 implies that λ is surjective. Hence we have $\text{QI}(G) = G/Z$ (this is Corollary 10.1).

Now if Γ is quasi-isometric to G then conjugation by the quasi-isometry Φ gives an isomorphism $\text{QI}(\Gamma) \cong \text{QI}(G)$ so we get a map $\lambda': \Gamma \rightarrow \text{QI}(G)$. Moreover $\ker \lambda'$ is finite: for each $\gamma \in \Gamma$, the quasi-isometry $\Phi L_\gamma \Phi^{-1}$ representing $\lambda'(\gamma)$ has uniformly bounded constants (depending on Φ), and hence by Theorem 1.1 is a *uniformly* bounded distance from its approximating element of G . Hence if $\gamma \in \ker \lambda'$, the approximating element is in Z , and so γ is restricted to a bounded set in Γ . Thus $\ker \lambda'$ is finite.

Finally, the image of λ' has finite index in $\text{QI}(G)$: this follows from the fact that the left-action of Γ on itself is transitive, and hence the conjugated action on G is cobounded. This gives the desired map $\Gamma \rightarrow G/Z$ with finite kernel and finite-index image.

If $G = \mathcal{MCG}(S_{1,2})$, we observe as in the proof of Theorem 1.1 that G/Z injects as a finite-index subgroup of $G' = \mathcal{MCG}(S_{0,5})$, and hence it inherits the rigidity property from G' with the additional cost of restricting to a finite-index subgroup of Γ .

References

- [1] Jorgen E. Andersen, *Mapping class groups do not have Kazhdan's property (T)*, Preprint, arXiv:0706.2184, 2007.
- [2] J. Behrstock, *Asymptotic geometry of the Mapping Class Group and Teichmüller space*, Ph.D. thesis, SUNY at Stony Brook, 2004.
- [3] ———, *Asymptotic geometry of the mapping class group and Teichmüller space*, *Geometry & Topology* **10** (2006), 1523–1578.
- [4] J. Behrstock, C. Druţu, and L. Mosher, *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, *Math. Ann.*, To appear, ARXIV:MATH.GT/0512592.
- [5] J. Behrstock, C. Druţu, and M. Sapir, *Median structures on asymptotic cones and homomorphisms into mapping class groups*, Preprint. ARXIV:0810.5376.
- [6] J. Behrstock and Y. Minsky, *Centroids and the rapid decay property in mapping class groups*, arXiv:0810.1969.
- [7] ———, *Dimension and rank for mapping class groups*, *Ann. of Math.* **167** (2008), no. 3, 1055–1077.
- [8] M. Bourdon and H. Pajot, *Rigidity of quasi-isometries for some hyperbolic buildings*, *Comment. Math. Helv.* **75** (2000), no. 4, 701–736.
- [9] J. Cannon and D. Cooper, *A characterization of cocompact hyperbolic and finite-volume hyperbolic groups in dimension three*, *Trans. AMS* **330** (1992), 419–431.
- [10] A. Casson and D. Jungreis, *Convergence groups and Seifert fibered 3-manifolds*, *Invent. Math.* **118** (1994), 441–456.

- [11] L. Van den Dries and A. Wilkie, *On Gromov's theorem concerning groups of polynomial growth and elementary logic*, J. Algebra **89** (1984), 349–374.
- [12] C. Druţu and M. Sapir, *Tree-graded spaces and asymptotic cones of groups*, Topology **44** (2005), 959–1058, With an appendix by Denis Osin and Mark Sapir.
- [13] T. Dymarz, *Large scale geometry of certain solvable groups*, Preprint. ARXIV:0712.2214.
- [14] A. Eskin, *Quasi-isometric rigidity of nonuniform lattices in higher rank symmetric spaces*, J. Amer. Math. Soc. **11** (1998), no. 2, 321–361.
- [15] Alex Eskin, David Fisher, and Kevin Whyte, *Quasi-isometries and rigidity of solvable groups*, Pure Appl. Math. Q. **3** (2007), no. 4, part 1, 927–947.
- [16] B. Farb, A. Lubotzky, and Y. Minsky, *Rank-1 phenomena for mapping class groups*, Duke Math. J. **106** (2001), no. 3, 581–597.
- [17] B. Farb and H. A. Masur, *Superrigidity and mapping class groups*, Topology **37** (1998), 1169–1176.
- [18] B. Farb and L. Mosher, *A rigidity theorem for the solvable Baumslag-Solitar groups*, Invent. Math. **131** (1998), no. 2, 419–451, Appendix by D. Cooper.
- [19] ———, *Quasi-isometric rigidity for the solvable Baumslag-Solitar groups, II*, Invent. Math. **137** (1999), no. 3, 613–649, ARXIV:MATH.GR/9809010.
- [20] B. Farb and R. Schwartz, *The large-scale geometry of Hilbert modular groups*, J. Differential Geom. **44** (1996), no. 3, 435–478.
- [21] D. Gabai, *Convergence groups are Fuchsian groups*, Ann. of Math. **136** (1992), 447–510.
- [22] M. Gromov, *Groups of polynomial growth and expanding maps*, IHES Sci. Publ. Math. **53** (1981), 53–73.
- [23] ———, *Infinite groups as geometric objects*, Proceedings of the International Congress of Mathematicians, Warsaw, Amer. Math. Soc., 1983, pp. 385–392.
- [24] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993, pp. 1–295.
- [25] M. Gromov and P. Pansu, *Rigidity of lattices: an introduction*, Geometric topology: recent developments, Lecture notes in Mathematics, vol. 1504, Springer, 1991, pp. 39–135.
- [26] U. Hamenstädt, *Geometry of the mapping class groups III: Geometric rank*, arXiv:math.GT/0512429.
- [27] J. Harer, *Stability of the homology of the mapping class group of an orientable surface*, Ann. of Math. **121** (1985), 215–249.
- [28] W. J. Harvey, *Boundary structure of the modular group*, Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference (I. Kra and B. Maskit, eds.), Ann. of Math. Stud. **97**, Princeton, 1981.
- [29] A. Hinkkanen, *The structure of certain quasisymmetric groups*, Mem. Amer. Math. Soc. **83** (1990), no. 422, iv+87.
- [30] N. V. Ivanov, *Algebraic properties of the Teichmüller modular group*, Dokl. Akad. Nauk SSSR **275** (1984), no. 4, 786–789.
- [31] N. V. Ivanov, *Complexes of curves and the Teichmüller modular group*, Uspekhi Mat. Nauk **42** (1987), 55–107.
- [32] N. V. Ivanov, *Teichmüller modular groups and arithmetic groups*, Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI) **167** (1988), no. Issled. Topol. 6, 95–110, 190–191.
- [33] N. V. Ivanov, *Automorphisms of complexes of curves and of Teichmüller spaces*, Internat. Math. Res. Notices (1997), no. 14, 651–666.
- [34] M. Kapovich, *Hyperbolic manifolds and discrete groups*, Progress in Mathematics, vol. 183, Birkhäuser Boston Inc., Boston, MA, 2001.
- [35] B. Kleiner and B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and euclidean buildings*, IHES Publ. Math. **86** (1997), 115–197.

- [36] M. Korkmaz, *Automorphisms of complexes of curves on punctured spheres and on punctured tori*, *Topology Appl.* **95** (1999), no. 2, 85–111.
- [37] F. Luo, *Automorphisms of the complex of curves*, *Topology* **39** (2000), no. 2, 283–298.
- [38] H. A. Masur and Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, *Invent. Math.* **138** (1999), 103–149.
- [39] ———, *Geometry of the complex of curves II: Hierarchical structure*, *Geom. Funct. Anal.* **10** (2000), 902–974.
- [40] J. McCarthy, *A "Tits-alternative" for subgroups of surface mapping class groups*, *Trans. Amer. Math. Soc.* **291** (1985), 583–612.
- [41] J. Milnor, *Growth of finitely generated solvable groups*, *J. Differential Geometry* **2** (1968), 447–449.
- [42] ———, *A note on curvature and fundamental group*, *J. Differential Geometry* **2** (1968), 1–7.
- [43] L. Mosher, *Mapping class groups are automatic*, *Ann. of Math.* **142** (1995), 303–384.
- [44] ———, *Fiber respecting quasi-isometries of surface group extensions*, Preprint: math.GR/0308067, 2003.
- [45] L. Mosher, M. Sageev, and K. Whyte, *Quasi-actions on trees II: Finite depth Bass-Serre trees*, To appear, arXiv:math.GR/0405237.
- [46] ———, *Quasi-actions on trees I: Bounded valence*, *Ann. of Math.* (2003), 116–154.
- [47] G. D. Mostow, *Strong rigidity of locally symmetric spaces*, *Ann. of Math. Stud.* no. 78, Princeton University Press, 1973.
- [48] P. Pansu, *Metriques de Carnot-Caratheodory et quasiisometries des espaces symmetriques de rang un*, *Ann. of Math.* **129** (1989), 1–60.
- [49] P. Papasoglu, *Quasi-isometry invariance of group splittings*, *Ann. of Math.* **161** (2005), no. 2, 759–830.
- [50] I. Peng, *Coarse differentiation and quasi-isometries of a class of solvable Lie groups I*, arXiv:0802.2596.
- [51] R. Schwartz, *The quasi-isometry classification of rank one lattices*, *Inst. Hautes Études Sci. Publ. Math.* (1995), no. 82, 133–168 (1996).
- [52] J. Stallings, *On torsion free groups with infinitely many ends*, *Ann. of Math.* **88** (1968), 312–334.
- [53] D. Sullivan, *On the ergodic theory at infinity of an arbitrary discrete group of hyperbolic motions*, *Riemann Surfaces and Related Topics: Proceedings of the 1978 Stony Brook Conference*, *Ann. of Math. Stud.* 97, Princeton, 1981.
- [54] A. S. Švarc, *A volume invariant of coverings*, *Dokl. Akad. Nauk SSSR (N.S.)* **105** (1955), 32–34.
- [55] Pekka Tukia, *Quasiconformal extension of quasisymmetric mappings compatible with a Möbius group*, *Acta Math.* **154** (1985), no. 3-4, 153–193.
- [56] K. Whyte, *The large scale geometry of the higher Baumslag-Solitar groups*, *Geom. Funct. Anal.* **11** (2001), no. 6, 1327–1343.
- [57] Xiangdong Xie, *Quasi-isometric rigidity of Fuchsian buildings*, *Topology* **45** (2006), no. 1, 101–169.

LEHMAN COLLEGE, CUNY

E-mail address: `jason@math.columbia.edu`

YALE UNIVERSITY

E-mail address: `bruce.kleiner@yale.edu`

YALE UNIVERSITY

E-mail address: `yair.minsky@yale.edu`

RUTGERS UNIVERSITY AT NEWARK

E-mail address: `mosher@andromeda.rutgers.edu`