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 virtually equal to it. (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 projection map $\mathcal{MCG}(S) \to \prod_{Y \subseteq S} \mathcal{C}(Y)$; and a construction of Σ-*hulls* in $\mathcal{MCG}(S)$, an analogue of convex hulls.

Contents

1. **Introduction**

In this article we will study the large-scale structure of the mapping class group $\mathcal{MCG}(S)$ of a finite-type oriented surface S. Our main tool is the action of $\mathcal{MCG}(S)$ on the complexes of curves of S and its subsurfaces. Our

Date: February 2, 2008.

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

first main result is that the group is *quasi-isometrically rigid* in the following sense.

Theorem 1.1. (Quasi-Isometric Rigidity) *If Γ is a finitely generated group quasi-isometric to $\mathcal{MCG}(S)$ then Γ is virtually $\mathcal{MCG}(S)$.*

That is, there exists a finite-index subgroup $\Gamma' < \Gamma$ and a homomorphism $\Gamma' \to \mathcal{MCG}(S)/Z(\mathcal{MCG}(S))$ with finite kernel and finite index image. $Z(\mathcal{MCG}(S))$ denotes the center of $\mathcal{MCG}(S)$.

This theorem was proved by Hamenstädt in [10]. The two proofs have a similar flavor in broad outline, although the underlying machinery supporting the outline is different.

Remarks on sporadic cases: Let $\xi(S) = \max\{0, 3g - 3 + k\}$, the number of curves in a pants decomposition of S, where g is the genus and k the number of punctures of S. In the cases $\xi \leq 1$, where S is a sphere with ≤ 4 punctures or a torus with ≤ 1 puncture, the group $\mathcal{MCG}(S)$ is virtually free or finite and has finite center, in which case the above theorem is already understood in different terms [23, 9, 6, 15], and is not covered by our methods.

When $\xi(S) \geq 2$, the center of $\mathcal{MCG}(S)$ is trivial, except when $S = S_{1,2}$ or $S_{2,0}$ in which case it is cyclic of order 2, generated by the hyperelliptic involution. We also have $\Gamma' = \Gamma$ except when $S = S_{1,2}$, in which case Γ' has index ≤ 5 in Γ.

Henceforth, the ambient surface S will always be an oriented, finite type surface with $\xi(S) \geq 2$. Nonetheless certain sporadic subsurfaces of S will play important roles, in particular annuli, once-punctured tori, and four-punctured spheres.

Theorem 1.1 is a consequence of our second main result, a classification of self-quasi-isometries of the group:

Theorem 1.2. (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 $f : \mathcal{MCG}(S) \to \mathcal{MCG}(S)$ 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 $\mathcal{MCG}(S_{1,2})/Z(S_{1,2}) \hookrightarrow \mathcal{MCG}(S_{0,5})$.

As a corollary to this theorem we obtain precise information about the natural homomorphism $\mathcal{MCG}(S)/Z(\mathcal{MCG}(S)) \to \mathrm{QI}(\mathcal{MCG}(S))$, the latter being the *quasi-isometry group* of $\mathcal{MCG}(S)$, obtained by identifying two quasi-isometries of $\mathcal{MCG}(S)$ when they differ by a bounded amount; the group operation is induced by composition, and the homomorphism is induced by left multiplication.

Corollary 1.3. *If $\xi(S) \geq 2$ then the homomorphism $\mathcal{MCG}(S)/Z(\mathcal{MCG}(S)) \to \mathrm{QI}(\mathcal{MCG}(S))$ is an isomorphism except when $S = S_{1,2}$, in which case it is an isomorphism to a subgroup of index 5.*

The method of studying the large scale geometry of groups via quasi-isometries, in particular classifying self quasi-isometries of a group and classifying groups quasi-isometric to a given group, originated with Gromov [8]. The particular form of quasi-isometry classification occurring in Corollary 1.3 first occurs for a different class of groups in Schwartz [22], and it is from this paper that we take the argument that reduces Theorem 1.1 to Theorem 1.2.

We note that this reduction is not used in Hamenstädt's proof of Theorem 1.1, and in particular that Theorem 1.2 and Corollary 1.3 are new here.

We will briefly sketch the proof of Theorem 1.2, before proceeding to definitions and preliminaries in Section 2.

Coarse preservation of Dehn twist flats. In order to control a quasi-isometry $f\colon G \to G$ 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 sets defined in terms of *markings* of S.

Theorem 10.2 states that a quasi-isometry $f\colon \mathcal{MCG}(S) \to \mathcal{MCG}(S)$ *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 this theorem is established, we can apply known results. The coarse permutation of flats induces an automorphism of the pants complex of S which by a theorem of Margalit [19] is induced by some mapping class $g \in \mathcal{MCG}(S)$ (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 g is uniformly close to f. By working slightly harder with the permutation of flats we can bypass Margalit's Theorem by showing directly that f induces an automorphism of the curve complex of S and then using the theorem of Ivanov–Korkmaz–Luo that each automorphism of the curve complex is induced by some mapping class [13, 17, 18].

Preservation of asymptotic Dehn twist flats. Theorem 10.2 will be proven, following Kleiner-Leeb [16], by reduction to the *asymptotic cone* of $\mathcal{MCG}(S)$. The asymptotic cone, which we denote $\mathcal{M}_\omega$, 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 §2.2 for details.

A quasi-isometry of $\mathcal{MCG}(S)$ converges after rescaling to a bilipschitz homeomorphsm of $\mathcal{M}_\omega$, 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. This is the statement of Theorem 10.1, whose proof will take up most of the paper.

The reduction of Theorem 10.2 to Theorem 10.1 will be discussed in detail in Section 10.2. Let us now sketch the proof of Theorem 10.1.

Structure of $\mathcal{MCG}(S)$ via hyperbolicity. We begin in Sections 3–5 by refining the tools developed in Masur–Minsky [20, 21] and Behrstock [1] to study the coarse structure of $\mathcal{MCG}(S)$ using properties of curve complexes.

A recurring theme is that of subsurface projections $\pi_W\colon \mathcal{MCG}(S) \to \mathcal{C}(W)$, where W is an essential subsurface of S, and $\mathcal{C}(W)$ is the curve complex of W (or the twist complex when W is an annulus). Combining these we obtain a map

$$\Pi\colon \mathcal{MCG}(S) \to \prod_W \mathcal{C}(W),$$

and the Quasidistance Formula of [21] (see Theorem 2.7 below) shows that this is, in a limited sense, like a quasi-isometric embedding to the ℓ^1 metric on the product.

Theorem 4.3, which may be of independent interest, characterizes the image of Π, up to bounded error, in terms of natural consistency inequalities. This theorem makes it possible to state and analyze many of our constructions simply in terms of what they look like in the projection image, where the hyperbolicity of the $\mathcal{C}(W)$'s [20] can be exploited. For instance, although many of the technical constructions in this paper were originally conceived using *hierarchy paths* in the sense of Masur–Minsky [21], these paths are difficult to define and to work with. Hence, for the most part we have avoided using them as much as possible and Theorem 4.3 is one of the main tools that allows us to do this.

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 map. 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 product structure comes from the different components of the complement of μ. 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 they are quasi-isometrically embedded subsets of $\mathcal{MCG}(S)$, generalizing a result of [21]. In Lemmas 3.2 and 3.3 we analyze the sets, called *junctures*, along which two of these regions are close — junctures are generalizations of the coarse intersections of quasiconvex sets in a hyperbolic space.

In Section 5, we define and study Σ-*hulls*, which function like convex hulls of finite sets in $\mathcal{MCG}(S)$. The Consistency Theorem 4.3 makes it easy to construct these hulls in terms of the hyperbolic hulls in the projections to subsurface complexes.

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 Lemmas 5.4 and 5.5 imply, in the cone, that ultralimits of Σ-hulls are contractible and have controlled geometry (Lemmas 6.1 and 6.2), which allows us to use them to build singular chains which are controlled in a sense that we describe with the term Σ-*compatible*. With this, and the results of Behrstock-Minsky [3], we prove local homological dimension bounds, a result originally established by Hamenstädt [10]. We also obtain Corollary 6.8, an analogue of a result of Kleiner-Leeb [16], 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 [1] and [3] to prove some separation properties of the asymptotic cone. In [3], it was shown that any pair of distinct points in the asymptotic cone of $\mathcal{MCG}(S)$ could be separated by an ultralimit of sets of the form $\mathcal{Q}(\mu)$, and this enabled an inductive argument to compute the compact topological dimension of the cone. Theorem 7.3 refines our understanding of the components of the complement of such a set, introducing the notion of a *microscopic jet*, which is an ultralimit of geodesic segments in subsurfaces of S, such that asymptotic behavior of projections to these segments determines the division of the complement into connected components.

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 $\pi_n \colon X \to g_n$ 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.

Theorem 7.3 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 example, one of our points x_ω or y_ω may equal 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 7.8. 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 Drutu-Sapir [5]. 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 8.4 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 8.8 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 8.4 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 6.8 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 $\boldsymbol{x}$. The starting point, using Theorem 8.8, is the fact that the germ of any top-dimensional manifold at $\boldsymbol{x}$ 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 9.6) 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. The global statement follows directly, and this gives the proof of Theorem 10.1.

2. Preliminaries

2.1. Curves, markings and projections

We review the foundations of curve complexes and marking complexes, in particular the quasidistance formula for the marking complex, expressed in terms of projections to curve complexes of essential subsurfaces. The main references are [21] and [3].

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. We denote $X = X_{g,n}$ where g is the genus and n the number of punctures. Let $\xi(X) = \max\{3g - 3 + n, 0\}$.

Throughout the paper we will consider a single "ambient" surface S of $\xi \geq 2$. We will also consider subsurfaces of S for which $\xi \geq 1$, as well as subannuli of S. That is, we will automatically exclude spheres with 0, 1 or 3 punctures and closed tori.

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

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

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

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 $\xi(X) \geq 1$, a curve system C on X is called a *pants decomposition* if it is pairwise nonisotopic and maximal with respect to this property — equivalently, each component of $X - C$ is a three-punctured sphere, a *pair of pants*. The number of curves in a pants decomposition is $\xi(X)$.

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' *essentially intersect* if, after putting them in efficient position, the intersection is nonempty.

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 $C \cap \overline{Y}$ 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 clases in S of essential simple closed curves contained 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.

On the set of essential subsurfaces define a relation $X \Subset Y$ to mean $\Gamma(X) \subset \Gamma(Y)$.

Lemma 2.1. *The relation $\Subset$ 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 $\Cap$, and $\Cup$, called* essential union *and* essential intersection, *which have the following properties:*

1. *$X \Cup Y$ is the unique $\Subset$-minimal essential subsurface Z such that $X \Subset Z$ and $Y \Subset Z$.*
2. *$X \Cap Y$ is the unique $\Subset$-maximal essential subsurface Z such that $Z \Subset X$ and $Z \Subset Y$.*

In other words we have a lattice whose partial order is $\Subset$ and whose meet and join operations are $\Cap$ and $\Cup$, 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 $\mathit{Fill}(C)$. One easily shows $X = \mathit{Fill}(\Gamma(X))$ provided $\Gamma(X) \neq \emptyset$, that is, if X is not a pair of pants. It follows immediately that $\Subset$ is a partial order.

Let us now show that $\Cup$ is defined. Given X and Y, let $Z = \mathit{Fill}(C)$ where $C = \Gamma(X) \cup \Gamma(Y)$. Any curve in C either essentially intersects 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 $X \Subset Z$ and $Y \Subset Z$. 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 $Z \Subset Z'$. Uniqueness follows from the fact that $\Subset$ is a partial order. We therefore set $X \Cup Y = Z$.

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

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 $\Gamma(X \Cap Y) = \Gamma(X) \cap \Gamma(Y)$. This is because any element γ in $\Gamma(X) \cap \Gamma(Y)$ is the core of an essential annulus A in both, hence $A \Subset X \Cap Y$ by (2), so $\gamma \in \Gamma(X \Cap Y)$. The other direction follows from the fact that $X \Cap Y$ 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 [3]). 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 $X \uplus X^c$ is usually not S; for example, if X is a regular neighborhood of a pants decomposition on S then $X^c = \emptyset$.

Curve complex. We associate a simplicial complex $\mathcal{C}(Y)$ with each connected surface Y with $\xi(Y) \geq 1$, as well as for annular subsurfaces 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 [11], Ivanov [12] and Masur-Minsky [20].

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 [21], 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 (isotopy class of) pair of 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 isotopy class of 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 $\mathrm{int}(\mathcal{A}_Y) \to S$ to a homeomorphism $\widetilde{\tau}_\gamma \colon \mathcal{A}_Y \to \mathcal{A}_Y$ 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 [21].

Lemma 2.2. *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 [20]:

Theorem 2.3. *For each surface S with $\xi(S) \geq 1$, the curve complex $\mathcal{C}(S)$ is an infinite diameter δ-hyperbolic metric space, with respect to the simplicial metric.*

Markings and partial markings. We define markings and the marking complex for any surface S with $\xi \geq 1$, and also for any essential subannulus of such a surface. We also introduce partial markings. (In [21], partial markings are called markings, and markings are called complete markings).

A *partial marking* $\mu = (\mathrm{base}(\mu), t)$ on S consists of a simplex $\mathrm{base}(\mu)$ in $\mathcal{C}(S)$ 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 $b \in \mathrm{base}(\mu)$; by convention we allow the empty set $\emptyset$ as a partial marking of S. 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 $\mathrm{base}(\mu)$ is a pants decomposition and every $b \in \mathrm{base}(\mu)$ is marked. Given two partial markings $\mu = (\mathrm{base}(\mu), t)$, $\mu' = (\mathrm{base}(\mu'), t')$ we write $\mu \subset \mu'$ to mean that $\mathrm{base}(\mu) \subset \mathrm{base}(\mu')$ and, for each $b \in \mathrm{base}(\mu)$, b is marked by μ only if it is marked by μ' in which case $t(b) = t'(b)$.

Given an essential subannulus F of S, a *marking of F* is just a vertex of $\mathcal{C}(F)$, and a *partial marking of F* is either a marking of F or $\emptyset$.

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

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

The vertices of $\mathcal{M}(S)$ are just the markings of S. To define the edges we first need this notion: If b and c are intersecting 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. This set has diameter bound in $\mathcal{C}(b)$ (see below for π_b in a more general setting).

Edges in $\mathcal{M}(S)$ correspond to *elementary moves* among markings on S, which come in two flavors: *twist moves* and *flip moves*. To define them, consider a marking μ on S and a curve $b \in \mathrm{base}(\mu)$.

A marking μ' is said to be obtained from μ by a *twist move* about b if $\mathrm{base}(\mu) = \mathrm{base}(\mu')$, μ, μ' have the same transversals to each curve in $\mathrm{base}(\mu) \setminus \{b\} = \mathrm{base}(\mu') \setminus \{b\}$, and the transversals $t(b)$ in μ and $t'(b)$ in μ' satisfy $d_{\mathcal{C}(b)}(t, t') \leq 2$.

A marking $\mu' = (\mathrm{base}(\mu'), t')$ is said to be obtained from μ by a *flip move* along b if there exists $b' \in \mathrm{base}(\mu')$ such that $\mathrm{base}(\mu) \setminus \{b\} = \mathrm{base}(\mu') \setminus \{b'\}$, $Fill(b, b')$ 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 $d_{\mathcal{C}(b')}(t'(b'), \pi_{b'}(b)) \leq 2$.

This complex is locally infinite because of the structure of transversals, but it is still quasi-isometric to the locally finite complex of *clean markings* defined in [21]: A *clean marking* is a marking $\mu = (\mathrm{base}(\mu), t)$ with the following properties: for each $b \in \mathrm{base}(\mu)$, $t(b)$ is $\pi_b(c)$ where $c = c(b)$ is an essential curve in the component F of $S \setminus (\mathrm{base}(\mu) \setminus \{b\})$ containing b, and the curves b and c have minimal intersection number in F. The complex of clean markings is a connected, $\mathcal{MCG}(S)$-invariant subcomplex of $\mathcal{M}(S)$ 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., [1] and [3]. Because the full complex is more convenient for our purposes, we record this quasi-isometry:

Proposition 2.4. *The marking complex $\mathcal{M}(S)$ is quasi-isometric to $\mathcal{MCG}(S)$ and to the subcomplex of clean markings. More precisely, for each $\mu_0 \in \mathcal{M}(S)$ the orbit map $\phi \mapsto \phi(\mu_0)$ is a quasi-isometry from $\mathcal{MCG}(S)$ to $\mathcal{M}(S)$.*

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

Proof. (Sketch.) As noted in [21], the clean marking complex is locally finite, the action of $\mathcal{MCG}(S)$ is properly discontinuous and cocompact, and so by Svarc-Milnor the orbit map is a quasi-isometry between $\mathcal{MCG}(S)$ and the complex of clean markings. The inclusion of the complex of clean markings into $\mathcal{M}(S)$ is an $\mathcal{MCG}(S)$-equivariant quasi-isometry, because for each marking there is a clean marking within a uniformly bounded distance C by Lemma 2.2 (and [21]), and for each edge of $\mathcal{M}(S)$ 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 [21] of the edges allowed between clean markings). $\square$

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

First, given an essential curve $\gamma \subset S$ and an essential subsurface $Y \subset S$, we define $\gamma \pitchfork Y$ and $Y \pitchfork \gamma$ to mean 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 an 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 $\overline{Y}$*, each a properly embedded arc $\alpha \subset \overline{Y}$ 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 $\gamma \pitchfork Y$.

Given two essential subsurfaces $X, Y \subset S$, define $X \pitchfork Y$ to mean that $\partial Y \pitchfork Y'$ and $\partial Y' \pitchfork Y$. Equivalently, after Y, Y' are isotoped so that ∂Y, $\partial Y'$ intersect efficiently, some component of $\overline{Y} \cap \partial Y'$ is an essential curve or arc in Y and some component of $\partial Y \cap \overline{Y}'$ is an essential curve or arc in Y'. It is also equivalent to say that neither of Y or Y' is isotopic into the other and no matter how Y, Y' are chosen in their isotopy classes their intersection is nonempty.

If Y is an essential subsurface and μ is a partial marking, we define $\mu \pitchfork Y$ and $Y \pitchfork \mu$ to mean that either $\text{base}(\mu) \pitchfork Y$ or (a component of) Y is an annulus neighborhood of some *marked* $b \in \text{base}(\mu)$.

Given a partial marking $\mu = (\text{base}(\mu), t)$ on S, its *open subsurface*, $\text{open}(\mu) = \text{open}_S(\mu)$, is defined to be the essential union of all subsurfaces Z such that $Z \not\pitchfork \mu$. (We usually drop the subscript unless we want to emphasize the surface in which the operation takes place, as in the proof of Lemma 3.2.) 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 $\text{base}(\mu)$.

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 $\text{open}(\mu)$. We note two properties of $\text{supp}(\mu)$: each component of $\partial\,\text{supp}(\mu)$ is isotopic to a component of $\text{base}(\mu)$; and μ restricts to a (full) marking of each component of $\text{supp}(\mu)$. Moreover $\text{supp}(\mu)$ is characterized up to isotopy as the maximal essential subsurface of S with respect to these properties. 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 $\mu \neq \emptyset$.

Subsurface projections. Following [21], [1], and [3], 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. 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 set of choices has uniformly bounded diameter, and in practice we will be able to abuse terminology and treat the maps as if their images are single points.

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 section in the marking complex discussion above, where b is the core of Y.

Note in both cases that $\pi_{\mathcal{C}(Y)}(\gamma)$ is nonempty if and only if $\gamma \pitchfork Y$. In [21] it is shown that $\operatorname{diam}\pi_{\mathcal{C}(Y)}(\gamma)$ is bounded by a constant depending only on the topology of S — that is, $\pi_{\mathcal{C}(Y)}$ is "coarsely well-defined".

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 [21] will be important for us:

Theorem 2.5. *Let Y be a connected essential subsurface of S satisfying $\xi(Y) \neq 3$ and let g be a geodesic segment in $\mathcal{C}(S)$ for which $v \pitchfork Y$ for every vertex v of g. Then*

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

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 $\operatorname{base}(\mu) \pitchfork Y$ we let $\pi_{\mathcal{C}(Y)}(\mu)$ be the union of $\pi_{\mathcal{C}(Y)}(b)$ over all $b \in \operatorname{base}(\mu)$. When Y is an annulus neighborhood of a marked $b \in \operatorname{base}(\mu)$ 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. When μ is a marking of S, $\pi_{\mathcal{M}(Y)}(\mu)$ will be a marking of Y, so that we will obtain a coarse Lipschitz map $\pi_{\mathcal{M}(Y)} \colon \mathcal{M}(S) \to \mathcal{M}(Y)$.

Write $\mu = (\operatorname{base}(\mu), t)$. 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.

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 denote the subsurface isotopic to the essential complement of A in Y (discarding 3-holed spheres), disjoint union with A. 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 that overlap μ.

There are choices at each stage of this construction, but the final output is coarsely well-defined, as proved in [1]. See Lemma 2.9 for a statement.

For a partial marking μ, recall that $\mathrm{open}(\mu)$ is the unique maximal essential subsurface Y such that $\mu \not\pitchfork Y$, 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 2.6. *If Y is an essential subsurface of S and μ a partial marking in S, then*

$$\mathrm{open}_Y(\pi_{\mathcal{M}(Y)}(\mu)) = \bigcup\!\!\!\!\!\!\circ\;\; \{Z \Subset Y : Z \not\pitchfork \mu\}.$$

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 $\mathrm{open}_Y(\mu')$ does not overlap μ.

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

We have shown that $\mathrm{open}_Y(\mu')$ 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 $\mathrm{open}_Y(\mu')$. Hence the two sides are equal. $\square$

Notation: 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

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

When the context is clear we occasionally abbreviate $d_{\mathcal{C}(Y)}$ to d_Y.

Extending markings by markings. Given two partial markings μ, ν on S, define the *extension of μ by ν* to be the partial marking given by:

$$\mu \rfloor \nu = \mu \cup \pi_{\mathcal{M}(\mathrm{open}(\mu))}(\nu).$$

This will be used prominently in Section 3.1.

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 $d \geq 0$, $A \geq 0$ let $\{\!\{d\}\!\}_A$ equal d if $d \geq A$ and 0 otherwise.

Given $r, s \geq 0$, $K \geq 1$, $C \geq 0$ we write $r \overset{K,C}{\approx} s$ to mean that $\frac{1}{K}s - C \leq r \leq Ks + C$. We also write $r \approx s$ to mean that $r \overset{K,C}{\approx} s$ 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 [21].

Theorem 2.7. *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 by*

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

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. In general any subsurface Y for which we consider $\mathcal{C}(Y)$ should be assumed connected, whereas if $\mathcal{M}(Y)$ is being considered Y can be disconnected, in which case $\mathcal{M}(Y)$ is a product over the components.

There are two ways that Theorem 2.7 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 2.8. *For any r there exists t such that, if for $\mu, \nu \in \mathcal{M}(S)$ we have $d_{\mathcal{C}(W)}(\mu, \nu) \leq r$ 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 [21] yields $t = O(r^{\xi(S)})$.

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

Lemma 2.9. *Subsurface projections are* coarsely 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.

Here "uniformly bounded" means bounded by a constant depending only on $\xi(S)$, and in fact for the $\mathcal{C}(Y)$ bounds in this lemma and the next a value of 3 will work (even 2 if $\xi(S) > 1$) – see [21] for details.

Lemma 2.10. *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 $\mathrm{diam}(\pi_{\mathcal{C}(Y)}(x) \cup \pi_{\mathcal{C}(Y)}(y))$ *is uniformly bounded.*

Similarly $d_{\mathcal{M}(Y)}(\mu, \nu)$ is uniformly bounded for any μ, ν in $\mathcal{M}(S)$ with $d(\mu, \nu) = 1$.

Lemma 2.11. *Subsurface projections for nested subsurfaces are* coarsely composable*:*

Let $X, Y \subset S$ be essential subsurfaces such that $X \Subset Y$. Any $\gamma \in \mathcal{C}_0(S)$ 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 $\pi_X \circ \pi_Y$, or in other words

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

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.

2.2. **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 in this paper. (See [4, 7] 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.

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, and Bolzano-Weierstrass in this lanugage means that when X is compact every sequence has an ultralimit. When two sequences coincide on a set of indices whose ω-measure is equal to 1, they have the same ultralimit, and for this reason we define the *ultraproduct* of a sequence of spaces X_n to be the quotient $\Pi_n X_n / \sim$ of the cartesian product identifying (x_n) with (y_n) if they agree for ω-a.e. n. We will often use the notation $\overline{X}$ for the ultraproduct and $\overline{x} = \langle x_n \rangle$ for the equivalence class of a sequence (x_n).

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

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

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

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

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

For the image of $\overline{y}$ in the asymptotic cone we use the notation either y_ω or $\boldsymbol{y}$.

The rescaling limit works equally well for a sequence $(X_n, x_n, \mathrm{dist}_n)$, and so we call $\lim_\omega (X_n, x_n, \mathrm{dist}_n / s_n)$ the asymptotic cone of the sequence.

Convention: For the remainder of the paper, we fix a non-principal ultrafilter ω, a sequence of scaling constants $s_n \to \infty$, and a basepoint μ_0 for $\mathcal{M}(S)$. 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 will do, and in the last section we will need the flexibility of varying the choice of s_n). Further, any mention of linear/sub-linear growth of a non-negative function $f(n)$ is with respect to this choice of ω and s_n, i.e., we say $f(n)$ has linear growth if $0 < \lim_\omega \frac{f(n)}{s_n} < \infty$ and sublinear if $\lim_\omega \frac{f(n)}{s_n} = 0$.

Note that since $\mathcal{M}$ is quasi-isometric to $\mathcal{MCG}$ with (any) word metric, which is homogeneous, 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 2.10, and we can use this to define its asymptotic cone $\mathcal{M}_\omega(W)$. For a disconnected subsurface $W = \sqcup_{i=1}^k W_i$ we have $\mathcal{M}(W) = \Pi_{i=1}^k \mathcal{M}(W_i)$ and we may similarly construct $\mathcal{M}_\omega(W)$ which can be identified with $\Pi_{i=1}^k \mathcal{M}_\omega(W_i)$ (this follows from the general fact that the process of taking asymptotic cones commutes with finite products). Note that for an annulus A we've defined $\mathcal{M}(A) = \mathcal{C}(A)$ which is quasi-isometric to $\mathbb{Z}$, so $\mathcal{M}_\omega(A)$ is $\mathbb{R}$.

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 $(S, S, \ldots)$ so that $\mathcal{M}(\overline{S})$ is the ultraproduct of $(\mathcal{M}(S), \ldots)$ 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 $\omega(\{n : a_n = a\}) = 1$ – so $\overline{A}$ is naturally identified with A. Hence, 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, and moreover this is the same for any equivalent sequence, 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.

For two sequences of sets (A_n) and (B_n), a sequence of functions $f_n \colon A_n \to B_n$ gives rise to a single function $\overline{f} \colon \overline{A} \to \overline{B}$, and $\overline{f}$ determines f_n up to equivalence in the ultraproduct of the sequence $(B_n^{A_n})$. Hence e.g. we can think of a sequence of projection maps $\pi_{\mathcal{M}(W_n)} \colon \mathcal{M}(S) \to \mathcal{M}(W_n)$ as a single map

$$\pi_{\mathcal{M}(\overline{W})} \colon \mathcal{M}(\overline{S}) \to \mathcal{M}(\overline{W})$$

which descends to a map of the asymptotic cones,

$$\pi_{\mathcal{M}_\omega(\overline{W})} \colon \mathcal{M}_\omega(S) \to \mathcal{M}_\omega(\overline{W}).$$

Note by Lemma 2.10 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 will be 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 *junctures*, the part of each set which comes closest to the other. Later, when we pass to the asymptotic cone, these junctures 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

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

the set of all (full) markings that extend μ. For instance, when μ is a curve system $\mathcal{Q}(\mu)$ is the collection of all markings whose base curves contain μ: such a set is quasi-isometric to a coset of a subgroup stabilizing a certain curve system. In particular, if μ is a pants decomposition, then $\mathcal{Q}(\mu)$ is quasi-isometric to a coset of a maximal-rank free-abelian subgroup generated by Dehn twists; accordingly, in this case, we refer to $\mathcal{Q}(\mu)$ as a *Dehn twist flat.*

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 $\mu' \bigm| Y$ (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 a second partial marking ν on S, recall from Section 2.1 the extension

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

If $\pi_{\mathcal{M}(\text{open}(\mu))}(\nu)$ is a full marking in $\text{open}(\mu)$ (for example if ν itself is a full marking), then $\mu \rfloor \nu \in \mathcal{Q}(\mu)$.

The following result generalizes the case considered in [3] 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*

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

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*

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

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

(3) *Moreover, again with uniform constants, for each $x \in \mathcal{M}(S)$*
$$d_{\mathcal{M}(S)}(x, \mathcal{Q}(\mu)) \approx d_{\mathcal{M}(S)}(x, \mu\rfloor x).$$

As a consequence of (1), combined with the ordinary quasidistance formula for components of $S \setminus \operatorname{base}(\mu)$, if we let Δ be the subset of $\operatorname{base}(\mu)$ consisting of those curves for which no transversal is defined then we have a quasi-isometry
$$\mathcal{Q}(\mu) \approx \mathbb{Z}^\Delta \times \prod_{\substack{Y \in |\operatorname{open}(\mu)| \\ \xi(Y) \geq 1}} \mathcal{M}(Y)$$

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 $\mathcal{Q}(\mu) = \{\mu\}$, $\operatorname{open}(\mu) = \emptyset$, $\operatorname{supp}(\mu) = S$, and (2) is just the ordinary quasidistance formula to μ.

Proof of Proposition 3.1. Recall that if F is an essential subsurface of S that is *not* essentially contained in $\operatorname{open}(\mu)$ then $\mu \pitchfork F$. 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 $\pi_{\mathcal{C}(F)}(\mu)$ (because of coarse well-definedness of $\pi_{\mathcal{C}(F)}$, Lemma 2.9), and so we can make the $\pi_{\mathcal{C}(F)}$ term drop out of the quasidistance formula by raising the threshold. We obtain (1) as an immediate consequence.

Now consider an essential (connected) subsurface $Y \subset S$ such that $\mu \pitchfork Y$. For each $\mu' \in \mathcal{Q}(\mu)$, from coarse well-definedness (Lemma 2.9) 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 $d_{\mathcal{C}(Y)}(x, \mu) \overset{C}{\approx} d_{\mathcal{C}(Y)}(x, \mu')$ where C depends only on the topology of S. Since 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 $d_{\mathcal{M}(S)}(x, \mu')$, 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 2.11 $\pi_{\mathcal{C}(Y)}(\mu')$ is within uniformly bounded distance of $\pi_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(\operatorname{open}(\mu))}(\mu'))$, which equals $\pi_{\mathcal{C}(Y)}(\pi_{\mathcal{C}(\operatorname{open}(\mu))}(x))$, 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 for the family $\{\mathcal{Q}(\mu)\}$.

Lemma 3.2. *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*

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

More precisely we have:

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

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

where we denote $\mathrm{supp}(\mu_0,\mu_1) = \mathrm{open}(\mu_0,\mu_1)^c$.
(3) *For $x_i \in \mathcal{Q}(\mu_i)$ we have*

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

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

$$\mathrm{open}(\mu \rfloor \nu) \stackrel{i}{\equiv} \mathrm{open}(\mu) \Cap \mathrm{open}(\nu) \tag{3.1}$$

for any two partial markings. $Z = \mathrm{open}(\mu \rfloor \nu)$ is the maximal essential subsurface that does not overlap $\mu \rfloor \nu$, hence Z does not overlap μ so $Z \Subset \mathrm{open}(\mu)$, and Z does not overlap $\pi_{\mathcal{M}(\mathrm{open}(\mu))}(\nu)$, so by Lemma 2.6, Z does not overlap ν. So $Z \Subset \mathrm{open}(\nu)$, and we conclude $\mathrm{open}(\mu \rfloor \nu) \Subset \mathrm{open}(\mu) \Cap \mathrm{open}(\nu)$.

Conversely letting $X = \operatorname{open}(\mu) \Cap \operatorname{open}(\nu)$, Lemma 2.6 implies that $X \subset \operatorname{open}_{\operatorname{open}(\mu)}(\pi_{\mathcal{M}(\operatorname{open}(\mu))}(\nu))$ so X does not overlap $\mu\rfloor\nu$. We conclude that $\operatorname{open}(\mu) \Cap \operatorname{open}(\nu) \Subset \operatorname{open}(\mu\rfloor\nu)$, 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 $\operatorname{supp}(\mu_0,\mu_1)$, the distance between these markings in the complex $\mathcal{M}(\operatorname{supp}(\mu_0,\mu_1))$ is defined, so that (2) makes sense.

To obtain the inequality $\gtrsim$ in (2), consider any term in the quasidistance formula for $d_{\mathcal{M}(\operatorname{supp}(\mu_0,\mu_1))}(\mu_0\rfloor\mu_1, \mu_1\rfloor\mu_0)$, indexed by $Y \subset \operatorname{supp}(\mu_0,\mu_1)$. 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 $x \mid \operatorname{open}(\mu,\nu)$ 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 $\partial\operatorname{open}(\mu,\nu)$ then $\pi_{\mathcal{C}(Y)}(x)$ and $\pi_{\mathcal{C}(Y)}(x')$ are uniformly close to $\pi_{\mathcal{C}(Y)}(\partial\operatorname{open}(\mu,\nu))$. If Y does not essentially intersect $\partial\operatorname{open}(\mu,\nu)$ then Y is isotopic into $\operatorname{open}(\mu,\nu)$ or its complement. If Y is isotopic into the complement of $\operatorname{open}(\mu,\nu)$ then either Y is an annulus component of $\operatorname{open}(\mu,\nu)$ or Y is an index for the right hand side of (2). If Y is an annulus component of $\operatorname{open}(\mu,\nu)$, or if Y is an essential subsurface of a component of $\operatorname{open}(\mu,\nu)$, then $\pi_{\mathcal{C}(Y)}(x)$ and $\pi_{\mathcal{C}(Y)}(x')$ are within uniformly bounded distance of the projection of $x \mid \operatorname{open}(\mu,\nu) = x' \mid \operatorname{open}(\mu,\nu)$. 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 $Y \Subset \operatorname{open}(\mu_0)$ and $Y \pitchfork \mu_1$ by Lemma 2.6; and the converse holds as well. Therefore by raising the threshold A we get

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

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 2.11; it follows that by raising the threshold above twice this bound, at the cost of another multiplicative factor of 2, we get

$$d_{\mathcal{M}(S)}(x_0, E_{01}) \approx \sum_{\substack{Y \Subset \mathrm{open}(\mu_0) \\ Y \pitchfork \mu_1}} \{\!\!\{ d_{\mathcal{C}(Y)}(x_0, \mu_1) \}\!\!\}_A \tag{3.2}$$

We obtain a similar expression for the second summand:

$$d_{\mathcal{M}(S)}(x_1, E_{10}) \approx \sum_{\substack{Y \Subset \mathrm{open}(\mu_1) \\ Y \pitchfork \mu_0}} \{\!\!\{ d_{\mathcal{C}(Y)}(x_1, \mu_0) \}\!\!\}_A \tag{3.3}$$

The third summand is given by

$$d_{\mathcal{M}(\mathrm{supp}(\mu_0,\mu_1))}(\mu_0\rfloor\mu_1, \mu_1\rfloor\mu_0) \approx \sum_{Y \Subset \mathrm{supp}(\mu_0,\mu_1)} \{\!\!\{ d_{\mathcal{C}(Y)}(\mu_0\rfloor\mu_1, \mu_1\rfloor\mu_0) \}\!\!\}_A$$

If $Y \Subset \mathrm{supp}(\mu_0, \mu_1)$ 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 $Y \Subset \mathrm{supp}(\mu_0, \mu_1)$ and $Y \not\pitchfork \mu_1$ then $Y \pitchfork \mu_0$ and these terms may be dropped.

If $Y \not\Subset \mathrm{supp}(\mu_0, \mu_1)$, and if $Y \pitchfork \mu_0$ and $Y \pitchfork \mu_1$, then $Y \pitchfork \partial\,\mathrm{supp}(\mu_0, \mu_1)$; for if not, Y would be isotopic to the complement of $\mathrm{supp}(\mu_0, \mu_1)$ and so to overlap μ_0 and μ_1, Y would have to be an annulus isotopic to a boundary curve of $\mathrm{supp}(\mu_0, \mu_1)$, which is marked by both μ_1 and μ_0. But in this case, by definition the annulus would be a component of $\mathrm{supp}(\mu_0, \mu_1)$ so $Y \Subset \mathrm{supp}(\mu_0, \mu_1)$ 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 $\pi_{\mathcal{C}(Y)}(\partial\,\mathrm{supp}(\mu_0, \mu_1))$, 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

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

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

$$d_{\mathcal{M}(\mathrm{open}(\mu_0,\mu_1))}(x_0, x_1) \approx \sum_{Y \Subset \mathrm{open}(\mu_0,\mu_1)} \{\!\!\{ d_{\mathcal{C}(Y)}(x_0, x_1) \}\!\!\}_A \tag{3.5}$$

Now putting these four sums (3.2), (3.3), (3.4), (3.5) together, and recalling that $Y \Subset \mathrm{open}(\mu_i)$ 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 $d_{\mathcal{M}(S)}(x_0, x_1)$. □

3.2. Cubes and their junctures

Consider a subset of $\mathcal{M}(S)$ formed as follows. Choose a partial marking μ such that the components $W_1, \ldots, W_m$ of $W = \mathrm{open}(\mu)$ satisfy $\xi(W_i) \leq 1$, so each W_i is an annulus, one-holed torus, or four-holed sphere. The marking complex $\mathcal{M}(W_i)$ is quasi-isometric to a tree which we denote $\mathcal{TM}(W_i)$: when W_i is an annulus then $\mathcal{TM}(W_i)$ is isometric to $\mathbf{R}$; and in the other two cases $\mathcal{TM}(W_i)$ is isometric to the dual tree of the usual modular diagram for $\mathrm{SL}_2\mathbb{Z}$. 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 $r = \{r_1, \ldots, r_m\}$. 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. It follows that 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 3.2. Here are the details.

Lemma 3.3. *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 $\mathrm{open}(\mu, \nu) = W \Cap V$, whose marking complex parameterizes the junctures of $\mathcal{Q}(\mu)$ and $\mathcal{Q}(\nu)$, by Lemma 3.2.

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 $W \Cap V$ are

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

and so that for each $i = 1, \dots, k$ one of the following holds: either $W_i \overset{i}{\equiv} V_i \overset{i}{\equiv} U_i$; 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 $U = X \Cap Y$ 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 $X \Cap Z$ 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 $W_i \Cap V_j \neq \emptyset$ 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 $Q(\nu)$ 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 a quasi-isometric embedding $\mathcal{M}(U_i) \to \mathcal{M}(W_i) \approx \mathcal{TM}(W_i)$ whose image is a subtree denoted $\tau_i \subset \mathcal{TM}(W_i)$: either $U_i \overset{i}{\equiv} W_i$ and $\tau_i = \mathcal{TM}(W_i)$; or U_i is an annulus and τ_i is the axis in $\mathcal{TM}(W_i)$ of the Dehn twist about U_i. Similarly, the inclusion $U_i \subset V_i$ induces a quasi-isometric embedding $\mathcal{M}(U_i) \to \mathcal{TM}(V_i)$ whose image is a subtree $\sigma_i \subset \mathcal{TM}(V_i)$.

By composing a coarse inverse of the map $\mathcal{M}(U_i) \to \tau_i$ with the map $\mathcal{M}(U_i) \to \sigma_i$, we obtain a quasi-isometry $g_i \colon \tau_i \to \sigma_i$. Notice that we may take g_i to be a simplicial isomorphism, as one can verify easily in either of two cases: if $W_i \overset{i}{\equiv} V_i \overset{i}{\equiv} U_i$ then these isotopies induce simplicial isomorphisms of marking complexes; and otherwise τ_i and σ_i are the axes in the trees $\mathcal{TM}(W_i)$ and $\mathcal{TM}(V_i)$, 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 $\mathcal{TM}(W_i)$ and $\mathcal{TM}(V_i)$ by gluing τ_i to σ_i isometrically using the map g_i. Let

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

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 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

$$j^\mu \colon \mathcal{Q}(\mu) \to \Upsilon$$

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 $\mathcal{TM}(V_j)$, $j = k+1,\ldots,n$, and to 0 in the $\mathbb{R}$ factor. Similarly we have

$$j^\nu\colon \mathcal{Q}(\nu) \to \Upsilon,$$

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 $\mathcal{TM}(W_j)$, $j = k+1,\ldots,m$, 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 3.2, that the images $j^\mu(E(\mu,\nu))$ and $j^\nu(E(\nu,\mu))$ 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 3.2 (and its interpretation in terms of projections in (3.2,3.3,3.4,3.5)), j^μ and j^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 j.

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

Lemma 3.4. *Let $T = T_1 \times \cdots \times T_N$ 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. □

Now it is easy to understand how the cubes $C(\mu,W,r)$ and $C(\nu,V,s)$ are situated by considering their j-images. $j(C(\mu,W,r))$ is a product of lines and points in the factors of Υ, with first coordinate 0, and $j(C(\nu,V,s))$ is a similar product with first coordinate D. Lemma 3.4 implies that the junctures of the images are again products of subintervals, and we conclude that the j-preimages, which are subcubes of the original cubes, are also junctures. □

The proof of Lemma 3.3 gives some more information about the structure of the junctures of two cubes, which we record here:

Lemma 3.5. *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.*

Moreover the components r_i and s_i that are not single points are associated to those pairs W_i, V_i for which $W_i \pitchfork V_i \neq \emptyset$, after the renumbering in the proof of Lemma 3.3

Proof. This is a consequence of the fact that the map j in the proof of Lemma 3.3 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 $X_1, \ldots, X_k$ and $\mathcal{TM}(W_{k+1}), \ldots, \mathcal{TM}(W_m)$, and points in the other factors, wherease $C(\nu, V, s)$ maps to a product of line segments in $X_1, \ldots, X_k$ and in $\mathcal{TM}(V_{k+1}), \ldots, \mathcal{TM}(V_n)$, and points in the rest. Thus, the only tree factors in which the junctures can be nondegenerate segments are the X_i, and these correspond to the pairs W_i, V_i that have nontrivial essential intersection. The pullbacks of the junctures by j are then subcubes respecting the product structures of the original cubes, and with nondegenerate segments only in the factors corresponding to the X_i. □

3.3. **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 $\mathcal{M}_\omega$. 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 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 $C^\omega(\overline{\mu}, \overline{W}, \overline{r})$ (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 a line in $\mathcal{M}_\omega(\overline{W}_i)$ whose length is in $[0, \infty]$. We will continue calling these objects cubes.

Lemma 3.3 on junctures for cubes implies, using the discussion in the beginning of this section that the intersection of two cubes is empty or is a cube (possibly a trivial cube, i.e., 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.

4. Consistency theorem

In this section we will derive a coarse characterization of the image of the map

$$\Pi\colon \mathcal{M}(S) \to \prod_{W\subseteq S} \mathcal{C}(W)$$

defined by curve-complex projections, $\Pi(\mu) = (\pi_W(\mu))_W$.

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 $W \pitchfork V$,

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

C2: Whenver $V \Subset W$ and $d_W(x_W, \partial V) > c_2$,

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

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 [1], namely

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

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

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 $V \Subset W$, $\pi_V \circ \pi_W$ is a bounded distance from π_V when both are defined (Lemma 2.11).

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 $W \subseteq S$.

In order to approach the proof we will first study more carefully the structure imposed by (C1) and (C2).

4.1. **Ordering induced by projections**

Let us fix a tuple (x_W) satisfying (C1-2). Without loss of generality, we will assume $c_1 > \max\{c_2, m_0, B\}$, here m_0 is the constant given by Lemma 4.2 and B is the constant given by Theorem 2.5.

If W, V are proper 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.

Although $\prec_k$ is not quite an order relation, 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

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

We then define $W \ll_k V$ to mean $W \ll_k \partial V$. This is a weaker relation than $\prec_k$ because $W \pitchfork \partial V$ allows the possibility that $V \subsetneq W$.

Note first 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 translates to the statement that

- $\prec_k$ is antisymmetric: 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 $U \prec_{k-1} W$.*
2. *If $U \ll_k V$ and $V \ll_2 W$ then $U \ll_{k-1} W$.*

(3) *If* $U \pitchfork V$ *and both* $U \ll_k \rho$ *and* $V \ll_k \rho$ *for a partial marking* ρ*, then* U *and* V *are* $\prec_{k-1}$*-ordered – that is, either* $U \prec_{k-1} V$ *or* $V \prec_{k-1} U$*.*

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 $\operatorname{diam}\pi_{\mathcal{C}(V)}(\gamma) \leq 2$ for any disjoint curve system γ (this is very similar to Lemmas 2.9 and 2.10),

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

In particular $d_V(\partial U, \partial W) > 2$, so ∂U and ∂W must intersect, so that $U \pitchfork W$. 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.

The proof of part (2) is similar. The case not covered by part (1) is when $V \subsetneq U$. Suppose first that $W \subsetneq V$. Then ∂W and ∂V form a curve system in U, and hence $\operatorname{diam}_U(\partial V \cup \partial W) = 1$. 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 $U \ll_{k-1} W$.

Now suppose $W \pitchfork V$. 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

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

But now by Theorem 2.5, this implies that any $\mathcal{C}(U)$-geodesic $[x_U, \pi_U(\partial W)]$ must pass within distance 1 of ∂V, and we conclude

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

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

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. □

Note that the weak transitivity proved in parts (1) and (2) of Lemma 4.4 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.

Let us now define

$$\mathcal{F}_k(\rho) = \{W \subset S : W \ll_k \rho\}.$$

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 $\mathcal{F}_k(Z)$.*

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

Suppose $U, V, W \in \mathcal{F}_k(Z)$, 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. □

We can also obtain this finiteness statement:

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

Proof. Suppose that $\mathcal{F}_k(Z)$ is infinite and let $\{Y_i\}$ be an infinite sequence within it. After extracting a subsequence we may assume that $\partial Y_i \to \lambda$ 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 i (eventually), and $\pi_U(\partial Y_i) \to \infty$ in $\mathcal{C}(U)$ – that is, $d_U(\partial Y_i, q) \to \infty$

for any fixed q. (This is a consequence of the Kobayashi/Luo argument that $\mathcal{C}(U)$ has infinite diameter, see [20, 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, $d_U(x_U, \partial Y_i) \to \infty$ means for any given p that eventually $U \ll_p Y_i$. However we have $Y_i \ll_k Z$ by assumption, so

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

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

4.2. **Proof of the consistency theorem**

As before we have (x_W) satisfying C1-2 with the same assumptions on c_1 and c_2. We will construct μ by induction.

Consider $\mathcal{F}_3(x_S)$. If it is empty, let $\mu_0 = x_S$. If not, by Lemmas 4.6 and 4.5 it is finite and partially ordered by $\prec_2$ – hence it contains minimal elements. Among these minimal elements, choose one, Y, of maximal complexity $\xi(Y)$, and let $\mu_0 = \partial Y$.

Now consider any $Z \subseteq S$ which intersects μ_0. We claim that

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

Suppose otherwise, so $d_Z(x_Z, \mu_0) \geq 4(c_1 + 4)$. If $\mathcal{F}_3(x_S) = \emptyset$ and $\mu_0 = x_S$ this would imply $Z \ll_4 x_S$ so $Z \in \mathcal{F}_4(x_S) \subset \mathcal{F}_3(x_S)$, 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 $Y \subset Z$. 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 $V \prec_1 Y$. Arguing as in the proof of Lemma 4.5, we obtain 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.

Our main goal in this section will be Proposition 5.2, in which we show that Σ-hulls admit coarse retractions. In Section 6 we will apply this to 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 $\mathrm{hull}_X(A)$ denote the union of geodesics $[a, a']$ with $a, a' \in A$. We will need the following properties of this construction, which are easy exercises.

Lemma 5.1. *The hyperbolic hull* $\mathrm{hull}_X(A)$ *is quasi-convex. There is a coarse-Lipschitz nearest-point retraction* $X \to \mathrm{hull}_X(A)$*. Moreover, if* $x \in X$ *and* $y \in \mathrm{hull}_X(A)$ *is a nearest point, and* $y' \in \mathrm{hull}_X(A)$ *then* $d(y, y')$ *is bounded just in terms of* $d(x, y') - d(x, y)$*. The map* $A \mapsto \mathrm{hull}(A)$ *is Lipschitz in the Hausdorff metric.*

The implicit constants in these statements depend only on the hyperbolicity constant of X *and the cardinality of* A*.*

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

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

Here W varies over all essential subsurfaces of S (including S) and $\mathrm{hull}_W(A)$ is the hyperbolic hull as defined in §5.1.

It is clear that $A \subset \Sigma_\epsilon(A)$, but a priori not much else. (For the reader familiar with the constructions in Masur-Minsky [21], 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. *Given a finite set $A \subset \mathcal{M}(S)$ there exists a map*

$$p_A \colon \mathcal{M} \to \Sigma_\epsilon(A)$$

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* $\mathrm{hull}_W(A)$ *to $\pi_W(x)$. Then*

$$d_W(x, y_W)$$

   is uniformly bounded.

The implicit constants depend only on ϵ, the topology of S, and the cardinality of A. The construction holds for any sufficently large choice of ϵ.

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 $\mathrm{hull}_W(A)$.

Lemma 5.3. *For any x and A, 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

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

so we may suppose $d_U(x, \partial V) < m_0$. Now if $d_U(x, y_U) < 2m_0 + 2$, we are done: we have $d_U(y_U, \partial V) < 3m_0 + 4$. (Recall that $\mathrm{diam}_U(\partial V)$ and $\mathrm{diam}_U(x)$ are bounded by 2.)

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

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

Now since $y_V \in \mathrm{hull}_V(A)$, there must be $a, b \in A$ such that $y_V \in [a', b']$ for $a' \in \pi_V(a)$, $b' \in \pi_V(b)$. Now $d_U(\partial V, a)$ and $d_U(\partial V, b)$ are at least m_0 by (5.1), and it follows by Lemma 4.2 that $d_V(\partial U, a) < m_0$ and $d_V(\partial U, b) < m_0$. Since $y_V \in [a', b']$ we conclude (again applying the triangle inequality appropriately) 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 2.5, 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 2.5 that $\operatorname{diam}_V(\gamma) < m_0$, and hence $\operatorname{diam}_V(A) < m_0$.

Moreover, since y_W itself is on such a geodesic, $d_V(y_W, A) < m_0$. Since $y_V \in \operatorname{hull}_V(A)$ 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 2.5 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 y_W.

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

Using this lemma, 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 define $p_A(x) \equiv \mu$, 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 $d_W(x, \operatorname{hull}_W(A)) \leq \epsilon$. In particular, if we choose ϵ larger than the above constant c_3, then $d_W(x, y_W) \leq \epsilon$. 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 $d_W(x, \mu)$. The quasidistance formula (Corollary 2.8) 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 2.10) 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

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

by Lemma 5.1, and if y'_W is a nearest point to $\pi_W(x')$ in $\operatorname{hull}_W(A')$, that Lemma also gives us 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, and 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 exist K and ϵ', depending on ϵ and $\#A$, such that*

(1) $\mathrm{diam}(\Sigma_\epsilon(A)) \le K\mathrm{diam}(A) + K$
(2) *If* $A' \subset \Sigma_\epsilon(A)$ *then* $\Sigma_\epsilon(A') \subset \Sigma_{\epsilon'}(A)$.
(3) $d_H(\Sigma(A'), \Sigma(A)) \le K d_H(A', A)$.

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

To see Part (3), note by Part (2) of Proposition 5.2, p_A and $p_{A'}$ differ by $Cd_H(A', A)$. Since $p_{A'}$ is uniformly close to the identity on $\Sigma_\epsilon(A')$, p_A must be within $C'd_H(A', A)$ of the identity on $\Sigma_\epsilon(A')$, for some C'. It follows that $\Sigma_\epsilon(A')$ is within $Kd_H(A', A)$ of $\Sigma_\epsilon(A)$, for some K. The opposite inclusion is obtained in the same way. □

We shall also have use for the following lemma.

Lemma 5.5. *If $A \subset \mathcal{M}(S)$ is any subset and a subsurface $W \subset S$ satisfies*

$$\mathrm{diam}_{\mathcal{C}(W)}(A) > 2m_0,$$

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

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

Proof. Since $U \pitchfork \partial W$ we either have $W \subset U$ or $W \pitchfork U$; we treat these two cases separately.

First, if $W \subset U$ then Theorem 2.5 immediately implies that

$$d_{\mathcal{C}(U)}(\partial W, \mathrm{hull}_U(A)) \le 1.$$

Now, consider the case that $W \pitchfork U$. Since $\mathrm{diam}_{\mathcal{C}(W)}(A) > 2m_0$, it follows that there exists $a \in A$ for which $d_{\mathcal{C}(W)}(a, \partial U) \ge m_0$. Then by Lemma 4.2 it follows that $d_{\mathcal{C}(U)}(a, \partial W) < m_0$. □

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 $\boldsymbol{A}$ is a finite set in $\mathcal{M}_\omega$ represented by a sequence (A_n), we denote by $\Sigma(\boldsymbol{A})$ the ultralimit of the coarse Σ-hulls $\Sigma_\epsilon(A_n)$, where ϵ is a fixed (sufficiently large) constant. Note that increasing ϵ does not change $\Sigma(\boldsymbol{A})$, nor does changing the representatives – this follows from Lemma 5.4 part (3). In fact Lemma 5.4 applied in the limit gives:

Lemma 6.1.

(1) $\mathrm{diam}(\Sigma(\boldsymbol{A})) \le K\mathrm{diam}(\boldsymbol{A})$ *where K is an a priori constant.*

(2) *If* $\boldsymbol{A}' \subset \Sigma(\boldsymbol{A})$ *then* $\Sigma(\boldsymbol{A}') \subset \Sigma(\boldsymbol{A})$.

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

$$p_{\boldsymbol{A}} : \mathcal{M}_\omega \to \Sigma(\boldsymbol{A}).$$

Moreover, Proposition 5.2 implies that $p_{\boldsymbol{A}}$ is jointly continuous in its arguments and in the points of $\boldsymbol{A}$. With this we can establish:

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

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

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

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

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 [3] 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 [16] (Theorem 6.7 below) and its corollary 6.8, 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 continuous map $f\colon P \to \mathcal{M}_\omega$ from a polyhedron to $\mathcal{M}_\omega$ 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. Note that if $f\colon P \to \mathcal{M}_\omega$ is Σ-compatible, then for every face $\tau \subset P$,

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

where $C = C(\dim \tau)$.

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* $f_0\colon Q \to \mathcal{M}_\omega$ *can be extended to a* Σ-*compatible map* $f\colon P \to \mathcal{M}_\omega$.

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 $f_0 \colon P \to \mathcal{M}_\omega$ is a map from a finite polyhedron to $\mathcal{M}_\omega$, then there is a map $f_1 \colon P \to \mathcal{M}_\omega$ 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 [3], there is an open cover $\mathbf{U} = \{U_i\}_{i \in I}$ of Y such that $P' := \mathrm{Nerve}(\mathcal{U})$ has dimension at most $\xi(S)$, and $\mathrm{diam}(U_i) < \rho$ for all $i \in I$. Let $\{\phi_i \colon Y \to [0,1]\}_{i \in I}$ be a partition of unity subordinate to $\mathcal{U}$, and $\phi \colon Y \to P'$ 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 $\alpha \colon P' \to \mathcal{M}_\omega$ with the property that $\alpha(U_i) = x_i$ (recall that the vertex set of $\mathrm{Nerve}(\mathcal{U})$ consists of elements of $\mathcal{U}$).

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

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 $U_{i_1}, \ldots, U_{i_k}$, then $f_0(x) \in U_{i_1} \cap \ldots \cap U_{i_k}$, and

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

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

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

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 $f_0, f_1 \colon P \to \mathcal{M}_\omega$, 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 $\rho > 0$. We may assume without loss of generality that for $i \in \{0,1\}$ and every face τ of P,

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

Let $P = P_1, P_2, \ldots, P_k, \ldots$ 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 $f_{i,k} \colon P_k \to \mathcal{M}_\omega$ 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 $k \to \infty$.

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

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

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 simplex τ of P_k. By the uniform continuity of f_i and the fact that $f_{i,k} \to f_i$ uniformly, it follows that the diameter of $f_{i,k}(\tau) \cup f_{i,k+1}(\mathrm{Sd}\,\tau)$ tends to zero as $k \to \infty$, where $\mathrm{Sd}\,\tau$ is the barycentric subdivision of τ. 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 a vertex in $\tau \times \{0\}$ or $\mathrm{Sd}\,\tau \times \{1\}$. By Lemma 6.3, we get $d(H_{i,k}(x,t), f_i(x)) < \delta_k$, where $\delta_k < \rho$ and $\delta_k \to 0$ as $k \to \infty$. It follows that the concatenation of

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

has tracks of diameter tending to zero as $j \to \infty$, 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. ☐

As a corollary of Lemmas 6.4 and 6.5, we obtain the following statement on homological dimension of the cone. Note that in Hamenstädt's approach [10] the homological statement comes directly.

Corollary 6.6. *If (U,V) is an open pair in $\mathcal{M}_\omega$, 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 $f_0 \colon (P,Q) \to (U,V)$ such that

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

Pick $\epsilon > 0$. Applying Lemma 6.4, we obtain a continuous map $f_1 \colon P \to \mathcal{M}_\omega$ 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 $(P,Q) \to (U,V)$, and the homotopy $\{f_t\}$ will be a homotopy of maps of pairs, so that f_0 and

f_1 induce the same map $H_k(P,Q) \to H_k(U,V)$. 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 $(P,Q) \to (P',Q') \to (U,V)$, where (P',Q') is a polyhedral pair of dimension $\leq \xi(S)$. This implies that $f_{1*} = 0$. Hence $[c] = 0$. □

We are now in a position to apply the following results of Kleiner-Leeb [16], which we will be using in the proof of Theorems 7.11 and 8.8.

Theorem 6.7. *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*

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

is injective for any $p \in M$.

Corollary 6.8. *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 6.6, we will be able to apply Corollary 6.8 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 7.3 and Theorem 7.8, which are concerned with separation properties of *microscopic* and *macroscopic* jets, respectively. Much of the technical work is done in Lemma 7.2. 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 [21]. If $\xi(S) \geq 2$, a *tight geodesic* in $\mathcal{C}(S)$ is a sequence of simplices $\{w_0, \ldots, w_n\}$ if any selection of vertices $v_i \in w_i$ yields a geodesic in the 1-skeleton of $\mathcal{C}(S)$, and if for each $1 \leq i \leq n-1$, the system w_i is the boundary of the subsurface filled by v_{i-1} and v_{i+1}. If $\xi(S) = 1$, every geodesic is considered to be tight. If $Y \subset S$ is an 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 [21] that any two vertices can be joined by a tight geodesic, and there are only finitely many possibilities.

Let $a, b \in \mathcal{M}(S)$, $W \subseteq S$ a connected essential subsurface, and let g be a (directed) tight geodesic in $\mathcal{C}(W)$ from $\pi_{\mathcal{C}(W)}(a)$ to $\pi_{\mathcal{C}(W)}(b)$. If $\sigma = [\alpha, \beta]$ is a subsegment of g, we call (σ, a, b) a *tight triple.* Let $|\sigma|$ denote the length of σ (measured in $\mathcal{C}(W)$).

We also associate to the triple (σ, a, b) the following pair of points in the marking graph of W: Let $\iota(\sigma, a, b)$ be $\alpha \rfloor \pi_{\mathcal{M}(W)}(a)$, the *initial marking,* and Let $\tau(\sigma, a, b)$ be $\beta \rfloor \pi_{\mathcal{M}(W)}(b)$, the *terminal marking.* (One can also, up to the usual bounded ambiguity, 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

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

Using Theorem 2.7 we can establish the following properties of this notion of size:

Lemma 7.1. *Given $\sigma = [\alpha, \beta], W, a, b$ as above, let $\Phi(\sigma)$ denote the set of subsurfaces $U \subset W$ that are disjoint from some simplex of σ.*

(1)
$$||\sigma||_{(W,a,b)} \stackrel{p,q}{\approx} d_{\mathcal{C}(W)}(\alpha, \beta) + \sum_{U \in \Phi(\sigma)} \left\{\!\!\left\{ d_{\mathcal{C}(U)}(a, b) \right\}\!\!\right\}_A.$$

(2) *If σ is written as a concatenation of subintervals $\sigma_1 * \cdots * \sigma_k$, then*
$$||\sigma||_{W,a,b} \stackrel{p}{\approx} \sum_i ||\sigma_i||_{W,a,b}$$

where the constants A, p, q depend only on the complexity of W.

Proof. Let $\iota = \iota(\sigma, a, b)$ and $\tau = \tau(\sigma, a, b)$. To prove (1), we first recall that Theorem 2.7 gives us, for large enough A and uniform p, q (depending on A)

$$d_{\mathcal{M}(W)}(\iota, \tau) \stackrel{p,q}{\approx} \sum_{Y \subseteq W} \left\{\!\!\left\{ d_{\mathcal{C}(Y)}(\iota, \tau) \right\}\!\!\right\}_A.$$

On the other hand, Theorem 2.5 gives a constant B such that, if $Y \subsetneq W$ intersects every simplex of σ, then

$$\mathrm{diam}_{\mathcal{C}(Y)}(\sigma) \leq B.$$

Now in particular Y intersects α and β, and since these are contained in ι and τ respectively, we get

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

Thus, if we require $A > B$ all of these terms drop out of the sum. leaving the single $Y = W$ term, and what is almost the summation in (1), except that distances are between ι and τ instead of a and b. Now consider a Y which is disjoint from a simplex v of σ. By tightness of the geodesic g containing

σ, the set $\phi(Y)$ of simplices disjoint from Y is a contiguous sequence of at most 3 simplices. Hence, if for example Y intersects β then the rest of g between β and $\pi_{\mathcal{C}(W)}(b)$ consists of simplices intersecting Y, and so (since τ contains β) Theorem 2.5 implies

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

If Y does not intersect β then $Y \subset W_\beta$ and so since by definition τ restricted to W_β is $\pi_{\mathcal{M}(W_\beta)}(b)$, 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. □

A *jet*, denoted J, is a quadruple $(\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$, where (σ_n, a_n, b_n) are tight triples with σ_n supported in W_n, and we assume that $\overline{a}$ and $\overline{b}$ have ultralimits in $\mathcal{M}_\omega(S)$ (i.e., that they do not go to ∞ faster than linearly). We refer to $\overline{W}$ as the *support surface* of the jet J. Call a jet *microscopic* if $||\sigma_n||_{(W_n, a_n, b_n)}$ grows sublinearly – that is, if

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

Often we write $||\sigma_n||_J$ to denote $||\sigma_n||_{(W_n, a_n, b_n)}$.

Any jet, J, has a corresponding sequence of initial points $\iota_n = \iota(\sigma_n, a_n, b_n)$. This sequence defines a point $\iota_\omega(\overline{\sigma}, \overline{a}, \overline{b}) \in \mathcal{M}_\omega(\overline{W})$, which we will call the *basepoint* of the jet and denote $\boldsymbol{\iota}(J)$ (or just $\boldsymbol{\iota}$ when J is understood); similarly one obtains $\boldsymbol{\tau}(J)$.

Sudden and gradual growth of jets. A jet J is *macroscopic* if it is not microscopic. Note that a jet is macroscopic if and only if $\boldsymbol{\iota}(J) \neq \boldsymbol{\tau}(J)$. For such jets, the way in which the linear growth happens turns out to be important:

Let α_n denote the initial vertex of σ_n. We will say that the jet J has *sudden growth* if the following happens: There are 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.

7.2. Projections and separation

Let $J = (\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$ be a microscopic jet with basepoint $\boldsymbol{\iota} \in \mathcal{M}_\omega(\overline{W})$. 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

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

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

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

If (ξ_n) is a sequence in $\mathcal{M}(S)$, let $\pi_{\sigma_n}(\xi_n)$ denote the composition of projection of ξ_n into $\mathcal{C}(W_n)$ with closest-point projection in $\mathcal{C}(W_n)$ into σ_n.

Lemma 7.2. *There exist $K, c > 0$ such that for any microscopic jet J: If (ξ_n) and (ξ'_n) are sequences in $\mathcal{M}(S)$ such that for ω-a.e. n we have $d_{\mathcal{C}(W_n)}(\pi_{\sigma_n}(\xi_n), \pi_{\sigma_n}(\xi'_n)) > K$, then*

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

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

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

Where A is any sufficiently large threshold and the constants a, b depend only on A.

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

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

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

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

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)

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

where p', p'' 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

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

and

$$a_n = \pi_{\sigma_n}(\xi_n), \qquad a'_n = \pi_{\sigma_n}(\xi'_n).$$

Let $h_n = [x_n, a_n]$ and $h'_n = [x'_n, a'_n]$ be $\mathcal{C}(W_n)$-geodesic segments. Because a_n is a nearest point to x_n on σ_n (and similarly for a'_n and x'_n), and $\mathcal{C}(W_n)$ is δ-hyperbolic, there is a constant K_δ such that, if $d(a_n, a'_n) > K_\delta$, 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'.

If $Y = W_n$, then we immediately find that

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

which is (7.1) in this case.

If $Y \subset W_n$, let B be the bound in Theorem 2.5. Suppose first that ∂Y is outside a 2-neighborhood of h_n. Then $a_n \cap Y \neq \emptyset$, and $d_{\mathcal{C}(Y)}(x_n, a_n) \leq B$. Moreover, ∂Y can only be disjoint from simplices on one side of a_n in σ_n (not both) since σ_n is a tight geodesic. It follows that $\pi_Y(a_n)$ 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 is within a 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 not within a 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 $Y \pitchfork \partial W_n$, and $d_{\mathcal{C}(Y)}(\xi_n, \partial W_n) > A$. Following the discussion in Section 4, define for $k > 0$ the relation $U \ll_k \gamma$ by

$$U \pitchfork \gamma \text{ and } d_{\mathcal{C}(U)}(\xi_n, \partial V) \geq kc$$

where $c > 6 + \max\{4, m_0, B\}$, and let $U \pitchfork V$ denote $U \pitchfork \partial V$. (In other words, $(\pi_U(\xi_n))$ is taking the role of (x_U) in Section 4). We may assume $A > 4c$, and setting $k = \lfloor A/c \rfloor$, we have

$$Y \ll_k W_n.$$

Moreover we have $d_{\mathcal{C}(W_n)}(\xi_n, \xi'_n) > d_{\mathcal{C}(W_n)}(a_n, a'_n) - \delta' > K - \delta'$, so assuming $K - \delta' > 2c$ we get

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

Now Lemma 4.4 implies that

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

so in particular

$$d_{\mathcal{C}(Y)}(\xi_n, \xi'_n) \geq (\lfloor A/c \rfloor - 1)c \geq d_{\mathcal{C}(Y)}(\xi_n, \partial W_n) - 2c. \tag{7.5}$$

This gives us (7.2).

To complete the proof, we note that, by the definition of ι_n and τ_n, whenever $Y \subseteq W_n$ we have $\operatorname{diam}_{\mathcal{C}(Y)}(\sigma_n) \geq d_{\mathcal{C}(Y)}(\iota_n, \tau_n)$. Thus the sum of all the $\operatorname{diam}_{\mathcal{C}(Y)}(\sigma_n)$ terms in (7.1) bounds $d_{\mathcal{M}(W_n)}(\iota_n, \tau_n)$ via Theorem 2.7 applied to $\mathcal{M}(W_n)$, and this establishes (7.3). Thus after rescaling and ultralimits we obtain the desired inequality. □

Projection equivalence. We can define a relation on sequences (x_n) in $\mathcal{M}(S)$ as follows: Say that $(x_n) \sim_{\overline{\sigma}} (x'_n)$ 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 $\mathcal{M}(\overline{S})$. We will deduce the following stronger result.

Theorem 7.3. *For any microscopic jet J, the relation $\sim_{\overline{\sigma}}$ descends to an equivalence relation on $\mathcal{M}_\omega(S) \setminus \mathcal{L}_\omega(J)$. Moreover, every equivalence class is open.*

In fact, the following more quantitative statement is an immediate consequence of Lemma 7.2:

Corollary 7.4. *There exists $c > 0$ such that for any microscopic jet J: If (ξ_n) and (ξ'_n) are sequences in $\mathcal{M}(S)$ which are* inequivalent *under $\sim_{\overline{\sigma}}$, then*

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

We now show how the corollary implies the theorem.

Proof of Theorem 7.3. If $\boldsymbol{\xi} = \boldsymbol{\xi}'$ then Corollary 7.4 implies either $(\xi_n) \sim_{\overline{\sigma}} (\xi'_n)$ or $\boldsymbol{\xi} \in \mathcal{L}_\omega(J)$. Hence in the complement of $\mathcal{L}_\omega(J)$ the equivalence relation $\sim_{\overline{\sigma}}$ descends to an equivalence relation in the asymptotic cone.

Further, Corollary 7.4 implies that a definite neighborhood of $\boldsymbol{\xi} \notin \mathcal{L}_\omega(J)$ consists of points represented by sequences which are $\sim_{\overline{\sigma}}$ equivalent to $\boldsymbol{\xi}$. Hence equivalence classes are open. □

Finding microscopic jets. The following counting argument will provide a useful way to construct microscopic jets.

Lemma 7.5. *Let $\overline{\mu}, \overline{\nu} \in \mathcal{M}(\overline{W})$, with $\mu_\omega, \nu_\omega \in \mathcal{M}_\omega(\overline{W})$. Suppose that*

$$d_{\mathcal{C}(W_n)}(\mu_n, \nu_n) \to_\omega \infty$$

and let ℓ_n be tight $\mathcal{C}(W_n)$-geodesics connecting $x_n \in \pi_{\mathcal{C}(W_n)}(\mu_n)$ to $y_n \in \pi_{\mathcal{C}(W_n)}(\nu_n)$. Then there exist a sequence $\sigma_n \subset \ell_n$ of subsegments such that $|\sigma_n| \to_\omega \infty$, but

$$\frac{1}{s_n}||\sigma_n||_{(W_n,\mu_n,\nu_n)} \to 0$$

Proof. Let $f(n)$ be a function going to ∞ more slowly than $|\ell_n|$; that is, such that $f(n) \to \infty$ but $|\ell_n|/f(n) \to \infty$. Divide ℓ_n into $f(n)$ equal segments, each of length $|\ell_n|/f(n)$ (the last one may be smaller). The sum of the $||\cdot||_{(W_n,\mu_n,\nu_n)}$-sizes of these is approximated by $||\ell_n||_{(W_n,\mu_n,\nu_n)}$ 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, one segment σ_n with $||\sigma_n||_{(W_n,\mu_n,\nu_n)} \leq Cs_n/f(n)$. □

This lemma has the following consequence which we utilize in Sections 7.4 and 8.

Lemma 7.6. *Let $a_\omega, b_\omega \in \mathcal{M}_\omega(\overline{S})$ be represented by $\overline{a}$ and $\overline{b}$, and suppose that $d_{\mathcal{C}(W_n)}(a_n, b_n) \to_\omega \infty$. Then there exists a microscopic jet $J = (\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$ such that*

$$\overline{a} \not\sim_{\overline{\sigma}} \overline{b}.$$

Proof. Let $\mu_n = \pi_{\mathcal{M}W_n}(a_n)$ and $\nu_n = \pi_{\mathcal{M}W_n}(b_n)$ and apply Lemma 7.5 to a sequence ℓ_n of tight $\mathcal{C}(W_n)$-geodesics between them. This yields a sequence of subsegments σ_n with $|\sigma_n| \to \infty$ but $||\sigma_n||_{(W_n,a_n,b_n)}$ growing sublinearly, hence $J = (\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$ is a microscopic jet. (The projections of μ_n, ν_n and a_n, b_n to each subsurface here are the same up to bounded error, by Lemma 2.11).

By construction, a_n and b_n project to opposite ends of σ_n, and are therefore inequivalent under $\sim_{\overline{\sigma}}$. □

7.3. Linear/sublinear decomposition

Now consider a macroscopic jet $J = (\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$. Let α_n be the initial vertices of σ_n, and ι_n the initial markings, as above.

Suppose that J has gradual growth. This allows us to make the following definition: We say that $\overline{z} \in \mathcal{M}(\overline{S})$ *escapes linearly* along J if

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

has linear growth. If this is not true we say it *escapes sublinearly.* Note that the gradual growth condition implies that (ι_n) itself escapes sublinearly. Again we will consider $\mathcal{L}_\omega(J)$.

Lemma 7.7. *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 $\boldsymbol{z} \in \Lambda_J$ implies that any sequence (z_n) representing $\boldsymbol{z}$ escapes linearly along J, and $\boldsymbol{z} \in \Omega_J$ implies any (z_n) escapes sublinearly.

A set $X \subset \mathcal{M}_\omega(S)$ is Σ-*convex* if, for any finite set $\boldsymbol{A} \subset X$ the hull $\Sigma(\boldsymbol{A})$ is in X as well. We will establish the following.

Theorem 7.8. *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 in particular 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 7.3.)

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

Lemma 7.9. *There exists $c > 0$ such that the following holds for any macroscopic jet J with gradual growth: Suppose that (ξ_n) and (ξ'_n) are such that*

(1) $\boldsymbol{\xi} \notin \mathcal{L}_\omega(J)$, *and*
(2) $d_{\mathcal{M}_\omega(S)}(\boldsymbol{\xi}, \boldsymbol{\xi}') < cd(\boldsymbol{\xi}, \mathcal{L}_\omega(J))$

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

Proof of Lemma 7.9: Write $J = (\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$, let c be the constant in Lemma 7.2, and suppose, by way of contradiction, that (1) and (2) hold but one of $\overline{\xi}, \overline{\xi}'$ escapes sublinearly and the other escapes linearly. After renaming the one that escapes linearly $\overline{\zeta}$ and the one that escapes sublinearly $\overline{\eta}$, 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 $\overline{\eta}$ escapes sublinearly, the restricted jet

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

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 7.2. In this new jet $J'' = (\overline{\sigma}'', \overline{W}, \overline{a}, \overline{b})$ 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 7.2 implies that $d(\boldsymbol{\xi}, \boldsymbol{\xi}') \geq cd(\boldsymbol{\xi}, \mathcal{L}_\omega(J))$. (Note that $\boldsymbol{\iota}(J) = \boldsymbol{\iota}(J'')$, so $\mathcal{L}_\omega(J) = \mathcal{L}_\omega(J'')$.) This contradicts our hypothesis that $d(\boldsymbol{\xi}, \boldsymbol{\xi}') < cd(\boldsymbol{\xi}, \mathcal{L}_\omega(J))$. Hence it must hold that $\overline{\xi}$ escapes linearly if and only if $\overline{\xi}'$ does. □

Lemma 7.7 follows immediately from Lemma 7.9 by considering the case $\boldsymbol{\xi} = \boldsymbol{\xi}'$.

Proof of Theorem 7.8. The openness of Λ_J and Ω_J is an easy consequence of Lemma 7.9.

It remains to prove the last part of the theorem, that Λ_J is Σ-convex and hence acyclic.

Let $\boldsymbol{A} \subset \Lambda_J$ be finite and let $\overline{A}$ represent it. Then each $\overline{a} \in \overline{A}$ 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 $\mathrm{hull}_{W_n}(A_n)$ 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.

Any compact singular chain $f\colon P \to \Lambda_J$ can be refined and then approximated, by a Σ-compatible chain, and openness allows us to do this within Λ_J: We 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 $f'\colon Q \to \mathcal{M}_\omega(S)$ is sufficiently close to the old to still be in Λ_J. 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 Λ_J.

Now we can cone off Q towards one vertex getting a polyhedron with the same 0-skeleton, and applying Lemma 6.3 extend f' to a Σ-compatible map on the cone. By Σ-convexity this chain lies in Λ_J, so our original chain bounds in Λ_J, and Λ_J is acyclic. $\square$

7.4. **Classification of pieces**

We now record an application of Theorem 7.8, which classifies the maximal subsets of $\mathcal{M}_\omega$ which can not be separated by a point. First, let us recall the notions of *pieces* and *tree-graded* spaces as defined by Druţu–Sapir [5].

A complete geodesic metric space X is called *tree-graded* if it 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 there exists a unique *finest* decomposition into pieces which can not be separated by a point [5].

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 [2] such pieces can not be realized as asymptotic cones of subgroups of $\mathcal{M}(S)$.

Theorem 7.10. *Fix a pair of points $\boldsymbol{\mu}, \boldsymbol{\nu} \in \mathcal{M}_\omega(S)$. If $\xi(S) \geq 2$, then the following are equivalent.*

(1) *No point of $\mathcal{M}_\omega(S)$ separates $\boldsymbol{\mu}$ from $\boldsymbol{\nu}$.*
(2) *There exist points $\boldsymbol{\mu}', \boldsymbol{\nu}'$ arbitarily close to $\boldsymbol{\mu}, \boldsymbol{\nu}$, resp., for which there exist representative sequences $(\mu_n'), (\nu_n')$ satisfying*

$$\lim_\omega d_{\mathcal{C}(S)}(\mu_n', \nu_n') < \infty.$$

Proof. We begin by showing (2) implies (1). Suppose first that $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ 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 $v_{n,0}, \ldots, v_{n,m}$ denote the simplices of a tight geodesic in $\mathcal{C}(S)$ connecting $v_{n,0} \in \mathrm{base}(\mu_n)$ to $v_{n,m} \in \mathrm{base}(\nu_n)$. For a fixed i let $\overline{v}_i = \langle v_{n,i} \rangle$. The regions $\mathcal{Q}(v_{n,i})$ have the structure described in Lemma 3.1, and in particular the cone $\mathcal{Q}_\omega(\overline{v}_i)$ 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 $v_{n,i} \rfloor v_{n,i+1} = v_{n,i} \cup v_{n,i+1}$, so by Lemma 3.2 the intersection $\mathcal{Q}_\omega(\overline{v}_i) \cap \mathcal{Q}_\omega(\overline{v}_{i+1})$ is equal to $\mathcal{Q}_\omega(\overline{v}_i \cup \overline{v}_{i+1})$. This again is not a singleton, and it follows that the union

$$\mathcal{Q}_\omega(\overline{v}_0) \cup \cdots \cup \mathcal{Q}_\omega(\overline{v}_m)$$

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

Now for general $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ satisfying (2), the above argument implies that $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ can be approximated arbitrarily closely by $\boldsymbol{\mu}'$ and $\boldsymbol{\nu}'$ which cannot be separated by a point. Since maximal subsets without cutpoints are closed [5], 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 $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$, so that there exists $r > 0$ such that whenever $d(\boldsymbol{\mu}, \boldsymbol{\mu}') \leq r$ and $d(\boldsymbol{\nu}, \boldsymbol{\nu}') \leq r$, we have $d_{\mathcal{C}(S)}(\mu'_n, \nu'_n) \to \infty$ for any representative sequences. We can assume $r < d_{\mathcal{M}_\omega}(\boldsymbol{\mu}, \boldsymbol{\nu})/2$.

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 $[\frac{1}{2} r s_n, r s_n]$ for all sufficiently large n. Fix such a pair of sequences (μ'_n) and (ν'_n).

Hence $d_{\mathcal{C}(S)}(\mu'_n, \nu'_n) \to \infty$, and by Lemma 7.6 there exists a microscopic jet $J = (\overline{\sigma}, \overline{S}, \overline{\mu}', \overline{\nu}')$ such that $\overline{\mu}' \not\sim_{\overline{\sigma}} \overline{\nu}'$.

Now 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 base(μ_n) and base(ν_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 $\overline{\mu} \not\sim_{\overline{\sigma}} \overline{\nu}$.

Now $\mathcal{L}_\omega(J) = \{\boldsymbol{\iota}(J)\}$ since J is built on the main surface S. We claim that $\boldsymbol{\mu} \neq \boldsymbol{\iota}(J)$ and $\boldsymbol{\nu} \neq \boldsymbol{\iota}(J)$. 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 2.7.

Second, if v is a vertex on a tight $\mathcal{C}(S)$–geodesic g from a to b, then $v \rfloor a$ is in $\Sigma_\epsilon(a, b)$, for a uniform ϵ. This will follow by showing, for all $W \subseteq S$, that $\pi_{\mathcal{C}(W)}(v \rfloor a)$ is uniformly close to $\mathrm{hull}_W(a, b)$. If $W \not\pitchfork v$, the projections of $v \rfloor a$ and a are by definition close. If $W \pitchfork v$ 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 2.5 implies that one of $d_W(v, a)$ or $d_W(v, b)$ is uniformly bounded. If $W = S$ then v is already in $\mathrm{hull}_S(a, b)$.

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

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

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 $\overline{\mu} \not\sim_{\overline{\sigma}} \overline{\nu}$, and that $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are different from $\boldsymbol{\iota}(J)$. Hence by Theorem 7.3, $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ can be separated by $\boldsymbol{\iota}(J)$.

□

7.5. Manifolds and jets

As an application of the decomposition associated to a macroscopic jet with the gradual growth property, and the homology lemma and its corollary, we can obtain:

Theorem 7.11. *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 $\overline{W}$ of J has $\xi(\overline{W}) > 1$. Then if $E \cap \mathcal{L}_\omega(J) \neq \emptyset$, we conclude*

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

Proof. Let $\boldsymbol{q} \in \mathcal{L}_\omega(J) \cap E$. Suppose on the contrary that $E \cap \Lambda_J \neq \emptyset$. Now $\mathcal{L}_\omega(J)$ has codimension at least 2 since $\xi(\overline{W}) > 1$ (here we are using the dimension theorem from [3]) and hence it cannot separate E which has dimension $\xi(S)$, even though it does separate Λ_J from Ω_J in $\mathcal{M}_\omega$, by Theorem 7.8. We conclude that $E \setminus \mathcal{L}_\omega(J)$ is contained in Λ_J.

Now let B be a ball in E containing $\boldsymbol{q}$ in its interior. Since $\mathcal{L}_\omega(J) \cap \partial B$ has codimension at least 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 7.8, 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 7.8 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 7.8 also implies that it bounds an r-chain B' in Λ_J.

Corollary 6.8 now implies, since B is embedded and $\partial B = \partial(B' + U)$, that $B \subset B' + U$. Assuming we have chosen ϵ so that $c\epsilon < d(\boldsymbol{q}, \partial B)/2$, we find that $\boldsymbol{q}$ cannot be in U. Hence $\boldsymbol{q} \in B' \subset \Lambda_J$. This is a contradiction. □

8. Local finiteness for manifolds

Our main goal in this section is Theorem 8.8, 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 8.4, 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**

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 $\tau\colon \overline{A} \to \overline{A}'$ 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 $\xi(\overline{W}) > 1$,*

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

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 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 e.g. $\tau(a_n) \in \tau(A_n) = A'_n$.

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 of subsurfaces with $\xi(\overline{W}) > 1$, and suppose that* $\operatorname{diam}_{\mathcal{C}(W_n)}(A_n) \to_\omega \infty$. *There exists a map $\tau\colon \overline{A} \to \Sigma_\epsilon(\overline{A})$ such that*

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

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

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

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

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

$$\tau_1(\overline{x}) = \begin{cases} \overline{x} & x_\omega \notin \mathcal{Q}_\omega(\partial\overline{W}) \\ \pi_{\mathcal{Q}(\partial\overline{W})}(\overline{x}) & x_\omega \in \mathcal{Q}_\omega(\partial\overline{W}). \end{cases} \tag{8.1}$$

The notation $\pi_{\mathcal{Q}(\mu)}\colon \mathcal{M}(S) \to \mathcal{Q}(\mu)$ 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 $\pi_{\mathcal{Q}(\partial\overline{W})}\colon \mathcal{M}(\overline{S}) \to \mathcal{Q}(\partial\overline{W})$.

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

(1) $\tau_1(\overline{x})_\omega = x_\omega$,
(2) Either $\tau_1(\overline{x})_\omega$ is not in $\mathcal{Q}_\omega(\partial\overline{W})$, or $\tau_1(x_n)$ is precisely in $\mathcal{Q}(\partial W_n)$ for ω-a.e. n.
(3) $\tau_1(\overline{x}) \in \Sigma_\epsilon(\overline{A})$ 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 $\tau_1(\overline{x}) = \overline{x}$. Hence, assume that $\tau_1(\overline{x}) = \pi_{\mathcal{Q}(\partial\overline{W})}(\overline{x}) \in \mathcal{Q}(\partial\overline{W})$. By definition of Σ_ϵ, it suffices to bound the distance from $\pi_{\mathcal{C}(U)}(\tau_1(x_n))$ to $\mathrm{hull}_U(A_n)$, for every $U \subseteq S$ and ω-a.e. n.

If U is disjoint from ∂W_n, then $d_{\mathcal{C}(U)}(x_n, \tau_1(x_n))$ is uniformly bounded: this can be seen easily from the definition of the projection $\pi_{\mathcal{Q}(\partial W_n)}$ and the fact that subsurface projection maps are coarsely Lipschitz (Lemma 2.10). Hence since $x_n \in A_n \subset \Sigma_\epsilon(A_n)$, this gives the desired bound. If U intersects ∂W_n, then $d_{\mathcal{C}(U)}(\partial W_n, \tau_1(x_n))$ is again bounded by Lemma 2.10. Although we know $\mathrm{diam}_{\mathcal{C}(W_n)}(A_n) \to_\omega \infty$, we only need the fact that it is greater than $2m_0$ for ω-a.e. n, which allows us to apply Lemma 5.5 to obtain $d_{\mathcal{C}(U)}(\partial W_n, \mathrm{hull}_U(A_n)) < m_0$. Thus we conclude that (3) holds.

Now for notational simplicity let us assume that we have replaced $\overline{A}$ by $\tau_1(\overline{A})$, so that property (2) holds for each $\overline{x} \in \overline{A}$ itself (and (1) and (3) become trivial).

Recalling that $\overline{A}$ has two elements, write $\overline{A} = \{\overline{a}, \overline{b}\}$. The discussion separates into several cases.

Case 1: The simplest case is when neither a_ω nor b_ω lies in $\mathcal{Q}_\omega(\partial\overline{W})$. 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 $d_{\mathcal{C}(W_n)}(a_n, b_n) \to_\omega \infty$. Then Lemma 7.6 yields a microscopic jet, J, built from $\mathcal{C}(W_n)$-geodesics σ_n and in which $(a_n) \not\sim_{\overline{\sigma}} (b_n)$. Moreover, by our assumption that a_ω and b_ω are not in $\mathcal{Q}_\omega(\partial\overline{W})$, we know that neither of them are contained in $\mathcal{L}_\omega(J)$. It follows from Lemma 7.3 that $\mathcal{L}_\omega(J)$ separates a_ω from b_ω.

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

From now on we may assume that at least one element of $\overline{A}$, say $\overline{a}$, lies in $\mathcal{Q}(\partial\overline{W})$. Now we consider the projections of a_ω and b_ω to the factor $\mathcal{M}_\omega(\overline{W})$ of $\mathcal{Q}_\omega(\partial\overline{W})$.

Case 2a: $\pi_{\mathcal{M}_\omega(\overline{W})}(a_\omega) = \pi_{\mathcal{M}_\omega(\overline{W})}(b_\omega)$.

In this case, we can simply adjust $\overline{a}$ so that the projections, before rescaling, are a bounded distance apart. That is, recall that $\mathcal{Q}(\partial W_n)$ can be identified with $\mathcal{M}(W_n) \times \mathcal{M}(W_n^c)$, and let $a_n = (\alpha_n, \beta_n)$ and $\pi_{\mathcal{Q}(\partial W_n)}(b_n) = (\gamma_n, \delta_n)$ in this product structure. Our assumption in this case means that $\alpha_\omega = \gamma_\omega$ in $\mathcal{M}_\omega(\overline{W})$, and so we 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 possible subsurfaces. Similarly $\tau_2(\overline{a})_\omega = a_\omega$, 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 $\overline{b} \in \mathcal{Q}(\partial\overline{W})$. If it is, then the roles of $\overline{a}$ and $\overline{b}$ can be reversed.)

Case 2b: $\pi_{\mathcal{M}_\omega(\overline{W})}(a_\omega) \neq \pi_{\mathcal{M}_\omega(\overline{W})}(b_\omega)$.

In this case, consider the jet $J = (\overline{\sigma}, \overline{W}, \overline{a}, \overline{b})$, where $\sigma_n = [x_n, y_n]$ with $x_n \in \pi_{\mathcal{C}(W_n)}(a_n)$ and $y_n \in \pi_{\mathcal{C}(W_n)}(b_n)$. Writing $\overline{a} = (\overline{\alpha}, \overline{\beta})$ in the product structure of $\mathcal{Q}(\partial\overline{W})$ as above, we get $\overline{\alpha} = \overline{\iota}(J)$. That is, recall $\iota_n(J) = \iota(\sigma_n, a_n, b_n) = x_n \rfloor \pi_{\mathcal{M}(W_n)}(a_n) = x_n \rfloor \alpha_n$, and since x_n is a vertex of α_n, this is just α_n.

Since a_ω and b_ω are contained in a connected top-dimensional manifold E, we can apply Lemma 7.11 to conclude that J cannot have the gradual growth property. That is $\mathcal{L}_\omega(J) = \mathcal{Q}_\omega(\overline{\alpha} \cup \partial\overline{W})$ contains a_ω, and if σ had gradual growth then b_ω would be in Λ_J, but then Lemma 7.11 would forbid b_ω from being in E.

Since J does not have gradual 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. (For short we write $|| \cdot ||$ to denote $|| \cdot ||_{(W_n, a_n, b_n)}$.) Let $\tau_3(a_n)$ be the marking obtained by projecting a_n to $\mathcal{Q}(p_n \cup \partial W_n)$ – hence $d_{\mathcal{M}(S)}(a_n, \tau_3(a_n))$ grows sublinearly, so $\tau_3(\overline{a})_\omega = a_\omega$, and as before we can show that $\tau_3(\overline{a}) \in \Sigma_\epsilon(\overline{A})$.

If $\overline{b} \in \mathcal{Q}(\partial\overline{W})$, we can do the same thing for $\overline{b}$, obtaining $\tau_3(\overline{b})$. In σ_n we now have a sequence $x_n \cdots p_n \cdots q_n \cdots v_n \cdots u_n \cdots y_n$, with u_n playing the role of p_n and v_n playing the role of q_n, so that $\tau_3(b_n) = \pi_{\mathcal{Q}(u_n \cup \partial W_n)}(b_n)$. If $b_\omega \notin \mathcal{Q}_\omega(\partial\overline{W})$, then we simply let $\tau_3(\overline{b}) = \overline{b}$, and let $u_n = v_n = y_n$.

We claim now that $d_{\mathcal{C}(W_n)}(\tau_3(a_n), \tau_3(b_n)) = d_{\mathcal{C}(W_n)}(p_n, u_n)$ is bounded ω-a.s. – for if it were not, we could again, as in case (1), extract a microscopic jet J' from the interval sequence $[q_n, v_n]$. The points a_ω and b_ω cannot 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 it gives, 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 $\overline{b} \in \mathcal{Q}(\partial\overline{W})$, and if not it is even easier for b_ω is not even in $\mathcal{Q}_\omega(\partial\overline{W})$.

Hence, Lemma 7.3 would imply that $\mathcal{L}_\omega(J')$ separates E, and again this would contradict the assumption that $\xi(\overline{W}) > 1$.

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

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 [21]. 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 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* $\operatorname{base}(a)$ *and* $\operatorname{base}(b)$.
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 a 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. *If $d_{\mathcal{C}(W)}(a, b) > m_0$, then there exists $h \in H(a, b)$ with $D(h) = W$.*

The level l *of a geodesic $h \in H$ is the number of applications of (2) needed to descend from h to g_H.*

To finish 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 $\operatorname{diam}_{\mathcal{C}(S)}(\tau(A_n))$ is ω-a.s. bounded. Again for notational convenience we replace 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 $\operatorname{diam}_{\mathcal{C}(W_n)}(A_n) \to_\omega \infty$, 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 $\operatorname{diam}_{\mathcal{C}(S)}(A_n)$. This bounds how many domains can occur as component domains in each g_{H_n} and hence there is (the same) bound on how many $\overline{W}$ exist with W_n a component domain of g_{H_n} for ω-a.e. n. (This is the general fact that the ultraproduct of a sequence of sets X_n of finite size k has size k.) For each such $\overline{W}$ successively, use Lemma 8.2 again to find $\tau(\overline{A})$ such that $\operatorname{diam}_{\mathcal{C}(W_n)}(\tau(\overline{A})_n)$ 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 $\operatorname{diam}_{\mathcal{C}(\overline{U})}$ 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 sequences $\overline{W}$ of component domains for the main geodesic sequence.

This procedure repeats $\xi(S)$ times: At each step we have bounds on the lengths of all geodesics that occur at level at most l in the hierarchy, and hence on the *number* of geodesics at level $l+1$. 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 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 $\operatorname{diam}_{\mathcal{C}(W_n)}(B_n)$ 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 $\operatorname{diam}_{\mathcal{C}(W_n)}(A_n) \to_\omega \infty$. 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 the same holds for $\overline{a}$. Then the same argument as Case 1 of Lemma 8.2 shows that $\operatorname{diam}_{\mathcal{C}(W_n)}(a_n, b_n)$ is bounded. Since $\operatorname{diam}_{\mathcal{C}(W_n)}(B_n)$ 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 $\overline{a} \in \mathcal{Q}(\partial \overline{W})$. We can apply the argument of Cases 2a and 2b of Lemma 8.2 to show 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 $\mathrm{diam}_{\mathcal{C}(W_n)}(B_n)$ is assumed bounded we are done.

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

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

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

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

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

If $a_\omega \notin \mathcal{M}_\omega(\partial \overline{W})$ we leave it unchanged, but if $\overline{a} \in \mathcal{M}(\partial \overline{W})$, we argue as in Case 2b of Lemma 8.2, first to show that a jet from $\overline{a}$ to $\overline{b}_1$ cannot have gradual growth, and then to modify $\overline{a}$: 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 $\tau_4(\overline{a}) = (p_n \cup \partial W_n) \rfloor a_n$.

Unlike Case 2b of 8.2, we do not attempt to modify $\overline{b}_1$. 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. a_n and b_{1n} project to opposite sides of σ_n so $\overline{a} \not\sim_{\overline{\sigma}} \overline{b}_1$. 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 $\mathcal{M}_\omega(\overline{W})$. Since $d_{\mathcal{C}(W_n)}(b_{1n}, b_{2n})$ is ω-a.s. bounded, we also have $\overline{a} \not\sim_{\overline{\sigma}} \overline{b}_2$, 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 $\overline{A} = \overline{B} \cup \{\overline{a}\}$. Now we finish the proof as we did before: we repeatedly apply this result, bounding first the lengths of the $\mathcal{C}(S)$-geodesics between elements of $\overline{A}$, and then lengths of geodesics in subsurfaces which are complementary components of the vertices of the $\mathcal{C}(S)$-geodesics, and so on until only domains of complexity 1 are left with unbounded diameters. □

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 8.4. *If $\boldsymbol{A}$ is a finite subset of a connected top-dimensional manifold E in $\mathcal{M}_\omega(S)$, then $\Sigma(\boldsymbol{A})$ is contained in a finite union of cubes.*

Proof. From Theorem 8.1 we may assume that $\boldsymbol{A}$ is represented by (A_n) such that, for ω-a.e. n, $\operatorname{diam}_{\mathcal{C}(W)}(A_n)$ is bounded by some fixed k_0 whenever $\xi(W) > 1$. Let us consider an arbitrary $A \subset \mathcal{M}(S)$, of fixed cardinality $\#A = \#\boldsymbol{A}$, satisfying this condition.

Fix ϵ large enough for Proposition 5.2 (on retractions of Σ-hulls) to apply. We may assume that $k_0 > \max\{3(m_0+4), 2m_0+\epsilon\}$, for later use. 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. Let the projections $(\pi_W(\mu))$ play the role of (x_W) in Section 4, so that the consistency conditions are satisfied as in Theorem 4.1 with any $c_1 \geq m_0$. Now define $\ll_k$ and $\prec_k$ as in section 4, that is, $V \ll_k W$ iff $V \pitchfork \partial W$ and

$$d_V(\mu, \partial W) > k(c_1 + 4), \tag{8.2}$$

whereas $V \prec_k W$ iff $V \ll_k W$ and $V \pitchfork W$. We choose c_1 so that $k_0 + \epsilon = 3(c_1 + 4)$. In particular $\mathcal{F}_3(a) = \{W : W \ll_3 a\}$ is the set

$$\{W \subset S : d_{\mathcal{C}(W)}(\mu, a) > k_0 + \epsilon\}.$$

Lemma 4.5 now tells us that $\prec_2$ is a partial order on $\mathcal{F}_3(a)$. Moreover, by Lemma 4.4, any V, W in $\mathcal{F}_3(a)$ such that $V \pitchfork W$ are in fact $\prec_2$-ordered.

$\mathcal{F}_3(a)$ 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, so let $\mathcal{U} = \mathcal{U}(\mu, a)$ be the subset of $\mathcal{V}$ consisting of elements maximal with respect to containment. Hence $\mathcal{U}$ enumerates the components of an essential subsurface, 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

$$G(A, \mathcal{U}, a) = \{\pi_{\mathcal{M}(\mathcal{U}^c)}(a)\} \times \prod_{U \in \mathcal{U}} \Sigma_{\epsilon'}(\pi_{\mathcal{M}(U)}(A)) \tag{8.3}$$

Here U varies over the connected 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 $\partial\mathcal{U}$. 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 $W \pitchfork \partial U$. Since $U \ll_3 a$, by Lemma 4.4 (2) we have $W \ll_3 a$, so that $W \in \mathcal{F}_3(a)$.

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 $Z \in \mathcal{F}_3(a)$ 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 $d_{\mathcal{C}(W)}(\mu, a) > 3(c_1 + 4)$ – then $W \in \mathcal{F}_3(a)$. 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 in $\mathcal{U}$ we have $\pi_W(\mu) \in \mathcal{N}_\epsilon(\mathrm{hull}_W(A))$. But $\mathrm{hull}_W(A)$ is within uniformly bounded distance of $\mathrm{hull}_W(\pi_{\mathcal{M}(\mathcal{U})}(A))$ by the coarse composition property of projections (Lemma 2.11). 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 (8.3). In other words we have proven:

Lemma 8.5. *$\Sigma_\epsilon(A)$ is contained in the union of sets*

$$G'(A, \mathcal{U}, a) = \mathcal{N}_{k_2}(G(A, \mathcal{U}, a))$$

where k_2 is a uniform bound, a is a fixed point in A, and $\mathcal{U} = \mathcal{U}(\mu, a)$ for $\mu \in \Sigma_\epsilon(A)$.

Now we would like to bound the number of $\mathcal{U}$ that can occur in this way.

If W occurs as a component of $\mathcal{U}$ for some μ, then $d_{\mathcal{M}(S)}(\mu,a) > k_0+\epsilon$. Since $\pi_W(\mu) \in \mathcal{N}_\epsilon(\mathrm{hull}_W(A))$,

$$\mathrm{diam}_{\mathcal{C}(W)}(A) > k_0.$$

By our assumptions about A, this means $\xi(W) \leq 1$. The following counting argument gives us a method for bounding the number of such surfaces:

Lemma 8.6. *If $d_{\mathcal{C}(W)}(a,b) \leq k$ for all subsurfaces W of $\xi(W) \geq t$, then the hierarchy $H(a,b)$ contains at most $O(k^{\xi(S)-t})$ geodesics of $\xi = t-1$.*

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 at least $s+1$. Hence the number of $\xi = s$ geodesics is bounded by the number of $\xi > s$ geodesics times the length bound on those geodesics. ☐

Now define

$$\mathcal{S}_1 = \{U \subset S : \xi(U) = 1 \text{ and } \mathrm{diam}_U(A) > k_0\},$$
$$\mathcal{S}_0 = \{U \subset S : \xi(U) = 0 \text{ and } \mathrm{diam}_U(A) > k_0\}.$$

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 Lemma 8.6 directly gives a bound on $\#\mathcal{S}_1$.

There is no uniform bound for $\#\mathcal{S}_0$, but we will see that those annuli U that arise in the minimal sets, $\mathcal{U}$, of the above construction form a restricted subset of $\mathcal{S}_0$, whose cardinality we can uniformly control. The following lemma is the main reason for this:

Lemma 8.7. *Suppose $\mathrm{diam}_{\mathcal{C}(W)}(A) > k_0$. If U is an annulus that appears as a component of a minimal set $\mathcal{U}(\mu,a)$ for some $\mu \in \Sigma_\epsilon(A)$ and $a \in A$, and $U \Subset W$, then $d_{\mathcal{C}(W)}(\partial U, A) \leq k_3$ for k_3 depending on k_0 and $\#A$.*

Proof. Define $\mathcal{F}_3(a)$ using μ as before. There are two possibilities for W:

(1) $W \notin \mathcal{F}_3(a)$: This means that $d_{\mathcal{C}(W)}(\mu,a) \leq k_0+\epsilon$. 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 2.5. It follows that $d_{\mathcal{C}(W)}(\partial U,a) \leq d_{\mathcal{C}(W)}(\mu,a) - 1 \leq k_0+\epsilon-1$.

(2) $W \in \mathcal{F}_3(a)$. In this case, since $U \Subset W$, W can't be $\prec_2$-minimal in $\mathcal{F}_3(a)$, because then W would have been included in $\mathcal{U}$ instead of U. Hence there is some element Y in $\mathcal{V}$ 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 $\pi_Y(\mu) \in \mathcal{N}_\epsilon(\mathrm{hull}_Y(A))$. 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 $d_{\mathcal{C}(W)}(\partial U, b) < m_0 + 1$.

□

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 $b \in A$ and a geodesic h in $H(a, b)$ with domain W such that $\xi(W) \geq 1$ and U is a component domain of h. By Lemma 8.7, ∂U is within k_3 of $\pi_W(c)$ for some $c \in A$. This restricts it, for each c, to a segment of length at most $2k_3$ 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 8.6, this gives us an a priori 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 $\overline{A}$. 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 $\Sigma(\boldsymbol{A})$ is contained in a finite union of asymptotic cones of sequences $G(A_n, \mathcal{U}_n, a_n)$, which by (8.3) must be sets of the form

$$\{x_\omega\} \times \prod_{\overline{U} \in |\overline{\mathcal{U}}|} T_{\overline{U}}$$

in $\mathcal{Q}_\omega(\partial\overline{\mathcal{U}})$, where $x_\omega \in \mathcal{M}_\omega(\overline{\mathcal{V}}^c)$, each component $\overline{U}$ of $\overline{\mathcal{U}}$ has $\xi \leq 1$, and each $T_{\overline{U}}$ is the convex hull of a finite set in the $\mathbb{R}$-tree $\mathcal{M}_\omega(\overline{U})$. Hence each $T_{\overline{U}}$ is a finite tree, so after breaking each tree into a finite union of segments, we obtain the desired finite union of cubes. □

8.3. Local finiteness

The main application of Theorem 8.4 is the following:

Theorem 8.8. *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 $B \subset int(B')$ where B' is a larger ball. Triangulate $\partial B'$ with simplices of diameter smaller than r, where r will be chosen shortly. Let $f_0 \colon B' \to E$ be the identity, and let $f_1 \colon \partial B' \to \mathcal{M}_\omega(S)$ be a Σ-compatible map with respect to the triangulation (it exists by Lemma 6.3), which agrees with f_0 on the 0-skeleton. By Lemma 6.1 we have $d(f_0, f_1) < Cr$ for a uniform C, and by Lemma 6.5, there exists a homotopy $h\colon \partial B' \times [0,1] \to \mathcal{M}_\omega(S)$ with track diameters at most $C'r$.

Choose r small enough that $Cr < \frac{1}{2}d(B, \partial B')$. 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 $F\colon B' \to \mathcal{M}_\omega(S)$ 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 6.8, 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 8.4, 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 $\boldsymbol{x} \in \mathcal{M}_\omega(S)$, by considering the *germs* of such manifolds, and using the Local Finiteness Theorem 8.8 to relate this to the *complex of orthants* through $\boldsymbol{x}$. The main result is Corollary 9.6, which states that germs of Dehn twist flats passing through $\boldsymbol{x}$ admit a purely topological characterization; a *Dehn twist flat* in $\mathcal{M}_\omega(S)$ is, by definition, an ultralimit of a sequence of Dehn twist flats (see Section 3.1). 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.

Poset of Germs

For the remainder of this section we fix a basepoint $\boldsymbol{x} \in \mathcal{M}_\omega$. We will consider the set of closed subsets of $\mathcal{M}_\omega$ containing $\boldsymbol{x}$, modulo the equivalence $C \sim C'$ if there exists an open neighborhood U of $\boldsymbol{x}$ such that $C \cap U = C' \cap U$. The equivalence classes are called *germs through* $\boldsymbol{x}$, and we let $\gamma(C)$ denote the germ of C through $\boldsymbol{x}$. Note that finite intersection and union yield well-defined operations on the set of germs, and the subset relation is well-defined as well. We let $\mathcal{G}$ denote the poset of germs at $\boldsymbol{x}$; this is a *lattice* (i.e., least upper bounds and greatest lower bounds exist for all pairs $C, C' \in \mathcal{G}$, namely $C \cup C'$ and $C \cap C'$).

Our goal for the next subsection is to study the set of germs of cubes in $\mathcal{M}_\omega(S)$ for which $\boldsymbol{x}$ is a *corner*. These germs will be called *orthants at* $\boldsymbol{x}$.

Structure of orthants

A cube with distinguished corner is a cube $C = C(\mu, W, r)$ for which each r_i has a distinguished endpoint $r_i(0)$. The corner is a 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 $C(\overline{\mu}, \overline{W}, \overline{r})$, we obtain a cube C^ω with corner κ^ω in the limit. We now define an *orthant* at $\boldsymbol{x}$ to be the germ $\gamma(C^\omega)$ for a cube C^ω with distinguished corner equal to $\boldsymbol{x}$. If an orthant O can be defined by a sequence $\overline{W}$ where all components are annuli, then we say that O is a Dehn twist orthant. The germ of a Dehn twist flat in $\mathcal{M}_\omega(S)$ is the union of Dehn twist orthants.

We recall that, up to ultraproduct equivalence, the sequence W^n of subsurfaces can be identified with a finite set $\overline{W}_1, \ldots, \overline{W}_k$ of sequences of connected subsurfaces such that $W_1^n, \ldots, W_k^n$ are the components of W^n for ω-a.e. n.

The *faces* of an orthant are the orthants obtained from the faces of the cube which meet $\boldsymbol{x}$. More precisely, if $O = \gamma(C^\omega(\overline{\mu}, \overline{W}, \overline{r}))$, let $\overline{W}'$ be a collection of components of $\overline{W}$. (Equivalently each W'^n is a collection of components of W^n.) Let μ'^n be markings in $(W'^n)^c$ obtained as the union of μ^n in $(W^n)^c$ and $r_i^n(0)$ for all components W_i^n of $W^n \setminus W'^n$. Let r'^n be the components of r^n supported in the components of W'^n. Then $C(\overline{\mu}', \overline{W}', \overline{r}')$ is a sequence of faces of $C(\overline{\mu}, \overline{W}, \overline{r})$ and their limit $C^\omega(\overline{\mu}', \overline{W}', \overline{r}')$ is a cube in the cone with corner still at $\boldsymbol{x}$. Hence its germ is an orthant at $\boldsymbol{x}$.

Let us define $\mathcal{O}$ to be the poset of all *nontrivial orthants at* $\boldsymbol{x}$, i.e. all orthants except for the singleton $\{\boldsymbol{x}\}$.

One complication that arises is that an orthant has many quite different representations by sequences of cubes, so that equality and the face relation are not trivial to detect. Nevertheless we will establish the following:

Lemma 9.1. *The poset $\mathcal{O}$ is isomorphic to the poset of simplices of a simplicial complex $\mathcal{K}$ of dimension $\xi(S) - 1$. Moreover, $\mathcal{K}$ is a flag complex.*

Recall that a simplicial complex Y 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 Y.

Proof. Recall from Section 3.3 that the dimension of an asymptotic cube $C^\omega(\overline{\mu}, \overline{W}, \overline{r})$ is equal to the number of components $\overline{r}_i$ of $\overline{r}$ for which the length of the limiting segment r_i^ω is positive (equivalently the lengths $l(r_i^n)$ grow linearly). We will call a germ of a cube of dimension k a k-orthant. Having excluded the unique 0-orthant $\{\boldsymbol{x}\}$ from $\mathcal{O}$, we will let 1-orthants be the vertices of our complex, and in general $k+1$-orthants will be k-simplices.

Since $\mathcal{O}$ is closed under extraction of positive-dimension faces, to obtain the structure of a simplicial complex it remains to show that the intersection of two orthants is either a common face, or the 0-orthant, in which case we will say they are disjoint in $\mathcal{K}$.

Thus consider two orthants $O_l = \gamma(C_l^\omega) \in \mathcal{O}$ (for $l = 1, 2$) where $C_1^\omega = C^\omega(\overline{\mu}, \overline{W}, \overline{r})$ and $C_2^\omega = C^\omega(\overline{\nu}, \overline{V}, \overline{s})$. To understand $C_1^\omega \cap C_2^\omega$, recall from Section 3.3 that this intersection is either empty or equal to the common ultralimit of the junctures of the approximating cubes. Since both cubes contain $\boldsymbol{x}$, the empty case cannot occur, and we are left to study the junctures.

Lemmas 3.3 and 3.5 show that the junctures of the approximating cubes $C_1^n = C(\mu^n, W^n, r^n)$ and $C_2^n = C(\nu^n, V^n, s^n)$ are themselves subcubes $C_{12}^n \subset C_1^n$ and $C_{21}^n \subset C_2^n$, which have the form

$$C_{12}^n = C(\mu^n, W^n, r'^n),$$

where r'^n denotes a collection of subintervals (or points) of the segments in r^n, and similarly

$$C_{21}^n = C(\nu^n, V^n, s'^n),$$

where s'^n are collections of subintervals or points of s^n.

Lemma 3.5 also states that for each n, after choosing an appropriate ordering of the components as in Lemma 3.3, two components $r_i'^n$ and $s_i'^n$ either both have 0 length or (the same) positive length, and the latter occurs only if the corresponding subsurfaces W_i^n and V_i^n have nontrivial essential intersection, $W_i^n \pitchfork V_i^n \neq \emptyset$. For each i, this either occurs or fails for ω-a.e. n, and accordingly we will say that $\overline{W}_i \pitchfork \overline{V}_i$ is nonempty or empty, following our usual ultraproduct convention.

We can and do parametrize each $r_i'^n$ in such a way that the corner $\kappa(C_{12}^n) = \{\mu^n\} \times \prod r_i'^n(0)$ is the nearest corner to $\kappa(C_1^n)$ (and similarly for $s_i'^n$. Because the limiting cube contains $\boldsymbol{x}$, it must be that the ultralimit $\kappa^\omega(C_{12}^\omega)$ of $\kappa(C_{12}^n)$ equals $\boldsymbol{x}$, and hence for each i $r_i'^\omega(0) = r_i^\omega(0)$, or equivalently the subsegments from $r_i^n(0)$ to $r_i'^n(0)$ 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 $l(\overline{r}_i')$ grow linearly. That is – we can replace the sublinearly growing segments with the basepoints of the original segments, and the linearly growing segments $\overline{r}_i'$ with the 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 $\boldsymbol{x}$, and hence its germ is equal to $O_1 \cap O_2$. Similarly $O_1 \cap O_2$ can be identified with a face of O_2.

We have shown that the intersection of two orthants is an orthant which is equal to a common face of the two. This completes the proof that $\mathcal{O}$ has the structure of a simplicial complex. The dimension is $\xi(S) - 1$ because $\xi(S)$ is the maximal dimension of a cube.

The structure of junctures also gives us a way to characterize when two sequences of cubes give rise to the same orthant. Letting O_1 and O_2 be as before, let us first assume that each r_i^ω and s_j^ω has positive length – otherwise we can restrict to subfaces of the cubes which have the same limit as the whole cubes. For equality, the dimensions must match so $\overline{W}$ and $\overline{V}$ must have the same cardinality k. Now $O_1 = O_2$ if and only if $O_1 \cap O_2 = O_1 = O_2$, which means that the junctures of the approximating cubes have to be large enough to give the same orthant in the limit. In particular $\overline{V}_i \Cap \overline{W}_i \neq \emptyset$ for each $i = 1, \dots, k$, and the juncture subsegments $r_i'^n$ and $s_i'^n$ must have linearly growing lengths.

When $\overline{V}_i = \overline{W}_i$, this means that r_i^ω and s_i^ω intersect on positive-length initial segments, corresponding to common subsegments $\overline{r}_i'$ and $\overline{s}_i'$ which are separated from the initial points by sublinearly growing segments. In the remaining cases $\overline{V}_i$ and $\overline{W}_i$ overlap on annuli $\overline{U}_i$ (which may be equal to one of them), and r_i^ω and s_i^ω have positive-length initial segments which are *twist segments*, i.e. are equal to limits of Dehn-twist lines along $\overline{U}_i$. 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 $r_i'^n(0)$ (and similarly for s_i^n and $s_i'^n$).

We also observe the following consequence of our discussion: suppose $O = \gamma(C(\overline{\mu}, \overline{W}, \overline{r}))$ is an orthant and suppose one of the r_i^ω has an initial Dehn-twist segment. Then if $\overline{W}_i$ is not already an annulus it can be replaced by one. That is, suppose $\overline{U} \subset \overline{W}_i$ is an annulus sequence and $\overline{s}$ a sequence of Dehn twist segments in $\mathcal{TM}(\overline{U})$ which, after the embedding $\mathcal{TM}(\overline{U}) \to \mathcal{TM}(\overline{W}_i)$ of Section 3.2, gives subsegments of $\overline{r}_i$ whose limit is an initial segment of r_i^ω. Then we can replace $\overline{W}_i$ by $\overline{U}$ and replace $\overline{r}_i$ by $\overline{s}$, and the resulting cube has the same germ. Note that this replacement may remove an initial segment sequence of sublinearly growing length from $\overline{r}_i$.

Now we are ready to show that $\mathcal{K}$ is a flag complex.

Let $O_1, \dots, O_k$ be 1-orthants which are the vertices of a complete graph in $\mathcal{O}$. Then for each i, j there is a 2-orthant O_{ij} whose faces are O_i and O_j.

Each O_i is the germ of $C^\omega(\overline{\mu}_i, \overline{W}_i, \overline{r}_i)$, where $\overline{W}_i$ has a single component. By the last observation above, we may assume that if r_i^ω has an initial twist segment then in fact $\overline{W}_i$ is an annulus sequence.

Now represent O_{ij} as $\gamma(C^\omega(\overline{\nu}_{ij}, \overline{V}_{ij}, \overline{s}_{ij}))$, where $\overline{V}_{ij}$ has two components. Since O_i is a face of O_{ij}, it must be equal to the germ of a sequence of faces of $C^n(\overline{\nu}_{ij}, \overline{V}_{ij}, \overline{s}_{ij})$, associated to one of the components of $\overline{V}_{ij}$. By the above discussion on when orthants are equal, this means that the essential intersection of $\overline{W}_i$ and one of the components of $\overline{V}_{ij}$ is nontrivial. However, the case where $\overline{W}_i \Cap \overline{V}_{ij}$ is a proper annulus of $\overline{W}_i$ *cannot occur*, because in that case there is an initial twist segment in r_i^ω, and so $\overline{W}_i$ is already an annulus sequence. The same holds for j. Hence $\overline{W}_i$ and $\overline{W}_j$ are *contained in* $\overline{V}_{ij}$, which means that they are disjoint from each other – recall this means that W_i^n and W_j^n are disjoint for n in a set I_{ij} of full ω-measure.

The intersection $\cap I_{ij}$ over all (i,j) still has full ω-meaure, so we conclude that $W_1^n, \ldots, W_k^n$ are pairwise disjoint for ω-a.e. n, and we let $\overline{W} = \overline{W}_1 \cup \cdots \cup \overline{W}_k$. Let $\overline{\sigma}$ be a marking sequence on $\overline{W}^c$ defined as the projection $\pi_{\mathcal{M}(\overline{W}^c)}(\overline{x})$, and let $\overline{r} = (\overline{r}_1, \ldots, \overline{r}_k)$. Then we obtain an orthant

$$O = \gamma(C^\omega(\overline{\sigma}, \overline{W}, \overline{r})).$$

We need to check that the corner of O, namely the limit of $\overline{\kappa} = \{\overline{\sigma}\} \times \prod \overline{r}_i(0)$, equals $\boldsymbol{x}$. But this is a consequence of the quasidistance formula: In the quasidistance formula for $d(\kappa^n, x^n)$, we separate the terms $\{\!\{d_Z(\kappa^n, x^n)\}\!\}_K$ according to whether $Z \Subset (W^n)^c$, $Z \Subset W^n$, or $Z \pitchfork \partial W^n$. The first type of term adds up to an estimate of $d_{\mathcal{M}(W^n)^c}(\sigma^n, x^n)$, which by definition is bounded. The second type adds up to estimate $\sum_i d_{\mathcal{M}(W_i^n)}(r_i^n(0), x^n)$ which grows sublinearly since the corner of each C_i^ω is $\boldsymbol{x}$. The third type is estimated, termwise, by $d_Z(\partial W^n, x^n)$, which sum up to estimate $d(x^n, \mathcal{Q}(\partial W^n))$, by Lemma 3.1. This again grows sublinearly since $\boldsymbol{x} \in \cap_i \mathcal{Q}_\omega(\partial \overline{W}_i) = \mathcal{Q}_\omega(\partial \overline{W})$. We conclude that $d(\kappa^n, x^n)$ grows sublinearly, so $\kappa^\omega = x^\omega$.

This tells us that $O \in \mathcal{O}$. Now it is clear that O_i are the vertices (i.e. 1-orthant faces) of O, since they are defined by the same surfaces and segments. This completes the proof. □

Applying local finiteness

Let us take $\widehat{\mathcal{O}} \subset \mathcal{G}$ to be the sublattice generated by $\mathcal{O}$ (i.e. by finite unions and intersections). Now consider the subset of $\mathcal{G}$ consisting of germs of submanifolds of $\mathcal{M}_\omega$ of dimension $\xi = \xi(S)$ which contain $\boldsymbol{x}$, and let $\mathcal{F} \subset \mathcal{G}$ denote the sublattice it generates.

The manifold local finiteness theorem, Theorem 8.8, implies:

Lemma 9.2. *$\mathcal{F} \subset \widehat{\mathcal{O}}$.*

Proof. Let M be a manifold of dimension $\xi = \xi(S)$ passing through $\boldsymbol{x}$. Theorem 8.8 states that there is a neighborhood U of $\boldsymbol{x}$ such that $M \cap U$ is contained in a finite union of cubes. After subdivision and possibly shrinking U we may assume that $\boldsymbol{x}$ is a corner of each of these cubes. Now invariance of domain, together with the fact that distinct cubes with a common corner intersect only along their boundaries (as in Lemma 9.1), implies that if M meets the interior of a ξ-dimensional cube then it contains the entire intersection of this cube with some neighborhood. Hence M also contains the intersection with some neighborhood of the closure of each such orthant. Finally if $M \cap U$ meets any cube P of lower dimension then (again by invariance of domain) it must meet the interior of a cube Q of which P is a face. It follows by the above that the intersection of P with a sufficiently small neighborhood is contained in 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.

Let $\overline{W} = (\overline{W}_1, \ldots, \overline{W}_\xi)$ be a decomposition associated to a top dimensional orthant O. In each $\overline{W}_i$ we have a ray $\boldsymbol{r}_i$ in the associated $\mathbb{R}$-tree $\mathcal{M}_\omega(\overline{W}_i)$. Actually we only need to consider the *germ* of a ray, but we will still denote it $\boldsymbol{r}_i$. A component $\overline{W}_i$ is called a *boundary annulus* if it is an annulus homotopic to the boundary of another $\overline{W}_j$ (necessarily of complexity 1) . Let $b(\overline{W})$ denote the number of boundary annuli. Note that $b(\overline{W}) = 0$ if all components of $\overline{W}$ are annuli, and is positive otherwise. If O' is an orthant meeting O along a codimension 1 face, then O' has a decomposition $\overline{W}'$ obtained from $\overline{W}$ by changing only one component $\overline{W}_j$, or the ray germ $\boldsymbol{r}_j$ in $\overline{W}_j$. If $\overline{W}_j$ is of complexity 1 then there are infinitely many choices for a different ray germ $\boldsymbol{r}'_j$, and hence infinitely many orthants adjacent to O along this face. If $\overline{W}_j$ is a non-boundary annulus then there are infinitely many annuli W'_j which can replace it. However, if $\overline{W}_j$ is a boundary annulus then the only change we can make is to replace $\boldsymbol{r}_j$ by the unique opposite ray germ $-\boldsymbol{r}_j$ in the same annulus complex.

We conclude that, along each of the $b(\overline{W})$ 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 number b of boundary annuli (in fact the same annuli) and for all the corresponding faces the unique neighbors 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 9.3. *Lunes are precisely the minimal ξ-dimensional elements of $\mathcal{F}$.*

Proof. Let L be a lune of rank b. 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

$$\boldsymbol{r}_1^{j_1}, \ldots, \boldsymbol{r}_b^{j_b}, \boldsymbol{r}_{b+1}^1, \ldots, \boldsymbol{r}_\xi^1,$$

where $\overline{W}_1, \ldots, \overline{W}_b$ are the boundary annuli in a decomposition $\overline{W}$, $j_i \in \{1, 2\}$, and the ray germs $\boldsymbol{r}_i^1$, $\boldsymbol{r}_i^2$ are opposite pairs for $i \in \{1, \ldots, b\}$.

For each $i \in \{b+1, \ldots, \xi\}$, 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 $\boldsymbol{r}_i^1$ as a ray germ in $\mathcal{M}_\omega(\overline{V}_i)$ via the natural embedding $\mathcal{M}_\omega(\overline{W}_i) \to \mathcal{M}_\omega(\overline{V}_i)$.

Let $\overline{W}[i]$ be the subsurface sequence obtained from $\overline{W}$ by replacing $\overline{W}_i$ by $\overline{V}_i$.

Let $\boldsymbol{r}_i^2$ and $\boldsymbol{s}_i^2$ be two ray germs in $\mathcal{M}_\omega(\overline{V}_i)$ which share a basepoint with $\boldsymbol{r}_i^1$ but are distinct from it and each other. If $\overline{V}_i \neq \overline{W}_i$, take care to choose $\boldsymbol{r}_i^2$ to be the opposite ray to $\boldsymbol{r}_i^1$ in the annulus line $\mathcal{M}_\omega(\overline{W}_i)$. Denote $\boldsymbol{s}_i^1 = \boldsymbol{r}_i^1$.

Now for each $i \in \{b+1, \ldots, \xi\}$ and each tuple $\hat{\jmath} = (j_1, \ldots, j_\xi) \in \{1,2\}^\xi$, consider the orthant $O[i](\hat{\jmath})$ formed from $\overline{W}[i]$ and the ray germs

$$\boldsymbol{r}_1^{j_1}, \ldots, \boldsymbol{r}_b^{j_b}, \ldots, \boldsymbol{s}_i^{j_i}, \ldots$$

in other words, we use $\boldsymbol{r}_k^{j_k}$ for all k except i, where we use $\boldsymbol{s}_i^{j_i}$. Let

$$M[i] = \bigcup_{\hat{\jmath}} O[i](\hat{\jmath}).$$

This is a manifold germ, and our lune L is the intersection

$$L = M[b+1] \cap \cdots \cap M[\xi].$$

This shows that L is in $\mathcal{F}$. Since any ξ-dimensional element of $\mathcal{F}$ contained in L must contain a top dimensional orthant $O \subset L$, the discussion above shows that L is minimal. Thus lunes are minimal ξ-dimensional elements of $\mathcal{F}$.

Now let C be a minimal ξ-dimensional element of $\mathcal{F}$. Then C must contain a ξ-dimensional orthant O, by Lemma 9.2. By the discussion above, C must contain the lune L determined by O, and by the minimality of C, we have $C = L$. □

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 *$\alpha$ 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. Our next goal is a characterization of the rank of lunes as a function on $\mathcal{L}$:

Lemma 9.4. *The rank is the unique function $f : \mathcal{L} \to \{0, \ldots, \xi\}$ with the following property. If $b \in \{0, \ldots, \xi\}$ and C is a lune with $f(C) \leq b$, then $f(C) = b$ if and only if C is part of a cycle*

$$\sum_{i=1}^{2^{\xi-b}} C_i,$$

where $f(C_i) \leq b$ *for all* i.

To prove this, we will need a lemma about flag complexes:

Lemma 9.5. *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 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 9.4. We will refer to the property stated in the lemma as *Property S.*

We first show that the rank function has Property S.

Suppose $b \in \{0, \ldots, \xi\}$, $C \in \mathcal{L}$, $\mathrm{rank}(C) \leq b$, and C is part of a cycle $\sum_{i=1}^{2^{\xi-b}} C_i$ where $\mathrm{rank}(C_i) \leq b$ for all i.

Since $\mathrm{rank}(C_i) \leq b$ for all i, each C_i is composed of $2^{\mathrm{rank}(C_i)} \leq 2^b$ orthants; nonzero cycles require at least 2^ξ orthants by Lemma 9.5, which implies that $\mathrm{rank}(C_i) = b$ 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 of C and C_j implies that $C = C_j$, and hence $\mathrm{rank}(C) = \mathrm{rank}(C_j) = b$.

The converse implication follows from the observation made earlier, that a lune of rank b is part of a cycle consisting of $2^{\xi-b}$ lunes of rank b. Thus rank has Property S.

Now suppose $f : \mathcal{L} \to \{0, \ldots, \xi\}$ has Property S, but is not equal to rank. Let b be the maximum of the integers $\bar{b} \in \{0, \ldots, \xi\}$ such that $f^{-1}(\bar{b}) \neq \mathrm{rank}^{-1}(\bar{b})$.

Suppose C is a lune of rank b. Then C belongs to a cycle $\sum_{i=1}^{2^{\xi-b}} C_i$ where $\mathrm{rank}(C_i) = b$. Then $f(C_i) \leq b$ by the choice of b. Hence by Property S, we get $f(C) = b$. Thus $\mathrm{rank}^{-1}(b) \subset f^{-1}(b)$.

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 $\mathrm{rank}(C_i) \leq b$ for all i, and by Lemma 9.5 we get $\mathrm{rank}\, C_i = b$ for all i. We conclude as above that $C = C_j$ for some j, and hence $\mathrm{rank}(C) = b$. Thus $f^{-1}(b) \subset \mathrm{rank}^{-1}(b)$. This contradicts the choice of b. □

Corollary 9.6. *There is a topological characterization of Dehn twist orthant germs, and Dehn twist flat germs in $\mathcal{M}_\omega(S)$.*

Proof. The lattice $\mathcal{F}$ is defined purely topologically, as is the dimension function on $\mathcal{F}$. Therefore Lemmas 9.3 and 9.4 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 characterized as intersections of Dehn twist ξ-orthants.

It remains to characterize Dehn-twist flat germs – the configurations of 2^ξ orthants associated to a Dehn-twist flat through $\boldsymbol{x}$. This boils down to characterizing when two Dehn-twist rays α and α' are opposites.

Consider a lune of rank $b > 0$. The corresponding decomposition $\overline{W}$ contains b boundary annuli, and the corresponding Dehn-twist rays appear in antipodal pairs and 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 are characterizable in terms of this simplicial structure.

Now, any Dehn-twist ray, i.e. a 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 opposite pairs of Dehn twist rays may be characterized topologically as those which may be embedded as antipodal rays in a triangulated lune boundary sphere. □

We conclude the section with this observation:

Lemma 9.7. *Suppose $E \subset \mathcal{M}_\omega(S)$ is a a connected top-dimensional manifold, and that for every $\boldsymbol{x} \in E$, the germ of E at $\boldsymbol{x}$ 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. Note that if $\boldsymbol{x} \in \mathcal{M}_\omega(S)$, and E' and E'' are both Dehn twist flats passing through $\boldsymbol{x}$, then they either have the same germ at $\boldsymbol{x}$, or the intersection of their germs has dimension strictly less than ξ at $\boldsymbol{x}$.

Now pick a Dehn twist flat $E' \subset \mathcal{M}_\omega$, such that $E \cap E'$ has nonempty interior in E. Let $U \subset E$ be the interior of $E \cap E'$ in E, and suppose $\boldsymbol{x} \in E$ lies in the closure of U in $\mathcal{M}_\omega$. Since E' is a closed subset of $\mathcal{M}_\omega(S)$, we have $\boldsymbol{x} \in E \cap E'$. By the definition of U, the germ of $E \cap E'$ at $\boldsymbol{x}$ has dimension ξ, and therefore by remark above, we conclude that $\boldsymbol{x} \in U$. 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 $E \cap E'$ will be open and closed in E', and hence $E \cap E' = E'$. □

10. Finishing the proofs

We are now ready to prove our main theorems on quasi-isometric rigidity and classification. The proof will follow the general sketch from the introduction.

10.1. Preservation of asymptotic Dehn twist flats

Theorem 10.1. *If $\xi(S) \geq 2$, any homeomorphism $f \colon \mathcal{M}_\omega(S) \to \mathcal{M}_\omega(S)$ permutes the Dehn twist flats in $\mathcal{M}_\omega(S)$.*

Proof. By Corollary 9.6, 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 9.7 therefore implies that $f(E)$ is itself a Dehn twist flat. □

10.2. Coarse preservation of Dehn twist flats

Theorem 10.2. *If $\xi(S) \geq 2$, then given $K \geq 1$ and $C \geq 0$ there exists A such that, if $f \colon \mathcal{MCG}(S) \to \mathcal{MCG}(S)$ is a (K, C)-quasi-isometry and E is a Dehn twist flat in $\mathcal{MCG}(S)$ then there exists a Dehn twist flat E' such that the Hausdorff distance between $f(E)$ and E' is at most A.*

Proof. The proof is essentially an argument by contradiction, using the Theorem 10.1. If there is no uniform control of the Hausdorff distance between quasi-isometric images of Dehn twist flats and Dehn twist flats, and 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 10.4, which is a consequence of Lemma 3.2).

We will find the following notation useful: Given subsets A, B of a metric space X and a point $p \in X$, we define a variant of the Hausdorff metric as follows: for each $r > 0$ we define

$$D_{r,p}(A,B) = \inf\{s \geq 0 \mid A \cap \mathcal{N}_r(p) \subset \mathcal{N}_s(B) \quad \text{and} \quad B \cap \mathcal{N}_r(p) \subset \mathcal{N}_s(A)\}$$

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 10.3. *Given (A_i), (B_i) two sequences of closed subsets, $A_\omega = B_\omega$ if and only if for each (p_i) and each $r \geq 0$ the ultralimit of $D_{r,p_i}(A_i, B_i)$ equals zero.*

Proof. Suppose that for some (p_i) and some $r > 0$ we have $D_{r,p_i}(A_i, B_i) \to_\omega \epsilon \in (0, \infty]$. Choose $\eta \in (0, \epsilon)$. It follows that, for ω-a.e. i, either

$$A_i \cap \mathcal{N}_r(p_i) \not\subset \mathcal{N}_\eta(B_i)$$

or

$$B_i \cap \mathcal{N}_r(p_i) \not\subset \mathcal{N}_\eta(A_i).$$

Exactly one of them holds for ω-a.e. i; let us assume the former. Then there exists a sequence

$$x_i \in A_i \cap \mathcal{N}_r(p_i) \setminus \mathcal{N}_\eta(B_i)$$

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 $\lim_\omega D_{r,p_i}(A_i, B_i) = 0$ for all (p_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 (p_i), so there exists some $r \geq 0$ such that $x_i \in A_i \cap \mathcal{N}_r(p_i)$ 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 $B_\omega \subset A_\omega$. □

Now suppose that Theorem 10.2 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 $f \colon \mathcal{MCG}(S) \to \mathcal{MCG}(S)$, 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 $f \colon \mathcal{MCG}(S) \to$

$\mathcal{MCG}(S)$ 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. Since 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, cf. [14]: it follows that π is uniformly onto, i.e., there exists a constant B depending only on K', C', n 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 $f_i\colon \mathcal{MCG}(S) \to \mathcal{MCG}(S)$, 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 precomposing 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 $f_\omega\colon \mathcal{M}_\omega(S) \to \mathcal{M}_\omega(S)$ and an Dehn twist flat $F_\omega = \lim_\omega(F)$, the asymptotic cone of F. Applying Theorem 10.1 we obtain a Dehn twist flat $F'_\omega = \lim_\omega F'_i$ such that $f_\omega(F_\omega) = F'_\omega$. It follows (applying Lemma 10.3) that, fixing any $R > 0$,

$$\frac{1}{s_i} D_{Rs_i,p_0}(f_i(F), F'_i) \to_\omega 0. \tag{10.2}$$

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 10.1 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 fact about Dehn twist flats:

Lemma 10.4. *There exist $\epsilon_1 \in (0, 1)$ 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 has nonempty* *intersection with* $\mathcal{N}_{\frac{r}{2}}(x)$ *and*

$$F_1 \cap \mathcal{N}_r(x) \subset \mathcal{N}_{\epsilon_1 r}(F_2)$$

then $F_1 = F_2$.

Proof. Express F_i as $\mathcal{Q}(P_i)$ for a pants decomposition P_i $(i = 1, 2)$. Lemma 3.2 implies that the junctures of $\mathcal{Q}(P_1)$ and $\mathcal{Q}(P_2)$ are $E_1 = \mathcal{Q}(P_1 \rfloor P_2)$ and $E_2 = \mathcal{Q}(P_2 \rfloor P_1)$. Note that $P_1 \rfloor P_2$ is a marking with base equal to P_1, and a transveral 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. Hence (F_i, E_i) is uniformly quasi-isometric to $(\mathbb{R}^n, \mathbb{R}^k)$ with $k < n$.

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 3.2 implies that $d(y, F_2)$ is bounded below (up to uniform coarse-Lipschitz error) by $d(y, E_1)$. Now use the elementary fact that, if $\mathbb{R}^k \subset \mathbb{R}^n$ with $k < n$ and $u \in \mathbb{R}^n$ then an R-ball $\mathcal{N}_R(u)$ is never contained in the $\frac{R}{2}$-neighborhood $\mathcal{N}_{\frac{R}{2}}(\mathbb{R}^k)$. After adjusting for the multiplicative errors, and setting r large enough to overcome the additive errors, we find that for suitable ϵ_1 we have

$$F_1 \cap \mathcal{N}_{\frac{r}{2}}(y) \not\subset \mathcal{N}_{\epsilon_1 r}(F_2).$$

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:

$$D_{Rs_i, p_0}(f_i(F), F'_i) \leq \epsilon_2 R s_i \leq s_i$$
$$D_{Rs_i, q_i}(f_i(F), F'_i) > s_i$$

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 $x'_i = p_{i,j+1}$,

$$D_{Rs_i, x_i}(f_i(F), F'_i) \leq \epsilon_2 R s_i, \tag{10.4}$$

but such that

$$D_{Rs_i, x'_i}(f_i(F), F'_i) \geq \epsilon_2 R s_i. \tag{10.5}$$

Now assuming $\epsilon_2 R s_i > 1$ (which is true for large enough i), we have $\mathcal{N}_{Rs_i}(x'_i) \subset \mathcal{N}_{Rs_i(1+\epsilon_2)}(x_i)$ and hence (10.5) implies

$$D_{Rs_i(1+\epsilon_2), x_i}(f_i(F), F'_i) \geq \epsilon_2 R s_i. \tag{10.6}$$

Now we apply Theorem 10.1 again using (x_i) as the basepoints, and we conclude via Lemma 10.3 that there exists a sequence (F''_i) of Dehn twist flats such that, for ω-almost every i,

$$D_{Rs_i(1+\epsilon_2), x_i}(f_i(F), F''_i) < \epsilon_2 R s_i \tag{10.7}$$

and in particular $F_i'' \neq F_i'$ for ω-a.e. i, by (10.6). Now (10.7) implies in particular that

$$f_i(F) \cap B_{Rs_i(1+\epsilon_2)}(x_i) \subset \mathcal{N}_{\epsilon_2 Rs_i}(F_i'').$$

Moreover by (10.4) we have

$$F_i' \cap \mathcal{N}_{Rs_i}(x_i) \subset \mathcal{N}_{\epsilon_2 Rs_i}(f_i(F))$$

and moreover (by triangle inequality)

$$F_i' \cap \mathcal{N}_{Rs_i}(x_i) \subset \mathcal{N}_{\epsilon_2 Rs_i}(f_i(F) \cap \mathcal{N}_{Rs_i(1+\epsilon_2)}(x_i)).$$

Putting this together we see

$$F_i' \cap \mathcal{N}_{Rs_i}(x_i) \subset \mathcal{N}_{2\epsilon_2 Rs_i}(F_i''). \tag{10.8}$$

Now since $x_i \in f_i(F_i)$ we note that $\mathcal{N}_{\epsilon_2 Rs_i}(x_i)$, which is contained in $\mathcal{N}_{\frac{Rs_i}{2}}(x_i)$, intersects F_i' nontrivially. Now (10.8) implies, by Lemma 10.4 (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 1.2, the classification of quasi-isometries of $\mathcal{MCG}(S)$.

Let $f\colon \mathcal{MCG}(S) \to \mathcal{MCG}(S)$ be a quasi-isometry. We first show that f induces an automorphism of the pants graph $\mathcal{P}(S)$. Recall that the *pants graph* of S is a simplicial graph, $\mathcal{P}(S)$, with vertices consisting of pants decompositions of S and with two vertices connected by an edge if the corresponding pants decompositions agree on all but one curve and those curves differ by an edge in the curve complex of the complexity one subsurface (complementary to the rest of the curves) in which they lie.

To do this we notice that the junctures discussed in Section 3 allow one to control the *coarse intersections* of Dehn twist flats. Given a metric space X and two subsets A, B, 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 $\{N_r(A) \cap N_r(B) \bigm| r \geq R\}$ have finite Hausdorff distance. For any subset $C \subset X$ 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.

Lemma 3.2 implies that, if P and P' are pants decompositions, then their associated Dehn twist flats $\mathcal{Q}(P)$ and $\mathcal{Q}(P')$ have well-defined coarse intersection, and moreover this intersection is represented equally well by $\mathcal{Q}(P \rfloor P')$ or $\mathcal{Q}(P' \rfloor P)$, which is itself a Dehn twist flat whose rank is the number of common curves of P and P'.

In particular, among pants decompositions the property of being adjacent in the pants graph is encoded by saying that their coarse intersection is quasi-isometric to Euclidean space of co-rank 1. This property is clearly preserved by quasi-isometries.

Since f coarsely preserves (maximal) Dehn twist flats by Theorem 10.2, it follows that it preserves the relation of adjacency among their definining pants decompositions, and hence induces an automorphism ϕ of $\mathcal{P}(S)$.

Now suppose that $S \neq S_{1,2}$, returning to this case in the end. Margalit proved that an automorphism of $\mathcal{P}(S)$ is induced by a homeomorphism of S [19]. Hence ϕ is induced by an element $\Phi \in \mathcal{MCG}(S)$. We finish by showing that there is a (uniform) bound on $d(f, \Phi)$.

Let P_1, P_2 be a pair of pants decompositions with no curves in common. Their junctures (or rather those of $\mathcal{Q}(P_i)$) are a pair of points, separated by some b_0 (we can uniformly control b_0 by making P_1 and P_2 as simple as possible). The $\mathcal{MCG}(S)$ orbit of the pair gives a family of such pairs whose junctures cover all of $\mathcal{M}(S)$, have diameter b_0, and satisfy the properties of junctures with the same constants. Now assuming $x \in \mathcal{M}(S)$ is one of the juncture points of P_1, P_2 (or any pair in its orbit), we see $\Phi(x)$ is a juncture point for $\phi(P_1), \phi(P_2)$. Moreover by Theorem 10.2, $f(x)$ satisfies the juncture properties for $\phi(P_1), \phi(P_2)$, but with worse constants depending on the quasi-isometry constants for f. This gives a bound between $\Phi(x)$ and $f(x)$, which is uniform as x varies in $\mathcal{M}(S)$. This gives the uniform bound on $d(f, \Phi)$.

The $S = S_{1,2}$ case: In this case we still get an automorphism ϕ as before, but it may no longer be induced by a mapping class. This is a finite-index problem caused by the hyperelliptic involution which we now recall. See Luo [18].

The quotient of $S_{1,2}$ by the hyperelliptic involution τ, minus 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 $\beta\colon \mathcal{MCG}(S_{1,2}) \to \mathcal{MCG}(S_{0,5})$ with kernel the center, or $\langle\tau\rangle = \mathbb{Z}/2\mathbb{Z}$. 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 punctures of $S_{1,2}$. Hence β is a quasi-isometry, and we let $\beta'\colon \mathcal{MCG}(S_{0,5}) \to \mathcal{MCG}(S_{1,2})$ be a quasi-inverse.

Now any quasi-isometry $f\colon \mathcal{MCG}(S_{1,2}) \to \mathcal{MCG}(S_{1,2})$ gives rise to a quasi-isometry $f' = \beta \circ f \circ \beta'$, and Theorem 1.2 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.

Remark about the proof. Margalit's Theorem on automorphisms of the pants graph is proved by reducing to Ivanov's theorem (and its completion by Korkmaz and Luo, [13, 17, 18]) that any automorphism of the curve complex of S is induced by a homeomorphism of that surface except for $\mathcal{C}(S_{1,2})$ which has automorphism that are not induced by homeomorphisms of $S_{1,2}$ (although they are all induced by homeomorphisms of $S_{1,2}$).

We do not make use of the full power of Margalit's Theorem. In our context, the automorphism of the pants graph is induced by an ambient quasi-isometry of $\mathcal{MCG}(S)$, and we were already making heavy use of this quasi-isometry in the proof of Theorem 10.2. In Margalit's proof there is no ambient quasi-isometry, and he must work hard to carry out the reduction to the theorem of Ivanov–Korkmaz–Luo.

Accordingly, by looking more closely at the flats we can carry out the proof by appealing directly to the theorem of Ivanov–Korkmaz–Luo: Using coarse intersections again, we can characterize intersection of flats in rank 1 or rank 2 subflats and this is preserved by the quasi-isometry. f therefore induces an action on rank 1 Dehn twist flats, i.e. curves of $\mathcal{C}(S)$, and preserves the property of two rank 1 flats belonging to a rank 2 flat – this is just adjacency in $\mathcal{C}(S)$. Hence f induces an automorphism of $\mathcal{C}(S)$, and from here we proceed similarly.

10.4. **Quasi-isometric rigidity**

We conclude with the proof of Theorem 1.1, quasi-isometric rigidity of $\mathcal{MCG}(S)$. The argument here is well-known, cf. [22].

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 $\lambda\colon G \to \mathrm{QI}(G)$, where $\mathrm{QI}(G)$ 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 1.2 implies that λ is surjective. Hence we have $\mathrm{QI}(G) = G/Z$.

Now if Γ is quasi-isometric to G then conjugation by the quasi-isometry Φ gives an isomorphism $\mathrm{QI}(\Gamma) \cong \mathrm{QI}(G)$ so we get a map $\lambda'\colon \Gamma \to \mathrm{QI}(G)$. 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 1.2 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 $\mathrm{QI}(G)$: 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 $\Gamma \to G/Z$ with finite kernel and finite-index image.

If $G = \mathcal{MCG}(S_{1,2})$, we observe as in the proof of Theorem 1.2 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] J. Behrstock, *Asymptotic geometry of the mapping class group and Teichmüller space*, Geometry & Topology **10** (2006), 1523–1578.

[2] J. Behrstock, C. Druţu, and L. Mosher, *Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity*, E-print arXiv:math.GT/0512592, 2005.

[3] J. Behrstock and Y. Minsky, *Dimension and rank for mapping class groups*, Ann. of Math., To appear. E-print arXiv:math.GT/0512352.

[4] 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.

[5] 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.

[6] M. J. Dunwoody, *The accessibility of finitely presented groups*, Invent. Math. **81** (1985), 449–457.

[7] M. Gromov, *Groups of polynomial growth and expanding maps*, IHES Sci. Publ. Math. **53** (1981), 53–73.

[8] ———, *Infinite groups as geometric objects*, Proceedings of the International Congress of Mathematicians, Warsaw, Amer. Math. Soc., 1983, pp. 385–392.

[9] 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.

[10] U. Hamenstädt, *Geometry of the mapping class groups III: Quasi-isometric rigidity, v.2*, Preprint, arXiv:math.GT/0512429, 2007.

[11] 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.

[12] N. V. Ivanov, *Complexes of curves and the Teichmüller modular group*, Uspekhi Mat. Nauk **42** (1987), 55–107.

[13] ———, *Automorphisms of complexes of curves and of Teichmüller spaces*, Internat. Math. Res. Notices (1997), no. 14, 651–666.

[14] M. Kapovich, *Lectures on the Geometric Group Theory*, Preprint. University of Utah. Available at `http://www.math.ucdavis.edu/~kapovich`, 2003.

[15] A. Karrass, A. Pietrowski, and D. Solitar, *Finite and infinite cyclic extensions of free groups*, J. Austral. Math. Soc. **16** (1973), 458–466, Collection of articles dedicated to the memory of Hanna Neumann, IV.

[16] B. Kleiner and B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, IHES Publ. Math. **86** (1997), 115–197.

[17] M. Korkmaz, *Automorphisms of complexes of curves on punctured spheres and on punctured tori*, Topology Appl. **95** (1999), no. 2, 85–111.

[18] F. Luo, *Automorphisms of the complex of curves*, Topology **39** (2000), no. 2, 283–298.

[19] D. Margalit, *Automorphisms of the pants complex*, Duke Math. J. **121** (2004), no. 3, 457–479.

[20] H. A. Masur and Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. **138** (1999), 103–149.

[21] ———, *Geometry of the complex of curves II: Hierarchical structure*, Geom. Funct. Anal. **10** (2000), 902–974.
[22] R. Schwartz, *The quasi-isometry classification of rank one lattices*, Inst. Hautes Études Sci. Publ. Math. (1995), no. 82, 133–168.
[23] J. Stallings, *On torsion free groups with infinitely many ends*, Ann. of Math. **88** (1968), 312–334.

Columbia University

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