

THE GEOMETRY OF THE DISK COMPLEX

HOWARD MASUR AND SAUL SCHLEIMER

ABSTRACT. We give a distance estimate for the metric on the disk complex and show that it is Gromov hyperbolic. As another application of our techniques, we find an algorithm which computes the Hempel distance of a Heegaard splitting, up to an error depending only on the genus.

CONTENTS

1. Introduction	2
2. Background on complexes	4
3. Background on coarse geometry	7
4. Natural maps	10
5. Holes in general and the lower bound on distance	13
6. Holes for the non-orientable surface	16
7. Holes for the arc complex	17
8. Background on three-manifolds	18
9. Holes for the disk complex	20
10. Holes for the disk complex – annuli	22
11. Holes for the disk complex – compressible	24
12. Holes for the disk complex – incompressible	27
13. Axioms for combinatorial complexes	34
14. Partition and the upper bound on distance	38
15. Background on Teichmüller space	49
16. Paths for the non-orientable surface	53
17. Paths for the arc complex	57
18. Background on train tracks	58
19. Paths for the disk complex	60
20. Hyperbolicity	62
21. Coarsely computing Hempel distance	67
References	71

Date: October 8, 2010.

This work is in the public domain.

1. INTRODUCTION

In this paper we initiate the study of the geometry of the disk complex of a handlebody V . The disk complex $\mathcal{D}(V)$ has a natural simplicial inclusion into the curve complex $\mathcal{C}(S)$ of the boundary of the handlebody. Surprisingly, this inclusion is not a quasi-isometric embedding; there are disks which are close in the curve complex yet very far apart in the disk complex. As we will show, any obstruction to joining such disks via a short path is a topologically meaningful subsurface of $S = \partial V$. We call such subsurfaces *holes*. A path in the disk complex must travel into and then out of these holes; paths in the curve complex may skip over a hole by using the vertex representing the boundary of the subsurface. We classify the holes:

Theorem 1.1. *Suppose V is a handlebody. If $X \subset \partial V$ is a hole for the disk complex $\mathcal{D}(V)$ of diameter at least 61 then:*

- X is not an annulus.
- If X is compressible then there are disks D, E with boundary contained in X so that the boundaries fill X .
- If X is incompressible then there is an I -bundle $\rho_F: T \rightarrow F$ so that T is a component of $V \setminus \partial_v T$ and X is a component of $\partial_h T$.

See Theorems 10.1, 11.6 and 12.1 for more precise statements. The I -bundles appearing in the classification lead us to study the complex of arcs embedded in the base surface F . Since the I -bundle may be twisted the surface F may be non-orientable.

Thus, as a necessary warm-up to the difficult case of the disk complex, we also analyse the holes for the curve complex of a non-orientable surface, as well as the holes for the arc complex $\mathcal{A}(S)$.

Topological application. It is a long-standing open problem to decide, given a Heegaard diagram, whether the underlying splitting surface is reducible. This question has deep connections to the geometry, topology, and algebra of the ambient three-manifold. For example, a resolution of this problem would give new solutions to both the three-sphere recognition problem and the triviality problem for three-manifold groups. The difficulty of deciding reducibility is underlined by its connection to the Poincaré conjecture: several approaches to the Poincaré Conjecture fell at essentially this point. See [10] for an entrance into the literature.

One generalization of deciding reducibility is to find an algorithm that, given a Heegaard diagram, computes the *distance* of the Heegaard splitting as defined by Hempel [20]. (For example, see [5, Section 2].)

The classification of holes for the disk complex leads to a coarse answer to this question.

Theorem 21.1. *In every genus g there is a constant K_g and an algorithm that, given a Heegaard diagram, computes the distance of the Heegaard splitting with error at most K_g .*

In addition to the classification of holes, the algorithm relies on the Gromov hyperbolicity of the curve complex [24] and the quasi-convexity of the disk set inside of the curve complex [26]. However the algorithm does not depend on our geometric applications of Theorem 1.1.

Geometric application. The hyperbolicity of the curve complex and the classification of holes allows us to prove:

Theorem 20.3. *The disk complex is Gromov hyperbolic.*

Again, as a warm-up to the proof of Theorem 20.3 we prove that $\mathcal{C}(F)$ and $\mathcal{A}(S)$ are hyperbolic in Corollary 6.4 and Theorem 20.2. Note that Bestvina and Fujiwara [4] have previously dealt with the curve complex of a non-orientable surface, following Bowditch [6].

These results cannot be deduced from the fact that $\mathcal{D}(V)$, $\mathcal{C}(F)$, and $\mathcal{A}(S)$ can be realized as quasi-convex subsets of $\mathcal{C}(S)$. This is because the curve complex is locally infinite. As simple example consider the Cayley graph of \mathbb{Z}^2 with the standard generating set. Then the cone $C(\mathbb{Z}^2)$ of height one-half is a Gromov hyperbolic space and \mathbb{Z}^2 is a quasi-convex subset. Another instructive example, very much in-line with our work, is the usual embedding of the three-valent tree T_3 into the Farey tessellation.

The proofs involve a distance estimate, as introduced by the first author and Yair Minsky [25]. We show that the distance in these complexes is coarsely equal to the sum of subsurface projection distances in holes. However, we do not use the hierarchy machine introduced in [25]. This is because hierarchies are too rigid for the disk complex and too flexible to respect a symmetry such as the involution giving a non-orientable surface. For the former we use the extremely flexible train track machine [27] and for the latter we use the highly rigid Teichmüller geodesic machine [33].

Theorems 19.1 and 20.3 are part of a more general framework. Namely, given a combinatorial complex \mathcal{G} we understand its geometry by classifying the holes: the geometric obstructions lying between \mathcal{G} and the curve complex. In Sections 13 and 14 we show that any complex \mathcal{G} satisfying certain axioms necessarily satisfies a distance estimate. That hyperbolicity follows from the axioms is proven in Section 20.

Our axioms are stated in terms of a path of markings, a path in the combinatorial complex, and their relationship. For the disk complex the combinatorial paths are surgery sequences of essential disks while the marking paths are provided by train track splitting sequences; both constructions are due to the first author and Minsky [26] (Section 18). The verification of the axioms (Section 19) relies on our work with Mosher, analysing train track splitting sequences in terms of subsurface projections [27].

We do not study non-orientable surfaces directly; instead we focus on symmetric multicurves in the double cover. This time marking paths are provided by Teichmüller geodesics, using the fact that the symmetric Riemann surfaces form a totally geodesic subset of Teichmüller space. The combinatorial path is given by the systole map. We use results of Rafi [33] to verify the axioms for the complex of symmetric curves. (See Section 16.) Section 17 verifies the axioms for the arc complex again using Teichmüller geodesics and the systole map. It is interesting to note that the axioms for the arc complex can also be verified using hierarchies or, indeed, train track splitting sequences.

The distance estimates for the marking graph and the pants graph, as given by the first author and Minsky [25], inspired the work here, but do not fit our framework. Indeed, neither the marking graph nor the pants graph are Gromov hyperbolic. It is crucial here that all holes *interfere*; this leads to hyperbolicity. When there are non-interfering holes, it is unclear how to partition the marking path to obtain the distance estimate.

Acknowledgments. We thank Jason Behrstock, Brian Bowditch, Yair Minsky, Lee Mosher, Hossein Namazi, and Kasra Rafi for many enlightening conversations.

We thank Tao Li for pointing out that our original bound inside of Theorem 12.1 of $O(\log g(V))$ could be reduced to a constant.

2. BACKGROUND ON COMPLEXES

We use $S_{g,b,c}$ to denote the compact connected surface of genus g with b boundary components and c cross-caps. If the surface is orientable we omit the subscript c and write $S_{g,b}$. The *complexity* of $S = S_{g,b}$ is $\xi(S) = 3g - 3 + b$. If the surface is closed and orientable we simply write S_g .

2.1. Arcs and curves. A simple closed curve $\alpha \subset S$ is *essential* if α does not bound a disk in S . The curve α is *non-peripheral* if α is not isotopic to a component of ∂S . A simple arc $\beta \subset S$ is proper if

$\beta \cap \partial S = \partial\beta$. An isotopy of S is proper if it preserves the boundary setwise. A proper arc $\beta \subset S$ is *essential* if β is not properly isotopic into a regular neighborhood of ∂S .

Define $\mathcal{C}(S)$ to be the set of isotopy classes of essential, non-peripheral curves in S . Define $\mathcal{A}(S)$ to be the set of proper isotopy classes of essential arcs. When $S = S_{0,2}$ is an annulus define $\mathcal{A}(S)$ to be the set of essential arcs, up to isotopies fixing the boundary pointwise. For any surface define $\mathcal{AC}(S) = \mathcal{A}(S) \cup \mathcal{C}(S)$.

For $\alpha, \beta \in \mathcal{AC}(S)$ the geometric intersection number $\iota(\alpha, \beta)$ is the minimum intersection possible between α and any β' equivalent to β . When $S = S_{0,2}$ we do not count intersection points occurring on the boundary. If α and β realize their geometric intersection number then α is *tight* with respect to β . If they do not realize their geometric intersection then we may *tighten* β until they do.

Define $\Delta \subset \mathcal{AC}(S)$ to be a *multicurve* if for all $\alpha, \beta \in \Delta$ we have $\iota(\alpha, \beta) = 0$. Following Harvey [18] we may impose the structure of a simplicial complex on $\mathcal{AC}(S)$: the simplices are exactly the multicurves. Also, $\mathcal{C}(S)$ and $\mathcal{A}(S)$ naturally span sub-complexes.

Note that the curve complexes $\mathcal{C}(S_{1,1})$ and $\mathcal{C}(S_{0,4})$ have no edges. It is useful to alter the definition in these cases. Place edges between all vertices with geometric intersection exactly one if $S = S_{1,1}$ or two if $S = S_{0,4}$. In both cases the result is the Farey graph. Also, with the current definition $\mathcal{C}(S)$ is empty if $S = S_{0,2}$. Thus for the annulus only we set $\mathcal{AC}(S) = \mathcal{C}(S) = \mathcal{A}(S)$.

Definition 2.1. For vertices $\alpha, \beta \in \mathcal{C}(S)$ define the *distance* $d_S(\alpha, \beta)$ to be the minimum possible number of edges of a path in the one-skeleton $\mathcal{C}^1(S)$ which starts at α and ends at β .

Note that if $d_S(\alpha, \beta) \geq 3$ then α and β *fill* the surface S . We denote distance in the one-skeleton of $\mathcal{A}(S)$ and of $\mathcal{AC}(S)$ by d_A and d_{AC} respectively. Recall that the geometric intersection of a pair of curves gives an upper bound for their distance.

Lemma 2.2. *Suppose that S is a compact connected surface which is not an annulus. For any $\alpha, \beta \in \mathcal{C}^0(S)$ with $\iota(\alpha, \beta) > 0$ we have $d_S(\alpha, \beta) \leq 2 \log_2(\iota(\alpha, \beta)) + 2$. \square*

This form of the inequality, stated for closed orientable surfaces, may be found in [20]. A proof in the bounded orientable case is given in [36]. The non-orientable case is then an exercise. When $S = S_{0,2}$ an induction proves

$$(2.3) \quad d_X(\alpha, \beta) = 1 + \iota(\alpha, \beta)$$

for distinct vertices $\alpha, \beta \in \mathcal{C}(X)$. See [25, Equation 2.3].

2.2. Subsurfaces. Suppose that $X \subset S$ is a connected compact subsurface. We say X is *essential* exactly when all boundary components of X are essential in S . We say that $\alpha \in \mathcal{AC}(S)$ *cuts* X if all representatives of α intersect X . If some representative is disjoint then we say α *misses* X .

Definition 2.4. An essential subsurface $X \subset S$ is *cleanly embedded* if for all components $\delta \subset \partial X$ we have: δ is isotopic into ∂S if and only if δ is equal to a component of ∂S .

Definition 2.5. Suppose $X, Y \subset S$ are essential subsurfaces. If X is cleanly embedded in Y then we say that X is *nested* in Y . If ∂X cuts Y and also ∂Y cuts X then we say that X and Y *overlap*.

A compact connected surface S is *simple* if $\mathcal{AC}(S)$ has finite diameter.

Lemma 2.6. *Suppose S is a connected compact surface. The following are equivalent:*

- S is not simple.
- The diameter of $\mathcal{AC}(S)$ is at least five.
- S admits an ending lamination or $S = S_1$ or $S_{0,2}$.
- S admits a pseudo-Anosov map or $S = S_1$ or $S_{0,2}$.
- $\chi(S) < -1$ or $S = S_{1,1}, S_1, S_{0,2}$.

Lemma 4.6 of [24] shows that pseudo-Anosov maps have quasi-geodesic orbits, when acting on the associated curve complex. A Dehn twist acting on $\mathcal{C}(S_{0,2})$ has geodesic orbits.

Note that Lemma 2.6 is only used in this paper when ∂S is non-empty. The closed case is included for completeness.

Proof sketch of Lemma 2.6. If S admits a pseudo-Anosov map then the stable lamination is an ending lamination. If S admits a filling lamination then, by an argument of Kobayashi [21], $\mathcal{AC}(S)$ has infinite diameter. (This argument is also sketched in [24], page 124, after the statement of Proposition 4.6.)

If the diameter of \mathcal{AC} is infinite then the diameter is at least five. To finish, one may check directly that all surfaces with $\chi(S) > -2$, other than $S_{1,1}$, S_1 and the annulus have $\mathcal{AC}(S)$ with diameter at most four. (The difficult cases, $S_{0,12}$ and $S_{0,03}$, are discussed by Scharlemann [35].) Alternatively, all surfaces with $\chi(S) < -1$, and also $S_{1,1}$, admit pseudo-Anosov maps. The orientable cases follow from Thurston's construction [38]. Penner's generalization [32] covers the non-orientable cases. \square

2.3. Handlebodies and disks. Let V_g denote the *handlebody* of genus g : the three-manifold obtained by taking a closed regular neighborhood of a polygonal, finite, connected graph in \mathbb{R}^3 . The genus of the boundary is the *genus* of the handlebody. A properly embedded disk $D \subset V$ is *essential* if $\partial D \subset \partial V$ is essential.

Let $\mathcal{D}(V)$ be the set of essential disks $D \subset V$, up to proper isotopy. A subset $\Delta \subset \mathcal{D}(V)$ is a *multidisk* if for every $D, E \in \Delta$ we have $\iota(\partial D, \partial E) = 0$. Following McCullough [28] we place a simplicial structure on $\mathcal{D}(V)$ by taking multidisks to be simplices. As with the curve complex, define $d_{\mathcal{D}}$ to be the distance in the one-skeleton of $\mathcal{D}(V)$.

2.4. Markings. A finite subset $\mu \subset \mathcal{AC}(S)$ *fills* S if for all $\beta \in \mathcal{C}(S)$ there is some $\alpha \in \mu$ so that $\iota(\alpha, \beta) > 0$. For any pair of finite subsets $\mu, \nu \subset \mathcal{AC}(S)$ we extend the intersection number:

$$\iota(\mu, \nu) = \sum_{\alpha \in \mu, \beta \in \nu} \iota(\alpha, \beta).$$

We say that μ, ν are *L-close* if $\iota(\mu, \nu) \leq L$. We say that μ is a *K-marking* if $\iota(\mu, \mu) \leq K$. For any K, L we may define $\mathcal{M}_{K,L}(S)$ to be the graph where vertices are filling K -markings and edges are given by L -closeness.

As defined in [25] we have:

Definition 2.7. A *complete clean marking* $\mu = \{\alpha_i\} \cup \{\beta_i\}$ consists of

- A collection of *base curves* $\text{base}(\mu) = \{\alpha_i\}$: a maximal simplex in $\mathcal{C}(S)$.
- A collection of *transversal curves* $\{\beta_i\}$: for each i define $X_i = S \setminus \bigcup_{j \neq i} \alpha_j$ and take $\beta_i \in \mathcal{C}(X_i)$ to be a Farey neighbor of α_i .

If μ is a complete clean marking then $\iota(\mu, \mu) \leq 2\xi(S) + 6\chi(S)$. As discussed in [25] there are two kinds of *elementary moves* which connected markings. There is a *twist* about a pants curve α , replacing its transversal β by a new transversal β' which is a Farey neighbor of both α and β . We can *flip* by swapping the roles of α_i and β_i . (In the case of the flip move, some of the other transversals must be *cleaned*.)

It follows that for any surface S there are choices of K, L so that $\mathcal{M}(S)$ is non-empty and connected. We use $d_{\mathcal{M}}(\mu, \nu)$ to denote distance in the marking graph.

3. BACKGROUND ON COARSE GEOMETRY

Here we review a few ideas from coarse geometry. See [8], [12], or [15] for a fuller discussion.

3.1. Quasi-isometry. Suppose r, s, A are non-negative real numbers, with $A \geq 1$. If $s \leq A \cdot r + A$ then we write $s \leq_A r$. If $s \leq_A r$ and $r \leq_A s$ then we write $s =_A r$ and call r and s *quasi-equal* with constant A . We also define the *cut-off function* $[r]_c$ where $[r]_c = 0$ if $r < c$ and $[r]_c = r$ if $r \geq c$.

Suppose that $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ are metric spaces. A relation $f: \mathcal{X} \rightarrow \mathcal{Y}$ is an A -*quasi-isometric embedding* for $A \geq 1$ if, for every $x, y \in \mathcal{X}$,

$$d_{\mathcal{X}}(x, y) =_A d_{\mathcal{Y}}(f(x), f(y)).$$

The relation f is a *quasi-isometry*, and \mathcal{X} is *quasi-isometric* to \mathcal{Y} , if f is an A -quasi-isometric embedding and the image of f is A -*dense*: the A -neighborhood of the image equals all of \mathcal{Y} .

3.2. Geodesics. Fix an interval $[u, v] \subset \mathbb{R}$. A *geodesic*, connecting x to y in \mathcal{X} , is an isometric embedding $f: [u, v] \rightarrow \mathcal{X}$ with $f(u) = x$ and $f(v) = y$. Often the exact choice of f is unimportant and all that matters are the endpoints x and y . We then denote the image of f by $[x, y] \subset \mathcal{X}$.

Fix now intervals $[m, n], [p, q] \subset \mathbb{Z}$. An A -quasi-isometric embedding $g: [m, n] \rightarrow \mathcal{X}$ is called an A -*quasi-geodesic* in \mathcal{X} . A function $g: [m, n] \rightarrow \mathcal{X}$ is an A -*unparameterized quasi-geodesic* in \mathcal{X} if

- there is an increasing function $\rho: [p, q] \rightarrow [m, n]$ so that $g \circ \rho: [p, q] \rightarrow \mathcal{X}$ is an A -*quasi-geodesic* in \mathcal{X} and
- for all $i \in [p, q - 1]$, $\text{diam}_{\mathcal{X}}(g[\rho(i), \rho(i + 1)]) \leq A$.

(Compare to the definition of (K, δ, s) -quasi-geodesics found in [24].)

A subset $\mathcal{Y} \subset \mathcal{X}$ is Q -*quasi-convex* if every \mathcal{X} -geodesic connecting a pair of points of \mathcal{Y} lies within a Q -neighborhood of \mathcal{Y} .

3.3. Hyperbolicity. We now assume that \mathcal{X} is a connected graph with metric induced by giving all edges length one.

Definition 3.1. The space \mathcal{X} is δ -*hyperbolic* if, for any three points x, y, z in \mathcal{X} and for any geodesics $k = [x, y]$, $g = [y, z]$, $h = [z, x]$, the triangle ghk is δ -*slim*: the δ -neighborhood of any two sides contains the third.

An important tool for this paper is the following theorem of the first author and Minsky [24]:

Theorem 3.2. *The curve complex of an orientable surface is Gromov hyperbolic.* \square

For the remainder of this section we assume that \mathcal{X} is δ -hyperbolic graph, $x, y, z \in \mathcal{X}$ are points, and $k = [x, y]$, $g = [y, z]$, $h = [z, x]$ are geodesics.

Definition 3.3. We take $\rho_k: \mathcal{X} \rightarrow k$ to be the *closest points relation*:

$$\rho_k(z) = \{w \in k \mid \text{for all } v \in k, d_{\mathcal{X}}(z, w) \leq d_{\mathcal{X}}(z, v)\}.$$

We now list several lemmas useful in the sequel.

Lemma 3.4. *There is a point on g within distance 2δ of $\rho_k(z)$. The same holds for h .* \square

Lemma 3.5. *The closest points $\rho_k(z)$ have diameter at most 4δ .* \square

Lemma 3.6. *The diameter of $\rho_g(x) \cup \rho_h(y) \cup \rho_k(z)$ is at most 6δ .* \square

Lemma 3.7. *Suppose that z' is another point in \mathcal{X} so that $d_{\mathcal{X}}(z, z') \leq R$. Then $d_{\mathcal{X}}(\rho_k(z), \rho_k(z')) \leq R + 6\delta$.* \square

Lemma 3.8. *Suppose that k' is another geodesic in X so that the endpoints of k' are within distance R of the points x and y . Then $d_X(\rho_k(z), \rho_{k'}(z)) \leq R + 11\delta$.* \square

We now turn to a highly useful consequence of the Morse stability of quasi-geodesics in hyperbolic spaces.

Lemma 3.9. *For every δ and A there is a constant C with the following property: If \mathcal{X} is δ -hyperbolic and $g: [0, N] \rightarrow \mathcal{X}$ is an A -unparametrized quasi-geodesic then for any $m < n < p$ in $[0, N]$ we have:*

$$d_{\mathcal{X}}(x, y) + d_{\mathcal{X}}(y, z) < d_{\mathcal{X}}(x, z) + C$$

where $x, y, z = g(m), g(n), g(p)$. \square

3.4. A hyperbolicity criterion. Here we give a hyperbolicity criterion tailored to our setting. We thank Brian Bowditch for both finding an error in our first proof of Theorem 3.11 and for informing us of Gilman's work [13, 14].

Suppose that \mathcal{X} is a graph with all edge-lengths equal to one. Suppose that $\gamma: [0, N] \rightarrow \mathcal{X}$ is a loop in \mathcal{X} with unit speed. Any pair of points $a, b \in [0, N]$ gives a *chord* of γ . If $a < b$, $N/4 \leq b - a$ and $N/4 \leq a + (N - b)$ then the chord is *1/4-separated*. The length of the chord is $d_{\mathcal{X}}(\gamma(a), \gamma(b))$.

Following Gilman [13, Theorem B] we have:

Theorem 3.10. *Suppose that \mathcal{X} is a graph with all edge-lengths equal to one. Then \mathcal{X} is Gromov hyperbolic if and only if there is a constant K so that every loop $\gamma: [0, N] \rightarrow \mathcal{X}$ has a 1/4-separated chord of length at most $N/7 + K$.* \square

Gilman's proof goes via the subquadratic isoperimetric inequality. We now give our criterion, noting that it is closely related to another paper of Gilman [14].

Theorem 3.11. *Suppose that \mathcal{X} is a graph with all edge-lengths equal to one. Then \mathcal{X} is Gromov hyperbolic if and only if there is a constant $M \geq 0$ and, for all unordered pairs $x, y \in \mathcal{X}^0$, there is a connected subgraph $g_{x,y}$ containing x and y with the following properties:*

- (Local) *If $d_{\mathcal{X}}(x, y) \leq 1$ then $g_{x,y}$ has diameter at most M .*
- (Slim triangles) *For all $x, y, z \in \mathcal{X}^0$ the subgraph $g_{x,y}$ is contained in an M -neighborhood of $g_{y,z} \cup g_{z,x}$.*

Proof. Suppose that $\gamma: [0, N] \rightarrow \mathcal{X}$ is a loop. If ϵ is the empty string let $I_{\epsilon} = [0, N]$. For any binary string ω let I_{ω_0} and I_{ω_1} be the first and second half of I_{ω} . Note that if $|\omega| \geq \lceil \log_2 N \rceil$ then $|I_{\omega}| \leq 1$.

Fix a string ω and let $[a, b] = I_{\omega}$. Let g_{ω} be the subgraph connecting $\gamma(a)$ to $\gamma(b)$. Note that $g_0 = g_1$ because $\gamma(0) = \gamma(N)$. Also, for any binary string ω the subgraphs $g_{\omega}, g_{\omega_0}, g_{\omega_1}$ form an M -slim triangle. If $|\omega| \leq \lceil \log_2 N \rceil$ then every $x \in g_{\omega}$ has some point $b \in I_{\omega}$ so that

$$d_{\mathcal{X}}(x, \gamma(b)) \leq M(\lceil \log_2 N \rceil - |\omega|) + 2M.$$

Since g_0 is connected there is a point $x \in g_0$ that lies within the M -neighborhoods both of g_{00} and of g_{01} . Pick some $b \in I_1$ so that $d_{\mathcal{X}}(x, \gamma(b))$ is bounded as in the previous paragraph. It follows that there is a point $a \in I_0$ so that a, b are $1/4$ -separated and so that

$$d_{\mathcal{X}}(\gamma(a), \gamma(b)) \leq 2M\lceil \log_2 N \rceil + 2M.$$

Thus there is an additive error K large enough so that \mathcal{X} satisfies the criterion of Theorem 3.10 and we are done. \square

4. NATURAL MAPS

There are several natural maps between the complexes and graphs defined in Section 2. Here we review what is known about their geometric properties, and give examples relevant to the rest of the paper.

4.1. Lifting, surgery, and subsurface projection. Suppose that S is not simple. Choose a hyperbolic metric on the interior of S so that all ends have infinite areas. Fix a compact essential subsurface $X \subset S$ which is not a peripheral annulus. Let S^X be the cover of S so that X lifts homeomorphically and so that $S^X \cong \text{interior}(X)$. For any $\alpha \in \mathcal{AC}(S)$ let α^X be the full preimage.

Since there is a homeomorphism between X and the Gromov compactification of S^X in a small abuse of notation we identify $\mathcal{AC}(X)$ with the arc and curve complex of S^X .

Definition 4.1. We define the *cutting relation* $\kappa_X: \mathcal{AC}(S) \rightarrow \mathcal{AC}(X)$ as follows: $\alpha' \in \kappa_X(\alpha)$ if and only if α' is an essential non-peripheral component of α^X .

Note that α cuts X if and only if $\kappa_X(\alpha)$ is non-empty. Now suppose that S is not an annulus.

Definition 4.2. We define the *surgeries relation* $\sigma_X: \mathcal{AC}(S) \rightarrow \mathcal{C}(S)$ as follows: $\alpha' \in \sigma_S(\alpha)$ if and only if $\alpha' \in \mathcal{C}(S)$ is a boundary component of a regular neighborhood of $\alpha \cup \partial S$.

With S and X as above:

Definition 4.3. We define the *subsurface projection relation* $\pi_X: \mathcal{AC}(S) \rightarrow \mathcal{C}(X)$ as follows: If X is not an annulus define then $\pi_X = \sigma_X \circ \kappa_X$. When X is an annulus $\pi_X = \kappa_X$.

If $\alpha, \beta \in \mathcal{AC}(S)$ both cut X we write $d_X(\alpha, \beta) = \text{diam}_X(\pi_X(\alpha) \cup \pi_X(\beta))$. This is the *subsurface projection distance* between α and β in X .

Lemma 4.4. *Suppose $\alpha, \beta \in \mathcal{AC}(S)$ are disjoint and cut X . Then $\text{diam}_X(\pi_X(\alpha)), d_X(\alpha, \beta) \leq 3$. \square*

See Lemma 2.3 of [25] and the remarks in the section Projection Bounds in [29].

Corollary 4.5. *Fix $X \subset S$. Suppose that $\{\beta_i\}_{i=0}^N$ is a path in $\mathcal{AC}(S)$. Suppose that β_i cuts X for all i . Then $d_X(\beta_0, \beta_N) \leq 3N + 3$. \square*

It is crucial to note that if some vertex of $\{\beta_i\}$ misses X then the projection distance $d_X(\beta_0, \beta_n)$ may be arbitrarily large compared to n . Corollary 4.5 can be greatly strengthened when the path is a geodesic [25]:

Theorem 4.6. *[Bounded Geodesic Image] There is constant M_0 with the following property. Fix $X \subset S$. Suppose that $\{\beta_i\}_{i=0}^n$ is a geodesic in $\mathcal{C}(S)$. Suppose that β_i cuts X for all i . Then $d_X(\beta_0, \beta_n) \leq M_0$. \square*

Here is a converse for Lemma 4.4.

Lemma 4.7. *For every $a \in \mathbb{N}$ there is a number $b \in \mathbb{N}$ with the following property: for any $\alpha, \beta \in \mathcal{AC}(S)$ if $d_X(\alpha, \beta) \leq a$ for all $X \subset S$ then $\iota(\alpha, \beta) \leq b$.*

Corollary D of [11] gives a more precise relation between projection distance and intersection number.

Proof of Lemma 4.7. We only sketch the contrapositive: Suppose we are given a sequence of curves α_n, β_n so that $\iota(\alpha_n, \beta_n)$ tends to infinity. Passing to subsequences and applying elements of the mapping class group we may assume that $\alpha_n = \alpha_0$ for all n . Setting $c_n = \iota(\alpha_0, \beta_n)$ and passing to subsequences again we may assume that β_n/c_n converges

to $\lambda \in \mathcal{PML}(S)$, the projectivization of Thurston's space of measured laminations. Let Y be any connected component of the subsurface filled by λ , chosen so that α_0 cuts Y . Note that $\pi_Y(\beta_n)$ converges to $\lambda|_Y$. Again applying Kobayashi's argument [21], the distance $d_Y(\alpha_0, \beta_n)$ tends to infinity. \square

4.2. Inclusions. We now record a well known fact:

Lemma 4.8. *The inclusion $\nu: \mathcal{C}(S) \rightarrow \mathcal{AC}(S)$ is a quasi-isometry. The surgery map $\sigma_S: \mathcal{AC}(S) \rightarrow \mathcal{C}(S)$ is a quasi-inverse for ν .*

Proof. Fix $\alpha, \beta \in \mathcal{C}(S)$. Since ν is an inclusion we have $d_{\mathcal{AC}}(\alpha, \beta) \leq d_S(\alpha, \beta)$. In the other direction, let $\{\alpha_i\}_{i=0}^N$ be a geodesic in $\mathcal{AC}(S)$ connecting α to β . Since every α_i cuts S we apply Corollary 4.5 and deduce $d_S(\alpha, \beta) \leq 3N + 3$.

Note that the composition $\sigma_S \circ \nu = \text{Id}|_{\mathcal{C}(S)}$. Also, for any arc $\alpha \in \mathcal{A}(S)$ we have $d_{\mathcal{AC}}(\alpha, \nu(\sigma_S(\alpha))) = 1$. Finally, $\mathcal{C}(S)$ is 1-dense in $\mathcal{AC}(S)$, as any arc $\gamma \subset S$ is disjoint from the one or two curves of $\sigma_S(\gamma)$. \square

Brian Bowditch raised the question, at the Newton Institute in August 2003, of the geometric properties of the inclusion $\mathcal{A}(S) \rightarrow \mathcal{AC}(S)$. The natural assumption, that this inclusion is again a quasi-isometric embedding, is false. In this paper we will exactly characterize how the inclusion distorts distance.

We now move up a dimension. Suppose that V is a handlebody and $S = \partial V$. We may take any disk $D \in \mathcal{D}(V)$ to its boundary $\partial D \in \mathcal{C}(S)$, giving an inclusion $\nu: \mathcal{D}(V) \rightarrow \mathcal{C}(S)$. It is important to distinguish the disk complex from its image $\nu(\mathcal{D}(V))$; thus we will call the image the *disk set*.

The first author and Minsky [26] have shown:

Theorem 4.9. *The disk set is a quasi-convex subset of the curve complex.* \square

It is natural to ask if this map is a quasi-isometric embedding. If so, the hyperbolicity of $\mathcal{C}(V)$ immediately follows. In fact, the inclusion again badly distorts distance and we investigate exactly how, below.

4.3. Markings and the mapping class group. Once the connectedness of $\mathcal{M}(S)$ is in hand, it is possible to use local finiteness to show that $\mathcal{M}(S)$ is quasi-isometric to the Cayley graph of the mapping class group [25].

Using subsurface projections the first author and Minsky [25] obtained a *distance estimate* for the marking complex and thus for the mapping class group.

Theorem 4.10. *There is a constant $C_0 = C_0(S)$ so that, for any $c \geq C_0$ there is a constant A with*

$$d_{\mathcal{M}}(\mu, \mu') =_A \sum [d_X(\mu, \mu')]_c$$

independent of the choice of μ and μ' . Here the sum ranges over all essential, non-peripheral subsurfaces $X \subset S$.

This, and their similar estimate for the pants graph, is a model for the distance estimates given below. Notice that a filling marking $\mu \in \mathcal{M}(S)$ cuts all essential, non-peripheral subsurfaces of S . It is not an accident that the sum ranges over the same set.

5. HOLES IN GENERAL AND THE LOWER BOUND ON DISTANCE

Suppose that S is a compact connected surface. In this paper a *combinatorial complex* $\mathcal{G}(S)$ will have vertices being isotopy classes of certain multicurves in S . We will assume throughout that vertices of $\mathcal{G}(S)$ are connected by edges only if there are representatives which are disjoint. This assumption is made only to simplify the proofs — all arguments work in the case where adjacent vertices are allowed to have uniformly bounded intersection. In all cases \mathcal{G} will be connected. There is a natural map $\nu: \mathcal{G} \rightarrow \mathcal{AC}(S)$ taking a vertex of \mathcal{G} to the isotopy classes of the components. Examples in the literature include the marking complex [25], the pants complex [9] [2], the Hatcher-Thurston complex [19], the complex of separating curves [7], the arc complex and the curve complexes themselves.

For any combinatorial complex \mathcal{G} defined in this paper *other than the curve complex* we will denote distance in the one-skeleton of \mathcal{G} by $d_{\mathcal{G}}(\cdot, \cdot)$. Distance in $\mathcal{C}(S)$ will always be denoted by $d_S(\cdot, \cdot)$.

5.1. Holes, defined. Suppose that S is non-simple. Suppose that $\mathcal{G}(S)$ is a combinatorial complex. Suppose that $X \subset S$ is a cleanly embedded subsurface. A vertex $\alpha \in \mathcal{G}$ *cuts* X if some component of α cuts X .

Definition 5.1. We say $X \subset S$ is a *hole* for \mathcal{G} if every vertex of \mathcal{G} cuts X .

Almost equivalently, if X is a hole then the subsurface projection $\pi_X: \mathcal{G}(S) \rightarrow \mathcal{C}(X)$ never takes the empty set as a value. Note that the entire surface S is always a hole, regardless of our choice of \mathcal{G} . A boundary parallel annulus cannot be cleanly embedded (unless S is also an annulus), so generally cannot be a hole. A hole $X \subset S$ is *strict* if X is not homeomorphic to S .

We now classify the holes for $\mathcal{A}(S)$.

Example 5.2. Suppose that $S = S_{g,b}$ with $b > 0$ and consider the arc complex $\mathcal{A}(S)$. The holes, up to isotopy, are exactly the cleanly embedded surfaces which contain ∂S . So, for example, if S is planar then only S is a hole for $\mathcal{A}(S)$. The same holds for $S = S_{1,1}$. In these cases it is an easy exercise to show that $\mathcal{C}(S)$ and $\mathcal{A}(S)$ are quasi-isometric. In all other cases the arc complex admits infinitely many holes.

Definition 5.3. If X is a hole and if $\pi_X(\mathcal{G}) \subset \mathcal{C}(X)$ has diameter at least R we say that the hole X has *diameter* at least R .

Example 5.4. Continuing the example above: Since the mapping class group acts on the arc complex, all non-simple holes for $\mathcal{A}(S)$ have infinite diameter.

Suppose now that $X, X' \subset S$ are disjoint holes for \mathcal{G} . In the presence of symmetry there can be a relationship between $\pi_X|_{\mathcal{G}}$ and $\pi_{X'}|_{\mathcal{G}}$ as follows:

Definition 5.5. Suppose that X, X' are holes for \mathcal{G} , both of infinite diameter. Then X and X' are *paired* if there is a homeomorphism $\tau: X \rightarrow X'$ and a constant L_4 so that

$$d_{X'}(\pi_{X'}(\gamma), \tau(\pi_X(\gamma))) \leq L_4$$

for every $\gamma \in \mathcal{G}$. Furthermore, if $Y \subset X$ is a hole then τ pairs Y with $Y' = \tau(Y)$. Lastly, pairing is required to be symmetric; if τ pairs X with X' then τ^{-1} pairs X' with X .

Definition 5.6. Two holes X and Y *interfere* if either $X \cap Y \neq \emptyset$ or X is paired with X' and $X' \cap Y \neq \emptyset$.

Examples arise in the symmetric arc complex and in the discussion of twisted I -bundles inside of a handlebody.

5.2. Projection to holes is coarsely Lipschitz. The following lemma is used repeatedly throughout the paper:

Lemma 5.7. *Suppose that $\mathcal{G}(S)$ is a combinatorial complex. Suppose that X is a hole for \mathcal{G} . Then for any $\alpha, \beta \in \mathcal{G}$ we have*

$$d_X(\alpha, \beta) \leq 3 + 3 \cdot d_{\mathcal{G}}(\alpha, \beta).$$

The additive error is required only when $\alpha = \beta$.

Proof. This follows directly from Corollary 4.5 and our assumption that vertices of \mathcal{G} connected by an edge represent disjoint multicurves. \square

5.3. Infinite diameter holes. We may now state a first answer to Bowditch's question.

Lemma 5.8. *Suppose that $\mathcal{G}(S)$ is a combinatorial complex. Suppose that there is a strict hole $X \subset S$ having infinite diameter. Then $\nu: \mathcal{G} \rightarrow \mathcal{AC}(S)$ is not a quasi-isometric embedding. \square*

This lemma and Example 5.2 completely determines when the inclusion of $\mathcal{A}(S)$ into $\mathcal{AC}(S)$ is a quasi-isometric embedding. It quickly becomes clear that the set of holes tightly constrains the intrinsic geometry of a combinatorial complex.

Lemma 5.9. *Suppose that $\mathcal{G}(S)$ is a combinatorial complex invariant under the natural action of $\text{MCG}(S)$. Then every non-simple hole for \mathcal{G} has infinite diameter. Furthermore, if $X, Y \subset S$ are disjoint non-simple holes for \mathcal{G} then there is a quasi-isometric embedding of \mathbb{Z}^2 into \mathcal{G} . \square*

We will not use Lemmas 5.8 or 5.9 and so omit the proofs. Instead our interest lies in proving the far more powerful distance estimate (Theorems 5.10 and 13.1) for $\mathcal{G}(S)$.

5.4. A lower bound on distance. Here we see that the sum of projection distances in holes gives a lower bound for distance.

Theorem 5.10. *Fix S , a compact connected non-simple surface. Suppose that $\mathcal{G}(S)$ is a combinatorial complex. Then there is a constant C_0 so that for all $c \geq C_0$ there is a constant A satisfying*

$$\sum [d_X(\alpha, \beta)]_c \leq_A d_{\mathcal{G}}(\alpha, \beta).$$

Here $\alpha, \beta \in \mathcal{G}$ and the sum is taken over all holes X for the complex \mathcal{G} . \square

The proof follows the proof of Theorems 6.10 and 6.12 of [25], practically word for word. The only changes necessary are to

- replace the sum over *all* subsurfaces by the sum over all holes,
- replace Lemma 2.5 of [25], which records how markings differing by an elementary move project to an essential subsurface, by Lemma 5.7 of this paper, which records how \mathcal{G} projects to a hole.

One major goal of this paper is to give criteria sufficient to obtain the reverse inequality; Theorem 13.1.

6. HOLES FOR THE NON-ORIENTABLE SURFACE

Fix F a compact, connected, and non-orientable surface. Let S be the orientation double cover with covering map $\rho_F: S \rightarrow F$. Let $\tau: S \rightarrow S$ be the associated involution; so for all $x \in S$, $\rho_F(x) = \rho_F(\tau(x))$.

Definition 6.1. A multicurve $\gamma \subset \mathcal{AC}(S)$ is *symmetric* if $\tau(\gamma) \cap \gamma = \emptyset$ or $\tau(\gamma) = \gamma$. A multicurve γ is *invariant* if there is a curve or arc $\gamma' \subset F$ so that $\gamma = \rho_F^{-1}(\gamma')$. The same definitions holds for subsurfaces $X \subset S$.

Definition 6.2. The *invariant complex* $\mathcal{C}^\tau(S)$ is the simplicial complex with vertex set being isotopy classes of invariant multicurves. There is a k -simplex for every collection of $k + 1$ distinct isotopy classes having pairwise disjoint representatives.

Notice that $\mathcal{C}^\tau(S)$ is simplicially isomorphic to $\mathcal{C}(F)$. There is also a natural map $\nu: \mathcal{C}^\tau(S) \rightarrow \mathcal{C}(S)$. We will prove:

Lemma 6.3. $\nu: \mathcal{C}^\tau(S) \rightarrow \mathcal{C}(S)$ is a quasi-isometric embedding.

It thus follows from the hyperbolicity of $\mathcal{C}(S)$ that:

Corollary 6.4 ([4]). $\mathcal{C}(F)$ is Gromov hyperbolic. \square

We begin the proof of Lemma 6.3: since ν sends adjacent vertices to adjacent edges we have

$$(6.5) \quad d_S(\alpha, \beta) \leq d_{\mathcal{C}^\tau}(\alpha, \beta),$$

as long as α and β are distinct in $\mathcal{C}^\tau(S)$. In fact, since the surface S itself is a hole for $\mathcal{C}^\tau(S)$ we may deduce a slightly weaker lower bound from Lemma 5.7 or indeed from Theorem 5.10.

The other half of the proof of Lemma 6.3 consists of showing that S is the *only* hole for $\mathcal{C}^\tau(S)$ with large diameter. After a discussion of Teichmüller geodesics we will prove:

Lemma 16.4. *There is a constant K with the following property: Suppose that α, β are invariant multicurves in S . Suppose that $X \subset S$ is an essential subsurface where $d_X(\alpha, \beta) > K$. Then X is symmetric.*

From this it follows that:

Corollary 6.6. *With K as in Lemma 16.4: If $X \subset S$ is a hole for $\mathcal{C}^\tau(S)$ with diameter greater than K then $X = S$.*

Proof. Suppose that $X \subset S$ is a strict subsurface, cleanly embedded. Suppose that $\text{diam}_X(\mathcal{C}^\tau(S)) > K$. Thus X is symmetric. It follows that $\partial X \setminus \partial S$ is also symmetric. Since ∂X does not cut X deduce that X is not a hole for $\mathcal{C}^\tau(S)$. \square

This corollary, together with the upper bound (Theorem 13.1), proves Lemma 6.3.

7. HOLES FOR THE ARC COMPLEX

Here we generalize the definition of the arc complex and classify its holes.

Definition 7.1. Suppose that S is a non-simple surface with boundary. Let Δ be a non-empty collection of components of ∂S . The *arc complex* $\mathcal{A}(S, \Delta)$ is the subcomplex of $\mathcal{A}(S)$ spanned by essential arcs $\alpha \subset S$ with $\partial\alpha \subset \Delta$.

Note that $\mathcal{A}(S, \partial S)$ and $\mathcal{A}(S)$ are identical.

Lemma 7.2. *Suppose $X \subset S$ is cleanly embedded. Then X is a hole for $\mathcal{A}(S, \Delta)$ if and only if $\Delta \subset \partial X$.* \square

This follows directly from the definition of a hole. We now have an straight-forward observation:

Lemma 7.3. *If $X, Y \subset S$ are holes for $\mathcal{A}(S, \Delta)$ then $X \cap Y \neq \emptyset$.* \square

The proof follows immediately from Lemma 7.2. Lemma 5.9 indicates that Lemma 7.3 is essential to proving that $\mathcal{A}(S, \Delta)$ is Gromov hyperbolic.

In order to prove the upper bound theorem for \mathcal{A} we will use pants decompositions of the surface S . In an attempt to avoid complications in the non-orientable case we must carefully lift to the orientation cover.

Suppose that F is non-simple, non-orientable, and has non-empty boundary. Let $\rho_F: S \rightarrow F$ be the orientation double cover and let $\tau: S \rightarrow S$ be the induced involution. Fix $\Delta' \subset \partial F$ and let $\Delta = \rho_F^{-1}(\Delta')$.

Definition 7.4. We define $\mathcal{A}^\tau(S, \Delta)$ to be the *invariant arc complex*: vertices are invariant multi-arcs and simplices arise from disjointness.

Again, $\mathcal{A}^\tau(S, \Delta)$ is simplicially isomorphic to $\mathcal{A}(F, \Delta')$. If $X \cap \tau(X) = \emptyset$ and $\Delta \subset X \cup \tau(X)$ then the subsurfaces X and $\tau(X)$ are paired holes, as in Definition 5.5. Notice as well that all non-simple symmetric holes $X \subset S$ for $\mathcal{A}^\tau(S, \Delta)$ have infinite diameter.

Unlike $\mathcal{A}(F, \Delta')$ the complex $\mathcal{A}^\tau(S, \Delta)$ may have disjoint holes. Nonetheless, we have:

Lemma 7.5. *Any two non-simple holes for $\mathcal{A}^\tau(S, \Delta)$ interfere.*

Proof. Suppose that X, Y are holes for the τ -invariant arc complex, $\mathcal{A}^\tau(S, \Delta)$. It follows from Lemma 16.4 that X is symmetric with $\Delta \subset X \cup \tau(X)$. The same holds for Y . Thus Y must cut either X or $\tau(X)$. \square

8. BACKGROUND ON THREE-MANIFOLDS

Before discussing the holes in the disk complex, we record a few facts about handlebodies and I -bundles.

Fix M a compact connected irreducible three-manifold. Recall that M is *irreducible* if every embedded two-sphere in M bounds a three-ball. Recall that if N is a closed submanifold of M then $\text{fr}(N)$, the frontier of N in M , is the closure of $\partial N \setminus \partial M$.

8.1. Compressions. Suppose that F is a surface embedded in M . Then F is *compressible* if there is a disk B embedded in M with $B \cap \partial M = \emptyset$, $B \cap F = \partial B$, and ∂B essential in F . Any such disk B is called a *compression* of F .

In this situation form a new surface F' as follows: Let N be a closed regular neighborhood of B . First remove from F the annulus $N \cap F$. Now form F' by gluing on both disk components of $\partial N \setminus F$. We say that F' is obtained by *compressing* F along B . If no such disk exists we say F is *incompressible*.

Definition 8.1. A properly embedded surface F is *boundary compressible* if there is a disk B embedded in M with

- $\text{interior}(B) \cap \partial M = \emptyset$,
- ∂B is a union of connected arcs α and β ,
- $\alpha \cap \beta = \partial\alpha = \partial\beta$,
- $B \cap F = \alpha$ and α is properly embedded in F ,
- $B \cap \partial M = \beta$, and
- β is essential in $\partial M \setminus \partial F$.

A disk, like B , with boundary partitioned into two arcs is called a *bigon*. Note that this definition of boundary compression is slightly weaker than some found in the literature; the arc α is often required to be essential in F . We do not require this additional property because, for us, F will usually be a properly embedded disk in a handlebody.

Just as for compressing disks we may *boundary compress* F along B to obtain a new surface F' : Let N be a closed regular neighborhood of B . First remove from F the rectangle $N \cap F$. Now form F' by gluing on both bigon components of $\text{fr}(N) \setminus F$. Again, F' is obtained by *boundary compressing* F along B . Note that the relevant boundary components of F and F' cobound a pair of pants embedded in ∂M . If no boundary compression exists then F is *boundary incompressible*.

Remark 8.2. Recall that any surface F properly embedded in a handlebody V_g , $g \geq 2$, is either compressible or boundary compressible.

Suppose now that F is properly embedded in M and Γ is a multicurve in ∂M .

Remark 8.3. Suppose that F' is obtained by a boundary compression of F performed in the complement of Γ . Suppose that $F' = F_1 \cap F_2$ is disconnected and each F_i cuts Γ . Then $\iota(\partial F_i, \Gamma) < \iota(\partial F, \Gamma)$ for $i = 1, 2$.

It is often useful to restrict our attention to boundary compressions meeting a single subsurface of ∂M . So suppose that $X \subset \partial M$ is an essential subsurface. Suppose that ∂F is tight with respect to ∂X . Suppose B is a boundary compression of F . If $B \cap \partial M \subset X$ we say that F is *boundary compressible into X* .

Lemma 8.4. *Suppose that M is irreducible. Fix X a connected essential subsurface of ∂M . Let $F \subset M$ be a properly embedded, incompressible surface. Suppose that ∂X and ∂F are tight and that X compresses in M . Then either:*

- $F \cap X = \emptyset$,
- F is boundary compressible into X , or
- F is a disk with $\partial F \subset X$.

Proof. Suppose that X is compressible via a disk E . Isotope E to make ∂E tight with respect to ∂F . This can be done while maintaining $\partial E \subset X$ because ∂F and ∂X are tight. Since M is irreducible and F is incompressible we may isotope E , rel ∂ , to remove all simple closed curves of $F \cap E$. If $F \cap E$ is non-empty then an outermost bigon of E gives the desired boundary compression lying in X .

Suppose instead that $F \cap E = \emptyset$ but F does cut X . Let $\delta \subset X$ be a simple arc meeting each of F and E in exactly one endpoint. Let N be a closed regular neighborhood of $\delta \cup E$. Note that $\text{fr}(N) \setminus F$ has three components. One is a properly embedded disk parallel to E and the other two B, B' are bigons attached to F . At least one of these, say B' is trivial in the sense that $B' \cap \partial M$ is a trivial arc embedded in $\partial M \setminus \partial F$. If B is non-trivial then B provides the desired boundary compression.

Suppose that B is also trivial. It follows that ∂E and one component $\gamma \subset \partial F$ cobound an annulus $A \subset X$. So $D = A \cup E$ is a disk with $(D, \partial D) \subset (M, F)$. As $\partial D = \gamma$ and F is incompressible and M is irreducible deduce that F is isotopic to E . \square

8.2. Band sums. A *band sum* is the inverse operation to boundary compression: Fix a pair of disjoint properly embedded surfaces $F_1, F_2 \subset M$. Let $F' = F_1 \cup F_2$. Fix a simple arc $\delta \subset \partial M$ so that δ meets each of F_1 and F_2 in exactly one point of $\partial \delta$. Let $N \subset M$ be a closed regular

neighborhood of δ . Form a new surface by adding to $F' \setminus N$ the rectangle component of $\text{fr}(N) \setminus F'$. The surface F obtained is the result of *band summing* F_1 to F_2 along δ . Note that F has a boundary compression dual to δ yielding F' : that is, there is a boundary compression B for F so that $\delta \cap B$ is a single point and compressing F along B gives F' .

8.3. Handlebodies and I-bundles. Recall that handlebodies are irreducible.

Suppose that F is a compact connected surface with at least one boundary component. Let T be the orientation I -bundle over F . If F is orientable then $T \cong F \times I$. If F is not orientable then T is the unique I -bundle over F with orientable total space. We call T the *I-bundle* and F the *base space*. Let $\rho_F: T \rightarrow F$ be the associated bundle map. Note that T is homeomorphic to a handlebody.

If $A \subset T$ is a union of fibres of the map ρ_F then A is *vertical* with respect to T . In particular take $\partial_v T = \rho_F^{-1}(\partial F)$ to be the *vertical boundary* of T . Take $\partial_h T$ to be the union of the boundaries of all of the fibres: this is the *horizontal boundary* of T . Note that $\partial_h T$ is always incompressible in T while $\partial_v T$ is incompressible in T as long as F is not homeomorphic to a disk.

Note that, as $|\partial_v T| \geq 1$, any vertical surface in T can be boundary compressed. However no vertical surface in T may be boundary compressed into $\partial_h T$.

We end this section with:

Lemma 8.5. *Suppose that F is a compact, connected surface with $\partial F \neq \emptyset$. Let $\rho_F: T \rightarrow F$ be the orientation I -bundle over F . Let X be a component of $\partial_h T$. Let $D \subset T$ be a properly embedded disk. If*

- ∂D is essential in ∂T ,
- ∂D and ∂X are tight, and
- D cannot be boundary compressed into X

then D may be properly isotoped to be vertical with respect to T . \square

9. HOLES FOR THE DISK COMPLEX

Here we begin to classify the holes for the disk complex, a more difficult analysis than that of the arc complex. To fix notation let V be a handlebody. Let $S = S_g = \partial V$. Recall that there is a natural inclusion $\nu: \mathcal{D}(V) \rightarrow \mathcal{C}(S)$.

Remark 9.1. The notion of a hole $X \subset \partial V$ for $\mathcal{D}(V)$ may be phrased in several different ways:

- every essential disk $D \subset V$ cuts the surface X ,

- $\overline{S \setminus X}$ is incompressible in V , or
- X is *disk-busting* in V .

The classification of holes $X \subset S$ for $\mathcal{D}(V)$ breaks roughly into three cases: either X is an annulus, is compressible in V , or is incompressible in V . In each case we obtain a result:

Theorem 10.1. *Suppose X is a hole for $\mathcal{D}(V)$ and X is an annulus. Then the diameter of X is at most 5.*

Theorem 11.6. *Suppose X is a compressible hole for $\mathcal{D}(V)$ with diameter at least 15. Then there are a pair of essential disks $D, E \subset V$ so that*

- $\partial D, \partial E \subset X$ and
- ∂D and ∂E fill X .

Theorem 12.1. *Suppose X is an incompressible hole for $\mathcal{D}(V)$ with diameter at least 61. Then there is an I -bundle $\rho_F: T \rightarrow F$ embedded in V so that*

- $\partial_h T \subset S$,
- X is isotopic in S to a component of $\partial_h T$,
- some component of $\partial_v T$ is boundary parallel into S ,
- F supports a pseudo-Anosov map.

As a corollary of these theorems we have:

Corollary 9.2. *If X is hole for $\mathcal{D}(V)$ with diameter at least 61 then X has infinite diameter.*

Proof. If X is a hole with diameter at least 61 then either Theorem 11.6 or Theorem 12.1 applies.

If X is compressible then Dehn twists, in opposite directions, about the given disks D and E yields an automorphism $f: V \rightarrow V$ so that $f|_X$ is pseudo-Anosov. This follows from Thurston's construction [38]. By Lemma 2.6 the hole X has infinite diameter.

If X is incompressible then $X \subset \partial_h T$ where $\rho_F: T \rightarrow F$ is the given I -bundle. Let $f: F \rightarrow F$ be the given pseudo-Anosov map. So g , the suspension of f , gives a automorphism of V . Again it follows that the hole X has infinite diameter. \square

Applying Lemma 5.8 we find another corollary:

Theorem 9.3. *If ∂V contains a strict hole with diameter at least 61 then the inclusion $\nu: \mathcal{D}(V) \rightarrow \mathcal{C}(S)$ is not a quasi-isometric embedding.*

\square

10. HOLES FOR THE DISK COMPLEX – ANNULI

The proof of Theorem 10.1 occupies the rest of this section. This proof shares many features with the proofs of Theorems 11.6 and 12.1. However, the exceptional definition of $\mathcal{C}(S_{0,2})$ prevents a unified approach. Fix V , a handlebody.

Theorem 10.1. *Suppose X is a hole for $\mathcal{D}(V)$ and X is an annulus. Then the diameter of X is at most 5.*

We begin with:

Claim. For all $D \in \mathcal{D}(V)$, $|D \cap X| \geq 2$.

Proof. Since X is a hole, every disk cuts X . Since X is an annulus, let α be a core curve for X . If $|D \cap X| = 1$, then we may band sum parallel copies of D along an subarc of α . The resulting disk misses α , a contradiction. \square

Assume, to obtain a contradiction, that X has diameter at least 6. Suppose that $D \in \mathcal{D}(V)$ is a disk chosen to minimize $|D \cap X|$. Among all disks $E \in \mathcal{D}(V)$ with $d_X(D, E) \geq 3$ choose one which minimizes $|D \cap E|$. Isotope D and E to make the boundaries tight and also tight with respect to ∂X . Tightening triples of curves is not canonical; nonetheless there is a tightening so that $S \setminus (\partial D \cup \partial E \cup X)$ contains no triangles. See Figure 1.




FIGURE 1. Triangles outside of X (see the left side) can be moved in (see the right side). This decreases the number of points of $D \cap E \cap (S \setminus X)$.

After this tightening we have:

Claim. Every arc of $\partial D \cap X$ meets every arc of $\partial E \cap X$ at least once.

Proof. Fix components arcs $\alpha \subset D \cap X$ and $\beta \subset E \cap X$. Let α', β' be the corresponding arcs in S^X the annular cover of S corresponding to

X . After the tightening we find that

$$|\alpha \cap \beta| \geq |\alpha' \cap \beta'| - 1.$$

Since $d_X(D, E) \geq 3$ Equation 2.3 implies that $|\alpha' \cap \beta'| \geq 2$. Thus $|\alpha \cap \beta| \geq 1$, as desired. \square

Claim. There is an outermost bigon $B \subset E \setminus D$ with the following properties:

- $\partial B = \alpha \cup \beta$ where $\alpha = B \cap D$, $\beta = \partial B \setminus \alpha \subset \partial E$,
- $\partial\alpha = \partial\beta \subset X$, and
- $|\beta \cap X| = 2$.

Furthermore, $|D \cap X| = 2$.

See the lower right of Figure 2 for a picture.

Proof. Consider the intersection of D and E , thought of as a collection of arcs and curves in E . Any simple closed curve component of $D \cap E$ can be removed by an isotopy of E , fixed on the boundary. (This follows from the irreducibility of V and an innermost disk argument.) Since we have assumed that $|D \cap E|$ is minimal it follows that there are no simple closed curves in $D \cap E$.

So consider any outermost bigon $B \subset E \setminus D$. Let $\alpha = B \cap D$. Let $\beta = \partial B \setminus \alpha = B \cap \partial V$. Note that β cannot completely contain a component of $E \cap X$ as this would contradict either the fact that B is outermost or the claim that every arc of $E \cap X$ meets some arc of $D \cap X$. Using this observation, Figure 2 lists the possible ways for B to lie inside of E .




FIGURE 2. The arc α cuts a bigon B off of E . The darker part of ∂E are the arcs of $E \cap X$. Either β is disjoint from X , β is contained in X , β meets X in a single subarc, or β meets X in two subarcs.

Let D' and D'' be the two essential disks obtained by boundary compressing D along the bigon B . Suppose α is as shown in one of the

first three pictures of Figure 2. It follows that either D' or D'' has, after tightening, smaller intersection with X than D does, a contradiction. We deduce that α is as pictured in lower right of Figure 2.

Boundary compressing D along B still gives disks $D', D'' \in \mathcal{D}(V)$. As these cannot have smaller intersection with X we deduce that $|D \cap X| \leq 2$ and the claim holds. \square

Using the same notation as in the proof above, let B be an outermost bigon of $E \setminus D$. We now study how $\alpha \subset \partial B$ lies inside of D .

Claim. The arc $\alpha \subset D$ connects distinct components of $D \cap X$.

Proof. Suppose not. Then there is a bigon $C \subset D \setminus \alpha$ with $\partial C = \alpha \cup \gamma$ and $\gamma \subset \partial D \cap X$. The disk $C \cup B$ is essential and intersects X at most once after tightening, contradicting our first claim. \square

We finish the proof of Theorem 10.1 by noting that $D \cup B$ is homeomorphic to $\Upsilon \times I$ where Υ is the simplicial tree with three edges and three leaves. We may choose the homeomorphism so that $(D \cup B) \cap X = \Upsilon \times \partial I$. It follows that we may properly isotope $D \cup B$ until $(D \cup B) \cap X$ is a pair of arcs. Recall that D' and D'' are the disks obtained by boundary compressing D along B . It follows that one of D' or D'' (or both) meets X in at most a single arc, contradicting our first claim. \square

11. HOLES FOR THE DISK COMPLEX – COMPRESSIBLE

The proof of Theorem 11.6 occupies the second half of this section.

11.1. Compression sequences of essential disks. Fix a multicurve $\Gamma \subset S = \partial V$. Fix also an essential disk $D \subset V$. Properly isotope D to make ∂D tight with respect to Γ .

If $D \cap \Gamma \neq \emptyset$ we may define:

Definition 11.1. A *compression sequence* $\{\Delta_k\}_{k=1}^n$ starting at D has $\Delta_1 = \{D\}$ and Δ_{k+1} is obtained from Δ_k via a boundary compression, disjoint from Γ , and tightening. Note that Δ_k is a collection of exactly k pairwise disjoint disks properly embedded in V . We further require, for $k \leq n$, that every disk of Δ_k meets some component of Γ . We call a compression sequence *maximal* if either

- no disk of Δ_n can be boundary compressed into $S \setminus \Gamma$ or
- there is a component $Z \subset S \setminus \Gamma$ and a boundary compression of Δ_n into $S \setminus \Gamma$ yielding an essential disk E with $\partial E \subset Z$.

We say that such maximal sequences *end essentially* or *end in Z* , respectively.

All compression sequences must end, by Remark 8.3. Given a maximal sequence we may relate the various disks in the sequence as follows:

Definition 11.2. Fix X , a component of $S \setminus \Gamma$. Fix $D_k \in \Delta_k$. A *disjointness pair* for D_k is an ordered pair (α, β) of essential arcs in X where

- $\alpha \subset D_k \cap X$,
- $\beta \subset \Delta_n \cap X$, and
- $d_{\mathcal{A}}(\alpha, \beta) \leq 1$.

If $\alpha \neq \alpha'$ then the two disjointness pairs (α, β) and (α', β) are distinct, even if α is properly isotopic to α' . A similar remark holds for the second coordinate.

The following lemma controls how subsurface projection distance changes in maximal sequences.

Lemma 11.3. *Fix a multicurve $\Gamma \subset S$. Suppose that D cuts Γ and choose a maximal sequence starting at D . Fix any component $X \subset S \setminus \Gamma$. Fix any disk $D_k \in \Delta_k$. Then either $D_k \in \Delta_n$ or there are four distinct disjointness pairs $\{(\alpha_i, \beta_i)\}_{i=1}^4$ for D_k in X where each of the arcs $\{\alpha_i\}$ appears as the first coordinate of at most two pairs.*

Proof. We induct on $n - k$. If D_k is contained in Δ_n there is nothing to prove. If D_k is contained in Δ_{k+1} we are done by induction. Thus we may assume that D_k is the disk of Δ_k which is boundary compressed at stage k . Let $D_{k+1}, D'_{k+1} \in \Delta_{k+1}$ be the two disks obtained after boundary compressing D_k along the bigon B . See Figure 3 for a picture of the pair of pants cobounded by ∂D_k and $\partial D_{k+1} \cup \partial D'_{k+1}$.




FIGURE 3. All arcs connecting D_k to itself or to $D_{k+1} \cup D'_{k+1}$ are arcs of $\Gamma \cap P$. The boundary compressing arc $B \cup S$ meets D_k twice and is parallel to the vertical arcs of $\Gamma \cap P$.

Let δ be a band sum arc dual to B (the dotted arc in Figure 3). We may assume that $|\Gamma \cap \delta|$ is minimal over all arcs dual to B . It follows that

the band sum of D_{k+1} with D'_{k+1} along δ is tight, without any isotopy. (This is where we use the fact that B is a boundary compression *in the complement of Γ* , as opposed to being a general boundary compression of D_k in V .)

There are now three possibilities: neither, one, or both points of $\partial\delta$ are contained in X .

First suppose that $X \cap \partial\delta = \emptyset$. Then every arc of $D_{k+1} \cap X$ is parallel to an arc of $D_k \cap X$, and similarly for D'_{k+1} . If D_{k+1} and D'_{k+1} are both components of Δ_n then choose any arcs β, β' of $D_{k+1} \cap X$ and of $D'_{k+1} \cap X$. Let α, α' be the parallel components of $D_k \cap X$. The four disjointness pairs are then (α, β) , (α, β') , (α', β) , (α', β') . Suppose instead that D_{k+1} is not a component of Δ_n . Then D_k inherits four disjointness pairs from D_{k+1} .

Second suppose that exactly one endpoint $x \in \partial\delta$ meets X . Let $\gamma \subset D_{k+1}$ be the component of $D_{k+1} \cap X$ containing x . Let X' be the component of $X \cap P$ that contains x and let α, α' be the two components of $D_k \cap X'$. Let β be any arc of $D'_{k+1} \cap X$.

If $D_{k+1} \notin \Delta_n$ and γ is not the first coordinate of one of D_{k+1} 's four pairs then D_k inherits disjointness pairs from D_{k+1} . If $D'_{k+1} \notin \Delta_n$ then D_k inherits disjointness pairs from D'_{k+1} .

Thus we may assume that both D_{k+1} and D'_{k+1} are in Δ_n or that only $D'_{k+1} \in \Delta_n$ while γ appears as the first arc of disjointness pair for D_{k+1} . In case of the former the required disjointness pairs are (α, β) , (α', β) , (α, γ) , and (α', γ) . In case of the latter we do not know if γ is allowed to appear as the second coordinate of a pair. However we are given four disjointness pairs for D_{k+1} and are told that γ appears as the first coordinate of at most two of these pairs. Hence the other two pairs are inherited by D_k . The pairs (α, β) and (α', β) give the desired conclusion.

Third suppose that the endpoints of δ meet $\gamma \subset D_{k+1}$ and $\gamma' \subset D'_{k+1}$. Let X' be a component of $X \cap P$ containing γ . Let α and α' be the two arcs of $D_k \cap X'$. Suppose both D_{k+1} and D'_{k+1} lie in Δ_n . Then the desired pairs are (α, γ) , (α', γ) , (α, γ') , and (α', γ') . If $D'_{k+1} \in \Delta_n$ while D_{k+1} is not then D_k inherits two pairs from D_{k+1} . We add to these the pairs (α, γ') , and (α', γ') . If neither disk lies in Δ_n then D_k inherits two pairs from each disk and the proof is complete. \square

Given a disk $D \in \mathcal{D}(V)$ and a hole $X \subset S$ our Lemma 11.3 allows us to adapt D to X .

Lemma 11.4. *Fix a hole $X \subset S$ for $\mathcal{D}(V)$. For any disk $D \in \mathcal{D}(V)$ there is a disk D' with the following properties:*

- ∂X and $\partial D'$ are tight.

- If X is incompressible then D' is not boundary compressible into X and $d_{\mathcal{A}}(D, D') \leq 3$.
- If X is compressible then $\partial D' \subset X$ and $d_{\mathcal{AC}}(D, D') \leq 3$.

Here $\mathcal{A} = \mathcal{A}(X)$ and $\mathcal{AC} = \mathcal{AC}(X)$.

Proof. If $\partial D \subset X$ then the lemma is trivial. So assume, by Remark 9.1, that D cuts ∂X . Choose a maximal sequence with respect to ∂X starting at D .

Suppose that the sequence is non-trivial ($n > 1$). By Lemma 11.3 there is a disk $E \in \Delta_n$ so that $D \cap X$ and $E \cap X$ contain disjoint arcs.

If the sequence ends essentially then choose $D' = E$ and the lemma is proved. If the sequence ends in X then there is a boundary compression of Δ_n , disjoint from ∂X , yielding the desired disk D' with $\partial D' \subset X$. Since $E \cap D' = \emptyset$ we again obtain the desired bound.

Assume now that the sequence is trivial ($n = 1$). Then take $E = D \in \Delta_n$ and the proof is identical to that of the previous paragraph. \square

Remark 11.5. Lemma 11.4 is unexpected: after all, any pair of curves in $\mathcal{C}(X)$ can be connected by a sequence of band sums. Thus arbitrary band sums can change the subsurface projection to X . However, the sequences of band sums arising in Lemma 11.4 are very special. Firstly they do not cross ∂X and secondly they are “tree-like” due to the fact every arc in D is separating.

When D is replaced by a surface with genus then Lemma 11.4 does not hold in general; this is a fundamental observation due to Kobayashi [21] (see also [17]). Namazi points out that even if D is only replaced by a planar surface Lemma 11.4 does not hold in general.

11.2. Proving the theorem. We now prove:

Theorem 11.6. *Suppose X is a compressible hole for $\mathcal{D}(V)$ with diameter at least 15. Then there are a pair of essential disks $D, E \subset V$ so that*

- $\partial D, \partial E \subset X$ and
- ∂D and ∂E fill X .

Proof. Choose disks D' and E' in $\mathcal{D}(V)$ so that $d_X(D', E') \geq 15$. By Lemma 11.4 there are disks D and E so that $\partial D, \partial E \subset X$, $d_X(D', D) \leq 6$, and $d_X(E', E) \leq 6$. It follows from the triangle inequality that $d_X(D, E) \geq 3$. \square

12. HOLES FOR THE DISK COMPLEX – INCOMPRESSIBLE

This section classifies incompressible holes for the disk complex.

Theorem 12.1. *Suppose X is an incompressible hole for $\mathcal{D}(V)$ with diameter at least 61. Then there is an I -bundle $\rho_F: T \rightarrow F$ embedded in V so that*

- $\partial_h T \subset \partial V$,
- X is a component of $\partial_h T$,
- some component of $\partial_v T$ is boundary parallel into ∂V ,
- F supports a pseudo-Anosov map.

Here is a short plan of the proof: We are given X , an incompressible hole for $\mathcal{D}(V)$. Following Lemma 11.4 we may assume that D, E are essential disks, without boundary compressions into X or $S \setminus X$, with $d_X(D, E) > 43$. Examine the intersection pattern of D and E to find two families of rectangles \mathcal{R} and \mathcal{Q} . The intersection pattern of these rectangles in V will determine the desired I -bundle T . The third conclusion of the theorem follows from standard facts about primitive annuli. The fourth requires another application of Lemma 11.4 as well as Lemma 2.6.

12.1. Diagonals of polygons. To understand the intersection pattern of D and E we discuss diagonals of polygons. Let D be a $2n$ sided regular polygon. Label the sides of D with the letters X and Y in alternating fashion. Any side labeled X (or Y) will be called an X side (or Y side).

Definition 12.2. An arc γ properly embedded in D is a *diagonal* if the points of $\partial\gamma$ lie in the interiors of distinct sides of D . If γ and γ' are diagonals for D which together meet three different sides then γ and γ' are *non-parallel*.

Lemma 12.3. *Suppose that $\Gamma \subset D$ is a collection of pairwise disjoint non-parallel diagonals. Then there is an X side of D meeting at most eight diagonals of Γ .*

Proof. A counting argument shows that $|\Gamma| \leq 4n - 3$. If every X side meets at least nine non-parallel diagonals then $|\Gamma| \geq \frac{9}{2}n > 4n - 3$, a contradiction. \square

12.2. Improving disks. Suppose now that X is an incompressible hole for $\mathcal{D}(V)$ with diameter at least 61. Note that, by Theorem 10.1, X is not an annulus. Let $Y = \overline{S \setminus X}$.

Choose disks D' and E' in V so that $d_X(D', E') \geq 61$. By Lemma 11.4 there are a pair of disks D and E so that both are essential in V , cannot be boundary compressed into X or into Y , and so that $d_{\mathcal{A}(X)}(D', D) \leq 3$ and $d_{\mathcal{A}(X)}(E', E) \leq 3$. Thus $d_X(D', D) \leq 9$ and $d_X(E', E) \leq 9$ (Lemma 5.7). By the triangle inequality $d_X(D, E) \geq 61 - 18 = 43$.

Recall, as well, that ∂D and ∂E are tight with respect to ∂X . We may further assume that ∂D and ∂E are tight with respect to each other. Also, minimize the quantities $|X \cap (\partial D \cap \partial E)|$ and $|D \cap E|$ while keeping everything tight. In particular, there are no triangle components of $\partial V \setminus (D \cup E \cup \partial X)$. Now consider D and E to be even-sided polygons, with vertices being the points $\partial D \cap \partial X$ and $\partial E \cap \partial X$ respectively. Let $\Gamma = D \cap E$. See Figure 4 for one *a priori* possible collection $\Gamma \subset D$.




FIGURE 4. In fact, $\Gamma \subset D$ cannot contain simple closed curves or non-diagonals.

From our assumptions and the irreducibility of V it follows that Γ contains no simple closed curves. Suppose now that there is a $\gamma \subset \Gamma$ so that, in D , both endpoints of γ lie in the same side of D . Then there is an outermost such arc, say $\gamma' \subset \Gamma$, cutting a bigon B out of D . It follows that B is a boundary compression of E which is disjoint from ∂X . But this contradicts the construction of E . We deduce that all arcs of Γ are diagonals for D and, via a similar argument, for E .

Let $\alpha \subset D \cap X$ be an X side of D meeting at most eight distinct types of diagonal of Γ . Choose $\beta \subset E \cap X$ similarly. As $d_X(D, E) \geq 43$ we have that $d_X(\alpha, \beta) \geq 43 - 6 = 37$.

Now break each of α and β into at most eight subarcs $\{\alpha_i\}$ and $\{\beta_j\}$ so that each subarc meets all of the diagonals of fixed type and only of that type. Let $R_i \subset D$ be the rectangle with upper boundary α_i and containing all of the diagonals meeting α_i . Let α'_i be the lower boundary of R_i . Define Q_j and β'_j similarly. See Figure 5 for a picture of R_i .

Call an arc α_i *large* if there is an arc β_j so that $|\alpha_i \cap \beta_j| \geq 3$. We use the same notation for β_j . Let Θ be the union of all of the large α_i and β_j . Thus Θ is a four-valent graph in X . Let Θ' be the union of the corresponding large α'_i and β'_j .

Claim 12.4. The graph Θ is non-empty.




FIGURE 5. The rectangle $R_i \subset D$ is surrounded by the dotted line. The arc α_i in $\partial D \cap X$ is indicated. In general the arc α'_i may lie in X or in Y .

Proof. If $\Theta = \emptyset$, then all α_i are small. It follows that $|\alpha \cap \beta| \leq 128$ and thus $d_X(\alpha, \beta) \leq 16$, by Lemma 2.2. As $d_X(\alpha, \beta) \geq 37$ this is a contradiction. \square

Let $Z \subset \partial V$ be a small regular neighborhood of Θ and define Z' similarly.

Claim 12.5. No component of Θ or of Θ' is contained in a disk $D \subset \partial V$. No component of Θ or of Θ' is contained in an annulus $A \subset \partial V$ that is peripheral in X .

Proof. For a contradiction suppose that W is a component of Z contained in a disk. Then there is some pair α_i, β_j having a bigon in ∂V . This contradicts the tightness of ∂D and ∂E . The same holds for Z' .

Suppose now that some component W is contained in an annulus A , peripheral in X . Thus W fills A . Suppose that α_i and β_j are large and contained in W . By the classification of arcs in A we deduce that either α_i and β_j form a bigon in A or ∂X , α_i and β_j form a triangle. Either conclusion gives a contradiction. \square

Claim 12.6. The graph Θ fills X .

Proof. Suppose not. Fix attention on any component $W \subset Z$. Since Θ does not fill, the previous claim implies that there is a component $\gamma \subset \partial W$ that is essential and non-peripheral in X . Note that any large α_i meets ∂W in at most two points, while any small α_i meets ∂W in at most 32 points. Thus $|\alpha \cap \partial W| \leq 256$ and the same holds for β . Thus $d_X(\alpha, \beta) \leq 36$ by the triangle inequality. As $d_X(\alpha, \beta) \geq 37$ this is a contradiction. \square

The previous two claims imply:

Claim 12.7. The graph Θ is connected. \square

There are now two possibilities: either $\Theta \cap \Theta'$ is empty or not. In the first case set $\Sigma = \Theta$ and in the second set $\Sigma = \Theta \cup \Theta'$. By the claims above, Σ is connected and fills X . Let $\mathcal{R} = \{R_i\}$ and $\mathcal{Q} = \{Q_j\}$ be the collections of large rectangles.

12.3. Building the I -bundle. We are given Σ , \mathcal{R} and \mathcal{Q} as above. Note that $\mathcal{R} \cup \mathcal{Q}$ is an I -bundle and Σ is the component of its horizontal boundary meeting X . See Figure 6 for a simple case.




FIGURE 6. $\mathcal{R} \cup \mathcal{Q}$ is an I -bundle: all arcs of intersection are parallel.

Let T_0 be a regular neighborhood of $\mathcal{R} \cup \mathcal{Q}$, taken in V . Again T_0 has the structure of an I -bundle. Note that $\partial_h T_0 \subset \partial V$, $\partial_h T_0 \cap X$ is a component of $\partial_h T_0$, and this component fills X due to Claim 12.6. We will enlarge T_0 to obtain the correct I -bundle in V .

Begin by enumerating all annuli $\{A_i\} \subset \partial_v T_0$ with the property that some component of ∂A_i is inessential in ∂V . Suppose that we have built the I -bundle T_i and are now considering the annulus $A = A_i$. Let $\gamma \cup \gamma' = \partial A \subset \partial V$ with γ inessential in ∂V . Let $B \subset \partial V$ be the disk which γ bounds. By induction we assume that no component of $\partial_h T_i$ is contained in a disk embedded in ∂V (the base case holds by Claim 12.5). It follows that $B \cap T_i = \partial B = \gamma$. Thus $B \cup A$ is isotopic, rel γ' , to be a properly embedded disk $B' \subset V$. As γ' lies in X or Y , both incompressible, γ' must bound a disk $C \subset \partial V$. Note that $C \cap T_i = \partial C = \gamma'$, again using the induction hypothesis.

It follows that $B \cup A \cup C$ is an embedded two-sphere in V . As V is a handlebody V is irreducible. Thus $B \cup A \cup C$ bounds a three-ball U_i in V . Choose a homeomorphism $U_i \cong B \times I$ so that B is identified with $B \times \{0\}$, C is identified with $B \times \{1\}$, and A is identified with $\partial B \times I$. We form $T_{i+1} = T_i \cup U_i$ and note that T_{i+1} still has the structure of an I -bundle. Recalling that $A = A_i$ we have $\partial_v T_{i+1} = \partial_v T_i \setminus A_i$. Also

$\partial_h T_{i+1} = \partial_h T_i \cup (B \cup C) \subset \partial V$. It follows that no component of $\partial_h T_{i+1}$ is contained in a disk embedded in ∂V . Similarly, $\partial_h T_{i+1} \cap X$ is a component of $\partial_h T_{i+1}$ and this component fills X .

After dealing with all of the annuli $\{A_i\}$ in this fashion we are left with an I -bundle T . Now all components of $\partial \partial_v T$ [sic] are essential in ∂V . All of these lying in X are peripheral in X . This is because they are disjoint from $\Sigma \subset \partial_h T$, which fills X , by induction. It follows that the component of $\partial_h T$ containing Σ is isotopic to X .

This finishes the construction of the promised I -bundle T and demonstrates the first two conclusions of Theorem 12.1. For future use we record:

Remark 12.8. Every curve of $\partial \partial_v T = \partial \partial_h T$ is essential in $S = \partial V$.

12.4. A vertical annulus parallel into the boundary. Here we obtain the third conclusion of Theorem 12.1: at least one component of $\partial_v T$ is boundary parallel in ∂V .

Fix T an I -bundle with the incompressible hole X a component of $\partial_h T$.

Claim 12.9. All components of $\partial_v T$ are incompressible in V .

Proof. Suppose that $A \subset \partial_v T$ was compressible. By Remark 12.8 we may compress A to obtain a pair of essential disks B and C . Note that ∂B is isotopic into the complement of $\partial_h T$. So $\overline{S \setminus X}$ is compressible, contradicting Remark 9.1. \square

Claim 12.10. Some component of $\partial_v T$ is boundary parallel.

Proof. Since $\partial_v T$ is incompressible (Claim 12.9) by Remark 8.2, we find that $\partial_v T$ is boundary compressible in V . Let B be a boundary compression for $\partial_v T$. Let A be the component of $\partial_v T$ meeting B . Let α denote the arc $A \cap B$.

The arc α is either essential or inessential in A . Suppose α is inessential in A . Then α cuts a bigon, C , out of A . Since B was a boundary compression the disk $D = B \cup C$ is essential in V . Since B meets $\partial_v T$ in a single arc, either $D \subset T$ or $D \subset \overline{V \setminus T}$. The former implies that $\partial_h T$ is compressible and the latter that X is not a hole. Either gives a contradiction.

It follows that α is essential in A . Now carefully boundary compress A : Let N be the closure of a regular neighborhood of B , taken in $V \setminus A$. Let A' be the closure of $A \setminus N$ (so A' is a rectangle). Let $B' \cup B''$ be the closure of $\text{fr}(N) \setminus A$. Both B' and B'' are bigons, parallel to B . Form $D = A' \cup B' \cup B''$: a properly embedded disk in V . If D is essential then, as above, either $D \subset T$ or $D \subset \overline{V \setminus T}$. Again, either gives a contradiction.

It follows that D is inessential in V . Thus D cuts a closed three-ball U out of V . There are two final cases: either $N \subset U$ or $N \cap U = B' \cup B''$. If U contains N then U contains A . Thus ∂A is contained in the disk $U \cap \partial V$. This contradicts Remark 12.8. Deduce instead that $W = U \cup N$ is a solid torus with meridional disk B . Thus W gives a parallelism between A and the annulus $\partial V \cap \partial W$, as desired. \square

Remark 12.11. Similar considerations prove that the multicurve

$$\{\partial A \mid A \text{ is a boundary parallel component of } \partial_v T\}$$

is disk-busting for V .

12.5. Finding a pseudo-Anosov map. Here we prove that the base surface F of the I -bundle T admits a pseudo-Anosov map.

As in Section 12.2, pick essential disks D' and E' in V so that $d_X(D', E') \geq 61$. Lemma 11.4 provides disks D and E which cannot be boundary compressed into X or into $\overline{S \setminus X}$ – thus D and E cannot be boundary compressed into $\partial_h T$. Also, as above, $d_X(D, E) \geq 61 - 18 = 43$.

After isotoping D to minimize intersection with $\partial_v T$ it must be the case that all components of $D \cap \partial_v T$ are essential arcs in $\partial_v T$. By Lemma 8.5 we conclude that D may be isotoped in V so that $D \cap T$ is vertical in T . The same holds of E . Choose A and B , components of $D \cap T$ and $E \cap T$. Each are vertical rectangles. Since $\text{diam}_X(\pi_X(D)) \leq 3$ (Lemma 4.4) we now have $d_X(A, B) \geq 43 - 6 = 37$.

We now begin to work in the base surface F . Recall that $\rho_F: T \rightarrow F$ is an I -bundle. Take $\alpha = \rho_F(A)$ and $\beta = \rho_F(B)$. Note that the natural map $\mathcal{C}(F) \rightarrow \mathcal{C}(X)$, defined by taking a curve to its lift, is distance non-increasing (see Equation 6.5). Thus $d_F(\alpha, \beta) \geq 37$. By Theorem 10.1 the surface F cannot be an annulus. Thus, by Lemma 2.6 the subsurface F supports a pseudo-Anosov map and we are done.

12.6. Corollaries. We now deal with the possibility of disjoint holes for the disk complex.

Lemma 12.12. *Suppose that X is a large incompressible hole for $\mathcal{D}(V)$ supported by the I -bundle $\rho_F: T \rightarrow F$. Let $Y = \partial_h T \setminus X$. Let $\tau: \partial_h T \rightarrow \partial_h T$ be the involution switching the ends of the I -fibres. Suppose that $D \in \mathcal{D}(V)$ is an essential disk.*

- *If F is orientable then $d_{\mathcal{A}(F)}(D \cap X, D \cap Y) \leq 6$.*
- *If F is non-orientable then $d_X(D, \mathcal{C}^\tau(X)) \leq 3$.*

Proof. By Lemma 11.4 there is a disk $D' \subset V$ which is tight with respect to $\partial_h T$ and which cannot be boundary compressed into $\partial_h T$

(or into the complement). Also, for any component $Z \subset \partial_h T$ we have $d_{\mathcal{A}(Z)}(D, D') \leq 3$.

Properly isotope D' to minimize $D' \cap \partial_v T$. Then $D' \cap \partial_v T$ is properly isotopic, in $\partial_v T$, to a collection of vertical arcs. Let $E \subset D' \cap T$ be a component. Lemma 8.5 implies that E is vertical in T , after an isotopy of D' preserving $\partial_h T$ setwise. Since E is vertical, the arcs $E \cap \partial_h T \subset D'$ are τ -invariant. The conclusion follows. \square

Recall Lemma 7.3: all holes for the arc complex intersect. This cannot hold for the disk complex. For example if $\rho_F: T \rightarrow F$ is an I -bundle over an orientable surface then take $V = T$ and notice that both components of $\partial_h T$ are holes for $\mathcal{D}(V)$. However, by the first conclusion of Lemma 12.12, X and Y are paired holes, in the sense of Definition 5.5. So, as with the invariant arc complex (Lemma 7.5), all holes for the disk complex interfere:

Lemma 12.13. *Suppose that $X, Z \subset \partial V$ are large holes for $\mathcal{D}(V)$. If $X \cap Z = \emptyset$ then there is an I -bundle $T \cong F \times I$ in V so that $\partial_h T = X \cup Y$ and $Y \cap Z \neq \emptyset$.*

Proof. Suppose that $X \cap Z = \emptyset$. It follows from Remark 9.1 that both X and Z are incompressible. Let $\rho_F: T \rightarrow F$ be the I -bundle in V with $X \subset \partial_h T$, as provided by Theorem 12.1. We also have a component $A \subset \partial_v T$ so that A is boundary parallel. Let U be the solid torus component of $V \setminus A$. Note that Z cannot be contained in $\partial U \setminus A$ because Z is not an annulus (Theorem 10.1).

Let $\alpha = \rho_F(A)$. Choose any essential arc $\delta \subset F$ with both endpoints in $\alpha \subset \partial F$. It follows that $\rho_F^{-1}(\delta)$, together with two meridional disks of U , forms an essential disk D in V . Let $W = \partial_h T \cup (U \setminus A)$ and note that $\partial D \subset W$.

If F is non-orientable then $Z \cap W = \emptyset$ and we have a contradiction. Deduce that F is orientable. Now, if Z misses Y then Z misses W and we again have a contradiction. It follows that Z cuts Y and we are done. \square

13. AXIOMS FOR COMBINATORIAL COMPLEXES

The goal of this section and the next is to prove, inductively, an upper bound on distance in a combinatorial complex $\mathcal{G}(S) = \mathcal{G}$. This section presents our axioms on \mathcal{G} : sufficient hypotheses for Theorem 13.1. The axioms, apart from Axiom 13.2, are quite general. Axiom 13.2 is necessary to prove hyperbolicity and greatly simplifies the recursive construction in Section 14.

Theorem 13.1. *Fix S a compact connected non-simple surface. Suppose that $\mathcal{G} = \mathcal{G}(S)$ is a combinatorial complex satisfying the axioms of Section 13. Let X be a hole for \mathcal{G} and suppose that $\alpha_X, \beta_X \in \mathcal{G}$ are contained in X . For any constant $c > 0$ there is a constant A satisfying:*

$$d_{\mathcal{G}}(\alpha_X, \beta_X) \leq A \sum [d_Y(\alpha_X, \beta_X)]_c$$

where the sum is taken over all holes $Y \subseteq X$ for \mathcal{G} .

The proof of the upper bound is more difficult than that of the lower bound, Theorem 5.10. This is because naturally occurring paths in \mathcal{G} between α_X and β_X may waste time in non-holes. The first example of this is the path in $\mathcal{C}(S)$ obtained by taking the short curves along a Teichmüller geodesic. The Teichmüller geodesic may spend time rearranging the geometry of a subsurface. Then the systole path in the curve complex must be much longer than the curve complex distance between the endpoints.

In Sections 16, 17, 19 we will verify these axioms for the curve complex of a non-orientable surface, the arc complex, and the disk complex.

13.1. The axioms. Suppose that $\mathcal{G} = \mathcal{G}(S)$ is a combinatorial complex. We begin with the axiom required for hyperbolicity.

Axiom 13.2 (Holes interfere). *All large holes for \mathcal{G} interfere, as given in Definition 5.6.*

Fix vertices $\alpha_X, \beta_X \in \mathcal{G}$, both contained in a hole X . We are given $\Lambda = \{\mu_n\}_{n=0}^N$, a path of markings in X .

Axiom 13.3 (Marking path). *We require:*

- (1) *The support of μ_{n+1} is contained inside the support of μ_n .*
- (2) *For any subsurface $Y \subseteq X$, if $\pi_Y(\mu_k) \neq \emptyset$ then for all $n \leq k$ the map $n \mapsto \pi_Y(\mu_n)$ is an unparameterized quasi-geodesic with constants depending only on \mathcal{G} .*

The second condition is crucial and often technically difficult to obtain.

We are given, for every essential subsurface $Y \subset X$, a perhaps empty interval $J_Y \subset [0, N]$ with the following properties.

Axiom 13.4 (Accessibility). *The interval for X is $J_X = [0, N]$. There is a constant B_3 so that*

- (1) *If $m \in J_Y$ then Y is contained in the support of μ_m .*
- (2) *If $m \in J_Y$ then $\iota(\partial Y, \mu_m) < B_3$.*
- (3) *If $[m, n] \cap J_Y = \emptyset$ then $d_Y(\mu_m, \mu_n) < B_3$.*

There is a combinatorial path $\Gamma = \{\gamma_i\}_{i=0}^K \subset \mathcal{G}$ starting with α_X ending with β_X and each γ_i is contained in X . There is a strictly increasing reindexing function $r: [0, K] \rightarrow [0, N]$ with $r(0) = 0$ and $r(K) = N$.

Axiom 13.5 (Combinatorial). *There is a constant C_2 so that:*

- $d_Y(\gamma_i, \mu_{r(i)}) < C_2$, for every $i \in [0, K]$ and every hole $Y \subset X$,
- $d_{\mathcal{G}}(\gamma_i, \gamma_{i+1}) < C_2$, for every $i \in [0, K - 1]$.

Axiom 13.6 (Replacement). *There is a constant C_4 so that:*

- (1) *If $Y \subset X$ is a hole and $r(i) \in J_Y$ then there is a vertex $\gamma' \in \mathcal{G}$ so that γ' is contained in Y and $d_{\mathcal{G}}(\gamma_i, \gamma') < C_4$.*
- (2) *If $Z \subset X$ is a non-hole and $r(i) \in J_Z$ then there is a vertex $\gamma' \in \mathcal{G}$ so that $d_{\mathcal{G}}(\gamma_i, \gamma') < C_4$ and so that γ' is contained in Z or in $X \setminus Z$.*

There is one axiom left: the axiom for straight intervals. This is given in the next subsection.

13.2. Inductive, electric, shortcut and straight intervals. We describe subintervals that arise in the partitioning of $[0, K]$. As discussed carefully in Section 13.3, we will choose a lower threshold $L_1(Y)$ for every essential $Y \subset X$ and a general upper threshold, L_2 .

Definition 13.7. Suppose that $[i, j] \subset [0, K]$ is a subinterval of the combinatorial path. Then $[i, j]$ is an *inductive interval* associated to a hole $Y \subsetneq X$ if

- $r([i, j]) \subset J_Y$ (for paired Y we require $r([i, j]) \subset J_Y \cap J_{Y'}$) and
- $d_Y(\gamma_i, \gamma_j) \geq L_1(Y)$.

When X is the only relevant hole we have a simpler definition:

Definition 13.8. Suppose that $[i, j] \subset [0, K]$ is a subinterval of the combinatorial path. Then $[i, j]$ is an *electric interval* if $d_Y(\gamma_i, \gamma_j) < L_2$ for all holes $Y \subsetneq X$.

Electric intervals will be further partitioned into shortcut and straight intervals.

Definition 13.9. Suppose that $[p, q] \subset [0, K]$ is a subinterval of the combinatorial path. Then $[p, q]$ is a *shortcut* if

- $d_Y(\gamma_p, \gamma_q) < L_2$ for all holes Y , including X itself, and
- there is a non-hole $Z \subset X$ so that $r([p, q]) \subset J_Z$.

Definition 13.10. Suppose that $[p, q] \subset [0, K]$ is a subinterval of the combinatorial path and is contained in an electric interval $[i, j]$. Then $[p, q]$ is a *straight interval* if $d_Y(\mu_{r(p)}, \mu_{r(q)}) < L_2$ for all non-holes Y .

Our final axiom is:

Axiom 13.11 (Straight). *There is a constant A depending only on X and \mathcal{G} so that for every straight interval $[p, q]$:*

$$d_{\mathcal{G}}(\gamma_p, \gamma_q) \leq_A d_X(\gamma_p, \gamma_q)$$

13.3. Deductions from the axioms. Axiom 13.3 and Lemma 3.9 imply that the reverse triangle inequality holds for projections of marking paths.

Lemma 13.12. *There is a constant C_1 so that*

$$d_Y(\mu_m, \mu_n) + d_Y(\mu_n, \mu_p) < d_Y(\mu_m, \mu_p) + C_1$$

for every essential $Y \subset X$ and for every $m < n < p$ in $[0, N]$. \square

We record three simple consequences of Axiom 13.4.

Lemma 13.13. *There is a constant C_3 , depending only on B_3 , with the follow properties:*

- (i) *If Y is strictly nested in Z and $m \in J_Y$ then $d_Z(\partial Y, \mu_m) \leq C_3$.*
- (ii) *If Y is strictly nested in Z then for any $m, n \in J_Y$, $d_Z(\mu_m, \mu_n) < C_3$.*
- (iii) *If Y and Z overlap then for any $m, n \in J_Y \cap J_Z$ we have $d_Y(\mu_m, \mu_n), d_Z(\mu_m, \mu_n) < C_3$.*

Proof. We first prove conclusion (i): Since Y is strictly nested in Z and since Y is contained in the support of μ_m (part (1) of Axiom 13.4), both ∂Y and μ_m cut Z . By Axiom 13.4, part (2), we have that $\iota(\partial Y, \mu_m) \leq B_3$. It follows that $\iota(\partial Y, \pi_Z(\mu_m)) \leq 2B_3$. By Lemma 2.2 we deduce that $d_Z(\partial Y, \mu_m) \leq 2 \log_2 B_3 + 3$. We take C_3 larger than this right hand side.

Conclusion (ii) follows from a pair of applications of conclusion (i) and the triangle inequality.

For conclusion (iii): As in (ii), to bound $d_Z(\mu_m, \mu_n)$ it suffices to note that ∂Y cuts Z and that ∂Y has bounded intersection with both of μ_m, μ_n . \square

We now have all of the constants C_1, C_3, C_2, C_4 in hand. Recall that L_4 is the pairing constant of Definition 5.5 and that M_0 is the constant of 4.6. We must choose a lower threshold $L_1(Y)$ for every essential $Y \subset X$. We must also choose the general upper threshold, L_2 and general lower threshold L_0 . We require, for all essential Z, Y in X , with

$\xi(Z) < \xi(Y) \leq \xi(X)$:

$$(13.14) \quad L_0 > C_3 + 2C_2 + 2L_4$$

$$(13.15) \quad L_2 > L_1(X) + 2L_4 + 6C_1 + 2C_2 + 14C_3 + 10$$

$$(13.16) \quad L_1(Y) > M_0 + 2C_3 + 4C_2 + 2L_4 + L_0$$

$$(13.17) \quad L_1(X) > L_1(Z) + 2C_3 + 4C_2 + 4L_4$$

14. PARTITION AND THE UPPER BOUND ON DISTANCE

In this section we prove Theorem 13.1 by induction on $\xi(X)$. The first stage of the proof is to describe the *inductive partition*: we partition the given interval $[0, K]$ into inductive and electric intervals. The inductive partition is closely linked with the hierarchy machine [25] and with the notion of antichains introduced in [34].

We next give the *electric partition*: each electric interval is divided into straight and shortcut intervals. Note that the electric partition also gives the base case of the induction. We finally bound $d_{\mathcal{G}}(\alpha_X, \beta_X)$ from above by combining the contributions from the various intervals.

14.1. Inductive partition. We begin by identifying the relevant surfaces for the construction of the partition. We are given a hole X for \mathcal{G} and vertices $\alpha_X, \beta_X \in \mathcal{G}$ contained in X . Define

$$B_X = \{Y \subsetneq X \mid Y \text{ is a hole and } d_Y(\alpha_X, \beta_X) \geq L_1(X)\}.$$

For any subinterval $[i, j] \subset [0, K]$ define

$$B_X(i, j) = \{Y \in B_X \mid d_Y(\gamma_i, \gamma_j) \geq L_1(X)\}.$$

We now partition $[0, K]$ into inductive and electric intervals. Begin with the partition of one part $\mathcal{P}_X = \{[0, K]\}$. Recursively \mathcal{P}_X is a partition of $[0, K]$ consisting of intervals which are either inductive, electric, or undetermined. Suppose that $[i, j] \in \mathcal{P}_X$ is undetermined.

Claim. If $B_X(i, j)$ is empty then $[i, j]$ is electric.

Proof. Since $B_X(i, j)$ is empty, every hole $Y \subsetneq X$ has either $d_Y(\gamma_i, \gamma_j) < L_1(X)$ or $Y \notin B_X$. In the former case, as $L_1(X) < L_2$, we are done.

So suppose the latter holds. Now, by the reverse triangle inequality (Lemma 13.12),

$$d_Y(\mu_{r(i)}, \mu_{r(j)}) < d_Y(\mu_0, \mu_N) + 2C_1.$$

Since $r(0) = 0$ and $r(K) = N$ we find:

$$d_Y(\gamma_i, \gamma_j) < d_Y(\alpha_X, \beta_X) + 2C_1 + 4C_2.$$

Deduce that

$$d_Y(\gamma_i, \gamma_j) < L_1(X) + 2C_1 + 4C_2 < L_2.$$

This completes the proof. \square

Thus if $B_X(i, j)$ is empty then $[i, j] \in \mathcal{P}_X$ is determined to be electric. Proceed on to the next undetermined element. Suppose instead that $B_X(i, j)$ is non-empty. Pick a hole $Y \in B_X(i, j)$ so that Y has maximal $\xi(Y)$ amongst the elements of $B_X(i, j)$.

Let $p, q \in [i, j]$ be the first and last indices, respectively, so that $r(p), r(q) \in J_Y$. (If Y is paired with Y' then we take the first and last indices that, after reindexing, lie inside of $J_Y \cap J_{Y'}$.)

Claim. The indices p, q are well-defined.

Proof. By assumption $d_Y(\gamma_i, \gamma_j) \geq L_1(X)$. By Equation 13.14,

$$L_1(X) > C_3 + 2C_2.$$

We deduce from Axiom 13.4 and Axiom 13.5 that $J_Y \cap r([i, j])$ is non-empty. Thus, if Y is not paired, the indices p, q are well-defined.

Suppose instead that Y is paired with Y' . Recall that measurements made in Y and Y' differ by at most the pairing constant L_4 given in Definition 5.5. By (13.16),

$$L_1(X) > C_3 + 2C_2 + 2L_4.$$

We deduce again from Axiom 13.4 that $J_{Y'} \cap r([i, j])$ is non-empty.

Suppose now, for a contradiction, that $J_Y \cap J_{Y'} \cap r([i, j])$ is empty. Define

$$h = \max\{\ell \in [i, j] \mid r(\ell) \in J_Y\}, \quad k = \min\{\ell \in [i, j] \mid r(\ell) \in J_{Y'}\}$$

Without loss of generality we may assume that $h < k$. It follows that $d_{Y'}(\gamma_i, \gamma_h) < C_3 + 2C_2$. Thus $d_Y(\gamma_i, \gamma_h) < C_3 + 2C_2 + 2L_4$. Similarly, $d_Y(\gamma_h, \gamma_j) < C_3 + 2C_2$. Deduce

$$d_Y(\gamma_i, \gamma_j) < 2C_3 + 4C_2 + 2L_4 < L_1(X),$$

the last inequality by (13.16). This is a contradiction to the assumption. \square

Claim. The interval $[p, q]$ is inductive for Y .

Proof. We must check that $d_Y(\gamma_p, \gamma_q) \geq L_1(Y)$. Suppose first that Y is not paired. Then by the definition of p, q , (2) of Axiom 13.4, and the triangle inequality we have

$$d_Y(\mu_{r(i)}, \mu_{r(j)}) \leq d_Y(\mu_{r(p)}, \mu_{r(q)}) + 2C_3.$$

Thus by Axiom 13.5,

$$d_Y(\gamma_i, \gamma_j) \leq d_Y(\gamma_p, \gamma_q) + 2C_3 + 4C_2.$$

Since by (13.17),

$$L_1(Y) + 2C_3 + 4C_2 < L_1(X) \leq d_Y(\gamma_i, \gamma_j)$$

we are done.

When Y is paired the proof is similar but we must use the slightly stronger inequality $L_1(Y) + 2C_3 + 4C_2 + 4L_4 < L_1(X)$. \square

Thus, when $B_X(i, j)$ is non-empty we may find a hole Y and indices p, q as above. In this situation, we subdivide the element $[i, j] \in \mathcal{P}_X$ into the elements $[i, p-1]$, $[p, q]$, and $[q+1, j]$. (The first or third intervals, or both, may be empty.) The interval $[p, q] \in \mathcal{P}_X$ is determined to be inductive and associated to Y . Proceed on to the next undetermined element. This completes the construction of \mathcal{P}_X .

As a bit of notation, if $[i, j] \in \mathcal{P}_X$ is associated to $Y \subset X$ we will sometimes write $I_Y = [i, j]$.

14.2. Properties of the inductive partition.

Lemma 14.1. *Suppose that Y, Z are holes and I_Z is an inductive element of \mathcal{P}_X associated to Z . Suppose that $r(I_Z) \subset J_Y$ (or $r(I_Z) \subset J_Y \cap J_{Y'}$, if Y is paired). Then*

- Z is nested in Y or
- Z and Z' are paired and Z' is nested in Y .

Proof. Let $I_Z = [i, j]$. Suppose first that Y is strictly nested in Z . Then by (ii) of Lemma 13.13, $d_Z(\mu_{r(i)}, \mu_{r(j)}) < C_3$. Then by Axiom 13.5

$$d_Z(\gamma_i, \gamma_j) < C_3 + 2C_2 < L_1(Z),$$

a contradiction. We reach the same contradiction if Y and Z overlap using (iii) of Lemma 13.13.

Now, if Z and Y are disjoint then there are two cases: Suppose first that Y is paired with Y' . Since all holes interfere, Y' and Z must meet. In this case we are done, just as in the previous paragraph.

Suppose now that Z is paired with Z' . Since all holes interfere, Z' and Y must meet. If Z' is nested in Y then we are done. If Y is strictly nested in Z' then, as $r([i, j]) \subset J_Y$, we find that as above by Axioms 13.5 and (ii) of Lemma 13.13 that

$$d_{Z'}(\gamma_i, \gamma_j) < C_3 + 2C_2$$

and so $d_Z(\gamma_i, \gamma_j) < C_3 + 2C_2 + 2L_4 < L_1(Z)$, a contradiction. We reach the same contradiction if Y and Z' overlap. \square

Proposition 14.2. *Suppose $Y \subsetneq X$ is a hole for \mathcal{G} .*

- (1) If Y is associated to an inductive interval $I_Y \in \mathcal{P}_X$ and Y is paired with Y' then Y' is not associated to any inductive interval in \mathcal{P}_X .
- (2) There is at most one inductive interval $I_Y \in \mathcal{P}_X$ associated to Y .
- (3) There are at most two holes Z and W , distinct from Y (and from Y' , if Y is paired) such that
 - there are inductive intervals $I_Z = [h, i]$ and $I_W = [j, k]$ and
 - $d_Y(\gamma_h, \gamma_i), d_Y(\gamma_j, \gamma_k) \geq L_0$.

Remark 14.3. It follows that for any hole Y there are at most three inductive intervals in the partition \mathcal{P}_X where Y has projection distance greater than L_0 .

Proof of Proposition 14.2. To prove the first claim: Suppose that $I_Y = [p, q]$ and $I_{Y'} = [p', q']$ with $q < p'$. It follows that $[r(p), r(q')] \subset J_Y \cap J_{Y'}$. If $q + 1 = p'$ then the partition would have chosen a larger inductive interval for one of Y or Y' . It must be the case that there is an inductive interval $I_Z \subset [q + 1, p' - 1]$ for some hole Z , distinct from Y and Y' , with $\xi(Z) \geq \xi(Y)$. However, by Lemma 14.1 we find that Z is nested in Y or in Y' . It follows that $Z = Y$ or $Z = Y'$, a contradiction.

The second statement is essentially similar.

Finally suppose that Z and W are the first and last holes, if any, satisfying the hypotheses of the third claim. Since $d_Y(\gamma_h, \gamma_i) \geq L_0$ we find by Axiom 13.5 that

$$d_Y(\mu_{r(h)}, \mu_{r(i)}) \geq L_0 - 2C_2.$$

By (13.14), $L_0 - 2C_2 > C_3$ so that

$$J_Y \cap r(I_Z) \neq \emptyset.$$

If Y is paired then, again by (13.14) we have $L_0 > C_3 + 2C_2 + 2L_4$, we also find that $J_{Y'} \cap r(I_Z) \neq \emptyset$. Symmetrically, $J_Y \cap r(I_W)$ (and $J_{Y'} \cap r(I_W)$) are also non-empty.

It follows that the interval between I_Z and I_W , after reindexing, is contained in J_Y (and $J_{Y'}$, if Y is paired). Thus for any inductive interval $I_V = [p, q]$ between I_Z and I_W the associated hole V is nested in Y (or V' is nested in Y'), by Lemma 14.1. If $V = Y$ or $V = Y'$ there is nothing to prove. Suppose instead that V (or V') is strictly nested in Y . It follows that

$$d_Y(\gamma_p, \gamma_q) < C_3 + 2C_2 < L_0.$$

Thus there are no inductive intervals between I_Z and I_W satisfying the hypotheses of the third claim. \square

The following lemma and proposition bound the number of inductive intervals. The discussion here is very similar to (and in fact inspired) the *antichains* defined in [34, Section 5]. Our situation is complicated by the presence of non-holes and interfering holes.

Lemma 14.4. *Suppose that X, α_X, β_X are given, as above. For any $\ell \geq (3 \cdot L_2)^{\xi(X)}$, if $\{Y_i\}_{i=1}^\ell$ is a collection of distinct strict sub-holes of X each having $d_{Y_i}(\alpha_X, \beta_X) \geq L_1(X)$ then there is a hole $Z \subseteq X$ such that $d_Z(\alpha_X, \beta_X) \geq L_2 - 1$ and Z contains at least L_2 of the Y_i . Furthermore, for at least $L_2 - 4(C_1 + C_3 + 2C_3 + 2)$ of these Y_i we find that $J_{Y_i} \subsetneq J_Z$. (If Z is paired then $J_{Y_i} \subsetneq J_Z \cap J_{Z'}$.) Each of these Y_i is disjoint from a distinct vertex $\eta_i \in [\pi_Z(\alpha_X), \pi_Z(\beta_X)]$.*

Proof. Let g_X be a geodesic in $\mathcal{C}(X)$ joining α_X, β_X . By the Bounded Geodesic Image Theorem (Theorem 4.6), since $L_1(X) > M_0$, for every Y_i there is a vertex $\omega_i \in g_X$ such that $Y_i \subset X \setminus \omega_i$. Thus $d_X(\omega_i, \partial Y_i) \leq 1$. If there are at least L_2 distinct ω_i , associated to distinct Y_i , then $d_X(\alpha_X, \beta_X) \geq L_2 - 1$. In this situation we take $Z = X$. Since $J_X = [0, N]$ we are done.

Thus assume there do not exist at least L_2 distinct ω_i . Then there is some fixed ω among these ω_i such that at least $\frac{\ell}{L_2} \geq 3(3 \cdot L_2)^{\xi(X)-1}$ of the Y_i satisfy

$$Y_i \subset (X \setminus \omega).$$

Thus one component, call it W , of $X \setminus \omega$ contains at least $(3 \cdot L_2)^{\xi(X)-1}$ of the Y_i . Let g_W be a geodesic in $\mathcal{C}(W)$ joining $\alpha_W = \pi_W(\alpha_X)$ and $\beta_W = \pi_W(\beta_X)$. Notice that

$$d_{Y_i}(\alpha_W, \beta_W) \geq d_{Y_i}(\alpha_X, \beta_X) - 8$$

because we are projecting to nested subsurfaces. This follows for example from Lemma 4.4. Hence $d_{Y_i}(\alpha_W, \beta_W) \geq L_1(W)$.

Again apply Theorem 4.6. Since $L_1(W) > M_0$, for every remaining Y_i there is a vertex $\eta_i \in g_W$ such that

$$Y_i \subset (W \setminus \eta_i)$$

If there are at least L_2 distinct η_i then we take $Z = W$. Otherwise we repeat the argument. Since the complexity of each successive subsurface is decreasing by at least 1, we must eventually find the desired Z containing at least L_2 of the Y_i , each disjoint from distinct vertices of g_Z .

So suppose that there are at least L_2 distinct η_i associated to distinct Y_i and we have taken $Z = W$. Now we must find at least $L_2 - 4(C_1 + C_3 + 2C_3 + 2)$ of these Y_i where $J_{Y_i} \subsetneq J_Z$.

To this end we focus attention on a small subset $\{Y^j\}_{j=1}^5 \subset \{Y_i\}$. Let η_j be the vertex of $g_Z = g_W$ associated to Y^j . We choose these Y^j so that

- the η_j are arranged along g_Z in order of index and
- $d_Z(\eta_j, \eta_{j+1}) > C_1 + C_3 + 2C_3 + 2$, for $j = 1, 2, 3, 4$.

This is possible by (13.15) because

$$L_2 > 4(C_1 + C_3 + 2C_3).$$

Set $J_j = J_{Y^j}$ and pick any indices $m_j \in J_j$. (If Z is paired then Y^j is as well and we pick $m_j \in J_{Y^j} \cap J_{(Y^j)'}$.) We use $\mu(m_j)$ to denote μ_{m_j} . Since ∂Y^j is disjoint from η_j , Axiom 13.4 and Lemma 2.2 imply

$$(14.5) \quad d_Z(\mu(m_j), \eta_j) \leq C_3 + 1.$$

Since the sequence $\pi_Z(\mu_n)$ satisfies the reverse triangle inequality (Lemma 13.12), it follows that the m_j appear in $[0, N]$ in order agreeing with their index. The triangle inequality implies that

$$d_Z(\mu(m_1), \mu(m_2)) > C_3.$$

Thus Axiom 13.4 implies that $J_Z \cap [m_1, m_2]$ is non-empty. Similarly, $J_Z \cap [m_4, m_5]$ is non-empty. It follows that $[m_2, m_4] \subset J_Z$. (If Z is paired then, after applying the symmetry τ to g_Z , the same argument proves $[m_2, m_4] \subset J_{Z'}$.)

Notice that $J_2 \cap J_3 = \emptyset$. For if $m \in J_2 \cap J_3$ then by (14.5) both $d_Z(\mu_m, \eta_2)$ and $d_Z(\mu_m, \eta_3)$ are bounded by $C_3 + 1$. It follows that

$$d_Z(\eta_2, \eta_3) < 2C_3 + 2,$$

a contradiction. Similarly $J_3 \cap J_4 = \emptyset$. We deduce that $J_3 \subsetneq [m_2, m_4] \subset J_Z$. (If Z is paired $J_3 \subset J_Z \cap J_{Z'}$.) Finally, there are at least

$$L_2 - 4(C_1 + C_3 + 2C_3 + 2)$$

possible Y_i 's which satisfy the hypothesis on Y^3 . This completes the proof. \square

Define

$$\mathcal{P}_{\text{ind}} = \{I \in \mathcal{P}_X \mid I \text{ is inductive}\}.$$

Proposition 14.6. *The number of inductive intervals is a lower bound for the projection distance in X :*

$$d_X(\alpha_X, \beta_X) \geq \frac{|\mathcal{P}_{\text{ind}}|}{2(3 \cdot L_2)^{\xi(X)-1} + 1} - 1.$$

Proof. Suppose, for a contradiction, that the conclusion fails. Let g_X be a geodesic in $\mathcal{C}(X)$ connecting α_X to β_X . Then, as in the proof of Lemma 14.4, there is a vertex ω of g_X and a component $W \subset X \setminus \omega$ where at least $(3 \cdot L_2)^{\xi(X)-1}$ of the inductive intervals in I_X have associated surfaces, Y_i , contained in W .

Since $\xi(X) - 1 \geq \xi(W)$ we may apply Lemma 14.4 inside of W . So we find a surface $Z \subseteq W \subsetneq X$ so that

- Z contains at least L_2 of the Y_i ,
- $d_Z(\alpha_X, \beta_X) \geq L_2$, and
- there are at least $L_2 - 4(C_1 + C_3 + 2C_3 + 2)$ of the Y_i where $J_{Y_i} \subsetneq J_Z$.

Since $Y_i \subsetneq Z$ and Y_i is a hole, Z is also a hole. Since $L_2 > L_1(X)$ it follows that $Z \in B_X$. Let $\mathcal{Y} = \{Y_i\}$ be the set of Y_i satisfying the third bullet. Let $Y^1 \in \mathcal{Y}$ and $\eta_1 \in g_Z$ satisfy $\partial Y^1 \cap \eta_1 = \emptyset$ and η_1 is the first such. Choose Y^2 and η_2 similarly, so that η_2 is the last such. By Lemma 14.4

$$(14.7) \quad d_Z(\eta_1, \eta_2) \geq L_2 - 4(C_1 + C_3 + 2C_3 + 2).$$

Let $p = \min I_{Y^1}$ and $q = \max I_{Y^2}$. Note that $[p, q] \subset J_Z$. (If Z is paired with Z' then $[p, q] \subset J_Z \cap J_{Z'}$.) Again by (1) of Axiom 13.4, and Lemma 2.2,

$$d_Z(\mu_{r(p)}, \partial Y^1) < C_3.$$

It follows that

$$d_Z(\mu_{r(p)}, \eta_1) \leq C_3 + 1$$

and the same bound applies to $d_Z(\mu_{r(q)}, \eta_2)$. Combined with (14.7) we find that

$$d_Z(\mu_{r(p)}, \mu_{r(q)}) \geq L_2 - 4C_1 - 4C_3 - 10C_3 - 10.$$

By the reverse triangle inequality (Lemma 13.12), for any $p' \leq p, q \leq q'$,

$$d_Z(\mu_{r(p')}, \mu_{r(q')}) \geq L_2 - 6C_1 - 4C_3 - 10C_3 - 10.$$

Finally by Axiom 13.5 and the above inequality we have

$$d_Z(\gamma_{p'}, \gamma_{q'}) \geq L_2 - 6C_1 - 4C_3 - 10C_3 - 10 - 2C_2.$$

By (13.15) the right-hand side is greater than $L_1(X) + 2L_4$ so we deduce that $Z \in B_X(p', q')$, for any such p', q' . (When Z is paired deduce also that $Z' \in B_X(p', q')$.)

Let I_V be the first inductive interval chosen by the procedure with the property that $I_V \cap [p, q] \neq \emptyset$. Note that, since I_{Y^1} and I_{Y^2} will also be chosen, $I_V \subset [p, q]$. Let p', q' be the indices so that V is chosen from $B_X(p', q')$. Thus $p' \leq p$ and $q \leq q'$. However, since $I_V \subset [p, q] \subset J_Z$, Lemma 14.1 implies that V is strictly nested in Z . (When pairing

occurs we may find instead that $V \subset Z'$ or $V' \subset Z$.) Thus $\xi(Z) > \xi(V)$ and we find that Z would be chosen from $B_X(p', q')$, instead of V . This is a contradiction. \square

14.3. Electric partition. The goal of this subsection is to prove:

Proposition 14.8. *There is a constant A depending only on $\xi(X)$, so that: if $[i, j] \subset [0, K]$ is a electric interval then*

$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_A d_X(\gamma_i, \gamma_j).$$

We begin by building a partition of $[i, j]$ into straight and shortcut intervals. Define

$$C_X = \{Y \subsetneq X \mid Y \text{ is a non-hole and } d_Y(\mu_{r(i)}, \mu_{r(j)}) \geq L_1(X)\}.$$

We also define, for all $[p, q] \subset [i, j]$

$$C_X(p, q) = \{Y \in C_X \mid J_Y \cap [r(p), r(q)] \neq \emptyset\}.$$

Our recursion starts with the partition of one part, $\mathcal{P}(i, j) = \{[i, j]\}$. Recursively $\mathcal{P}(i, j)$ is a partition of $[i, j]$ into shortcut, straight, or undetermined intervals. Suppose that $[p, q] \in \mathcal{P}(i, j)$ is undetermined.

Claim. If $C_X(p, q)$ is empty then $[p, q]$ is straight.

Proof. We show the contrapositive. Suppose that Y is a non-hole with $d_Y(\mu_{r(p)}, \mu_{r(q)}) \geq L_2$. Since $L_2 > C_3$, Axiom 13.4 implies that $J_Y \cap [r(p), r(q)]$ is non-empty. Also, the reverse triangle inequality (Lemma 13.12) gives:

$$d_Y(\mu_{r(p)}, \mu_{r(q)}) < d_Y(\mu_{r(i)}, \mu_{r(j)}) + 2C_1.$$

Since $L_2 > L_1(X) + 2C_1$, we find that $Y \in C_X$. It follows that $Y \in C_X(p, q)$. \square

So when $C_X(p, q)$ is empty the interval $[p, q]$ is determined to be straight. Proceed onto the next undetermined element of $\mathcal{P}(i, j)$. Now suppose that $C_X(p, q)$ is non-empty. Then we choose any $Y \in C_X(p, q)$ so that Y has maximal $\xi(Y)$ amongst the elements of $C_X(p, q)$. Notice that by the accessibility requirement that $J_Y \cap [r(p), r(q)]$ is non-empty.

There are two cases. If $J_Y \cap r([p, q])$ is empty then let $p' \in [p, q]$ be the largest integer so that $r(p') < \min J_Y$. Note that p' is well-defined. Now divide the interval $[p, q]$ into the two undetermined intervals $[p, p']$, $[p' + 1, q]$. In this situation we say Y is associated to a *shortcut of length one* and we add the element $[p' + \frac{1}{2}]$ to $\mathcal{P}(i, j)$.

Next suppose that $J_Y \cap r([p, q])$ is non-empty. Let $p', q' \in [p, q]$ be the first and last indices, respectively, so that $r(p'), r(q') \in J_Y$. (Note that it is possible to have $p' = q'$.) Partition $[p, q] = [p, p' - 1] \cup [p', q'] \cup [q' + 1, q]$.

The first and third parts are undetermined; either may be empty. This completes the recursive construction of the partition.

Define

$$\mathcal{P}_{\text{short}} = \{I \in \mathcal{P}(i, j) \mid I \text{ is a shortcut}\}$$

and

$$\mathcal{P}_{\text{str}} = \{I \in \mathcal{P}(i, j) \mid I \text{ is straight}\}.$$

Proposition 14.9. *With $\mathcal{P}(i, j)$ as defined above,*

$$d_X(\gamma_i, \gamma_j) \geq \frac{|\mathcal{P}_{\text{short}}|}{2(3 \cdot L_2)^{\xi(X)-1} + 1} - 1.$$

Proof. The proof is identical to that of Proposition 14.6 with the caveat that in Lemma 14.4 we must use the markings $\mu_{r(i)}$ and $\mu_{r(j)}$ instead of the endpoints γ_i and γ_j . \square

Now we “electrify” every shortcut interval using Theorem 13.1 recursively.

Lemma 14.10. *There is a constant $L_3 = L_3(X, \mathcal{G})$, so that for every shortcut interval $[p, q]$ we have $d_{\mathcal{G}}(\gamma_p, \gamma_q) < L_3$.*

Proof. As $[p, q]$ is a shortcut we are given a non-hole $Z \subset X$ so that $r([p, q]) \subset J_Z$. Let $Y = X \setminus Z$. Thus Axiom 13.6 gives vertices γ'_p, γ'_q of \mathcal{G} lying in Y or in Z , so that $d_{\mathcal{G}}(\gamma_p, \gamma'_p), d_{\mathcal{G}}(\gamma_q, \gamma'_q) \leq C_4$.

If one of γ'_p, γ'_q lies in Y while the other lies in Z then

$$d_{\mathcal{G}}(\gamma_p, \gamma_q) < 2C_4 + 1.$$

If both lie in Z then, as Z is a non-hole, there is a vertex $\delta \in \mathcal{G}(S)$ disjoint from both of γ'_p and γ'_q and we have

$$d_{\mathcal{G}}(\gamma_p, \gamma_q) < 2C_4 + 2.$$

If both lie in Y then there are two cases. If Y is not a hole for $\mathcal{G}(S)$ then we are done as in the previous case. If Y is a hole then by the definition of shortcut interval, Lemma 5.7, and the triangle inequality we have

$$d_W(\gamma'_p, \gamma'_q) < 6 + 6C_4 + L_2$$

for all holes $W \subset Y$. Notice that Y is strictly contained in X . Thus we may inductively apply Theorem 13.1 with $c = 6 + 6C_4 + L_2$. We deduce that all terms on the right-hand side of the distance estimate vanish and thus $d_{\mathcal{G}}(\gamma'_p, \gamma'_q)$ is bounded by a constant depending only on X and \mathcal{G} . The same then holds for $d_{\mathcal{G}}(\gamma_p, \gamma_q)$ and we are done. \square

We are now equipped to give:

Proof of Proposition 14.8. Suppose that $\mathcal{P}(i, j)$ is the given partition of the electric interval $[i, j]$ into straight and shortcut subintervals. As a bit of notation, if $[p, q] = I \in \mathcal{P}(i, j)$, we take $d_{\mathcal{G}}(I) = d_{\mathcal{G}}(\gamma_p, \gamma_q)$ and $d_X(I) = d_X(\gamma_p, \gamma_q)$. Applying Axiom 13.5 we have

$$(14.11) \quad d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq \sum_{I \in \mathcal{P}_{\text{str}}} d_{\mathcal{G}}(I) + \sum_{I \in \mathcal{P}_{\text{short}}} d_{\mathcal{G}}(I) + C_2 |\mathcal{P}(i, j)|$$

The last term arises from connecting left endpoints of intervals with right endpoints. We must bound the three terms on the right.

We begin with the third; recall that $|\mathcal{P}(i, j)| = |\mathcal{P}_{\text{short}}| + |\mathcal{P}_{\text{str}}|$, that $|\mathcal{P}_{\text{str}}| \leq |\mathcal{P}_{\text{short}}| + 1$, and that $|\mathcal{P}_{\text{short}}| \leq_A d_X(\gamma_i, \gamma_j)$. The second inequality follows from the construction of the partition while the last is implied by Proposition 14.9. Thus the third term of Equation 14.11 is quasi-bounded above by $d_X(\gamma_i, \gamma_j)$.

By Lemma 14.10, the second term of Equation 14.11 is at most $L_3 |\mathcal{P}_{\text{short}}|$. Finally, by Axiom 13.11, for all $I \in \mathcal{P}_{\text{str}}$ we have

$$d_{\mathcal{G}}(I) \leq_A d_X(I),$$

Also, it follows from the reverse triangle inequality (Lemma 13.12) that

$$\sum_{I \in \mathcal{P}_{\text{str}}} d_X(I) \leq d_X(\gamma_i, \gamma_j) + (2C_1 + 2C_2) |\mathcal{P}_{\text{str}}| + 2C_2.$$

We deduce that $\sum_{I \in \mathcal{P}_{\text{str}}} d_{\mathcal{G}}(I)$ is also quasi-bounded above by $d_X(\gamma_i, \gamma_j)$. Thus for a somewhat larger value of A we find

$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_A d_X(\gamma_i, \gamma_j).$$

This completes the proof. \square

14.4. The upper bound. We will need:

Proposition 14.12. *For any $c > 0$ there is a constant A with the following property. Suppose that $[i, j] = I_Y$ is an inductive interval in \mathcal{P}_X . Then we have:*

$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_A \sum_Z [d_Z(\gamma_i, \gamma_j)]_c$$

where Z ranges over all holes for \mathcal{G} contained in X .

Proof. Axiom 13.6 gives vertices $\gamma'_i, \gamma'_j \in \mathcal{G}$, contained in Y , so that $d_{\mathcal{G}}(\gamma_i, \gamma'_i) \leq C_4$ and the same holds for j . Since projection to holes is coarsely Lipschitz (Lemma 5.7) for any hole Z we have $d_Z(\gamma_i, \gamma'_i) \leq 3 + 3C_3$.

Fix any $c > 0$. Now, since

$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq d_{\mathcal{G}}(\gamma'_i, \gamma'_j) + 2C_3$$

to find the required constant A it suffices to bound $d_{\mathcal{G}}(\gamma'_i, \gamma'_j)$. Let $c' = c + 6C_3 + 6$. Since $Y \subsetneq X$, induction gives us a constant A so that

$$\begin{aligned} d_{\mathcal{G}}(\gamma'_i, \gamma'_j) &\leq_A \sum_Z [d_Z(\gamma'_i, \gamma'_j)]_{c'} \\ &\leq \sum_Z [d_Z(\gamma_i, \gamma_j) + 6C_3 + 6]_{c'} \\ &< (6C_3 + 6)N + \sum_Z [d_Z(\gamma_i, \gamma_j)]_c \end{aligned}$$

where N is the number of non-zero terms in the final sum. Also, the sum ranges over sub-holes of Y . We may take A somewhat larger to deal with the term $(6C_3 + 6)N$ and include all holes $Z \subset X$ to find

$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_A \sum_Z [d_Z(\gamma_i, \gamma_j)]_c$$

where the sum is over all holes $Z \subset X$. □

14.5. Finishing the proof. Now we may finish the proof of Theorem 13.1. Fix any constant $c \geq 0$. Suppose that X, α_X, β_X are given as above. Suppose that $\Gamma = \{\gamma_i\}_{i=0}^K$ is the given combinatorial path and \mathcal{P}_X is the partition of $[0, K]$ into inductive and electric intervals. So we have:

$$(14.13) \quad d_{\mathcal{G}}(\alpha_X, \beta_X) \leq \sum_{I \in \mathcal{P}_{\text{ind}}} d_{\mathcal{G}}(I) + \sum_{I \in \mathcal{P}_{\text{ele}}} d_{\mathcal{G}}(I) + C_2 |\mathcal{P}_X|$$

Again, the last term arises from adjacent right and left endpoints of different intervals.

We must bound the terms on the right-hand side; begin by noticing that $|\mathcal{P}_X| = |\mathcal{P}_{\text{ind}}| + |\mathcal{P}_{\text{ele}}|$, $|\mathcal{P}_{\text{ele}}| \leq |\mathcal{P}_{\text{ind}}| + 1$ and $|\mathcal{P}_{\text{ind}}| \leq_A d_X(\alpha_X, \beta_X)$. The second inequality follows from the way the partition is constructed and the last follows from Proposition 14.6. Thus the third term of Equation 14.13 is quasi-bounded above by $d_X(\alpha_X, \beta_X)$.

Next consider the second term of Equation 14.13:

$$\begin{aligned} \sum_{I \in \mathcal{P}_{\text{ele}}} d_{\mathcal{G}}(I) &\leq_A \sum_{I \in \mathcal{P}_{\text{ele}}} d_X(I) \\ &\leq d_X(\alpha_X, \beta_X) + (2C_1 + 2C_2) |\mathcal{P}_{\text{ele}}| + 2C_2 \end{aligned}$$

with the first inequality following from Proposition 14.8 and the second from the reverse triangle inequality (Lemma 13.12).

Finally we bound the first term of Equation 14.13. Let $c' = c + L_0$. Thus,

$$\begin{aligned} \sum_{I \in \mathcal{P}_{\text{ind}}} d_{\mathcal{G}}(I) &\leq \sum_{I_Y \in \mathcal{P}_{\text{ind}}} \left(A'_Y \left(\sum_{Z \subsetneq Y} [d_Z(I_Y)]_{c'} \right) + A'_Y \right) \\ &\leq A'' \left(\sum_{I \in \mathcal{P}_{\text{ind}}} \sum_{Z \subsetneq X} [d_Z(I)]_{c'} \right) + A'' \cdot |\mathcal{P}_{\text{ind}}| \\ &\leq A'' \left(\sum_{Z \subsetneq X} \sum_{I \in \mathcal{P}_{\text{ind}}} [d_Z(I)]_{c'} \right) + A'' \cdot |\mathcal{P}_{\text{ind}}| \end{aligned}$$

Here A'_Y and the first inequality are given by Proposition 14.12. Also $A'' = \max\{A'_Y \mid Y \subsetneq X\}$. In the last line, each sum of the form $\sum_{I \in \mathcal{P}_{\text{ind}}} [d_Z(I)]_{c'}$ has at most three terms, by Remark 14.3 and the fact that $c' > L_0$. For the moment, fix a hole Z and any three elements $I, I', I'' \in \mathcal{P}_{\text{ind}}$.

By the reverse triangle inequality (Lemma 13.12) we find that

$$d_Z(I) + d_Z(I') + d_Z(I'') < d_Z(\alpha_X, \beta_X) + 6C_1 + 8C_2$$

which in turn is less than $d_Z(\alpha_X, \beta_X) + L_0$.

It follows that

$$[d_Z(I)]_{c'} + [d_Z(I')]_{c'} + [d_Z(I'')]_{c'} < [d_Z(\alpha_X, \beta_X)]_c + L_0.$$

Thus,

$$\sum_{Z \subsetneq X} \sum_{I \in \mathcal{P}_{\text{ind}}} [d_Z(I)]_{c'} \leq L_0 \cdot N + \sum_{Z \subsetneq X} [d_Z(\alpha_X, \beta_X)]_c$$

where N is the number of non-zero terms in the final sum. Also, the sum ranges over all holes $Z \subsetneq X$.

Combining the above inequalities, and increasing A once again, implies that

$$d_{\mathcal{G}}(\alpha_X, \beta_X) \leq A \sum_Z [d_Z(\alpha_X, \beta_X)]_c$$

where the sum ranges over all holes $Z \subseteq X$. This completes the proof of Theorem 13.1. \square

15. BACKGROUND ON TEICHMÜLLER SPACE

Our goal in Sections 16, 17 and 19 will be to verify the axioms stated in Section 13 for the complex of curves of a non-orientable surface, for the arc complex, and for the disk complex. Here we give the necessary background on Teichmüller space.

Fix now a surface $S = S_{g,n}$ of genus g with n punctures. Two conformal structures on S are equivalent, written $\Sigma \sim \Sigma'$, if there is a conformal map $f: \Sigma \rightarrow \Sigma'$ which is isotopic to the identity. Let $\mathcal{T} = \mathcal{T}(S)$ be the *Teichmüller space* of S ; the set of equivalence classes of conformal structures Σ on S .

Define the Teichmüller metric by,

$$d_{\mathcal{T}}(\Sigma, \Sigma') = \inf_f \left\{ \frac{1}{2} \log K(f) \right\}$$

where the infimum ranges over all quasiconformal maps $f: \Sigma \rightarrow \Sigma'$ isotopic to the identity and where $K(f)$ is the maximal dilatation of f . Recall that the infimum is realized by a Teichmüller map that, in turn, may be defined in terms of a quadratic differential.

15.1. Quadratic differentials.

Definition 15.1. A *quadratic differential* $q(z) dz^2$ on Σ is an assignment of a holomorphic function to each coordinate chart that is a disk and of a meromorphic function to each chart that is a punctured disk. If z and ζ are overlapping charts then we require

$$q_z(z) = q_{\zeta}(\zeta) \left(\frac{d\zeta}{dz} \right)^2$$

in the intersection of the charts. The meromorphic function $q_z(z)$ has at most a simple pole at the puncture $z = 0$.

At any point away from the zeroes and poles of q there is a natural coordinate $z = x + iy$ with the property that $q_z \equiv 1$. In this natural coordinate the foliation by lines $y = c$ is called the *horizontal foliation*. The foliation by lines $x = c$ is called the *vertical foliation*.

Now fix a quadratic differential q on $\Sigma = \Sigma_0$. Let x, y be natural coordinates for q . For every $t \in \mathbb{R}$ we obtain a new quadratic differential q_t with coordinates

$$x_t = e^t x, \quad y_t = e^{-t} y.$$

Also, q_t determines a conformal structure Σ_t on S . The map $t \mapsto \Sigma_t$ is the Teichmüller geodesic determined by Σ and q .

15.2. Marking coming from a Teichmüller geodesic. Suppose that Σ is a Riemann surface structure on S and σ is the uniformizing hyperbolic metric in the conformal class of Σ . In a slight abuse of terminology, we call the collection of shortest simple non-peripheral closed geodesics the *systoles* of σ . Fix a constant ϵ smaller than the Margulis constant. The ϵ -thick part of Teichmüller space consists of

those Riemann surfaces such that the hyperbolic systole has length at least ϵ .

We define $P = P(\sigma)$, a *Bers pants decomposition* of S , as follows: pick α_1 , any systole for σ . Define α_i to be any systole of σ restricted to $S \setminus (\alpha_1 \cup \dots \cup \alpha_{i-1})$. Continue in this fashion until P is a pants decomposition. Note that any curve with length less than the Margulis constant will necessarily be an element of P .

Suppose that $\Sigma, \Sigma' \in \mathcal{T}(S)$. Suppose that P, P' are Bers pants decompositions with respect to Σ and Σ' . Suppose also that $d_{\mathcal{T}}(\Sigma, \Sigma') \leq 1$. Then the curves in P have uniformly bounded lengths in Σ' and conversely. By the Collar Lemma, the intersection $\iota(P, P')$ is bounded, solely in terms of $\xi(S)$.

Suppose now that $\{\Sigma_t \mid t \in [-M, M]\}$ is the Teichmüller geodesic defined by the quadratic differentials q_t . Let σ_t be the hyperbolic metric uniformizing Σ_t . Let $P_t = P(\sigma_t)$ be a Bers pants decomposition.

We now find transversals in order to complete P_t to a *Bers marking* ν_t . Suppose that $P_t = \{\alpha_i\}$. For each i , let A^i be the annular cover of S corresponding to α_i . Note that q_t lifts to a singular Euclidean metric q_t^i on A^i . Let α^i be a geodesic representative of the core curve of A^i with respect to the metric q_t^i . Choose $\gamma_i \in \mathcal{C}(A^i)$ to be any geodesic arc, also with respect to q_t^i , that is perpendicular to α^i . Let β_i be any curve in $S \setminus (\{\alpha_j\}_{j \neq i})$ which meets α_i minimally and so that $d_{A^i}(\beta_i, \gamma_i) \leq 3$. (See the discussion after the proof of Lemma 2.4 in [25].) Doing this for each i gives a complete clean marking $\nu_t = \{\alpha_i\} \cup \{\beta_i\}$.

We now have:

Lemma 15.2. [33, Remark 6.2 and Equation (3)] *There is a constant $B_0 = B_0(S)$ with the following property. For any Teichmüller geodesic and for any time t , there is a constant $\delta > 0$ so that if $|t - s| \leq \delta$ then*

$$\iota(\nu_t, \nu_s) < B_0.$$

Suppose that Σ_t and Σ_s are surfaces in the ϵ -thick part of $\mathcal{T}(S)$. We take B_0 sufficiently large so that if $\iota(\nu_t, \nu_s) \geq B_0$ then $d_{\mathcal{T}}(\Sigma_t, \Sigma_s) \geq 1$.

15.3. The marking axiom. We construct a sequence of markings μ_n , for $n \in [0, N] \subset \mathbb{N}$, as follows. Take $\mu_0 = \nu_{-M}$. Now suppose that $\mu_n = \nu_t$ is defined. Let $s > t$ be the first time that there is a marking with $\iota(\nu_t, \nu_s) \geq B_0$, if such a time exists. If so, let $\mu_{n+1} = \nu_s$. If no such time exists take $N = n$ and we are done.

We now show that $\mu_n = \nu_t$ and $\mu_{n+1} = \nu_s$ have bounded intersection. By the above lemma there is a marking ν_r with $t \leq r < s$ and

$$\iota(\nu_r, \nu_s) \leq B_0.$$

By construction

$$\iota(\nu_t, \nu_r) < B_0.$$

Since intersection number bounds distance in the marking complex we find that by the triangle inequality, ν_t and ν_s are bounded distance in the marking complex. Conversely, since distance bounds intersection in the marking complex we find that $\iota(\mu_n, \mu_{n+1})$ is bounded. It follows that $d_Y(\mu_n, \mu_{n+1})$ is uniformly bounded, independent of $Y \subset S$ and of $n \in [0, N]$.

It now follows from Theorem 6.1 of [33] that, for any subsurface $Y \subset S$, the sequence $\{\pi_Y(\mu_n)\} \subset \mathcal{C}(Y)$ is an unparameterized quasi-geodesic. Thus the marking path $\{\mu_n\}$ satisfies the second requirement of Axiom 13.3. The first requirement is trivial as every μ_n fills S .

15.4. The accessibility axiom. We now turn to Axiom 13.4. Since μ_n fills S for every n , the first requirement is a triviality.

In Section 5 of [33] Rafi defines, for every subsurface $Y \subset S$, an *interval of isolation* I_Y inside of the parametrizing interval of the Teichmüller geodesic. Note that I_Y is defined purely in terms of the geometry of the given quadratic differentials. Further, for all $t \in I_Y$ and for all components $\alpha \subset \partial Y$ the hyperbolic length of α in Σ_t is less than the Margulis constant. Furthermore, by Theorem 5.3 [33], there is a constant B_3 so that if $[s, t] \cap I_Y = \emptyset$ then

$$d_Y(\nu_s, \nu_t) \leq B_3.$$

So define $J_Y \subset [0, N]$ to be the subinterval of the marking path where the time corresponding to μ_n lies in I_Y . The third requirement follows. Finally, if $m \in J_Y$ then ∂Y is contained in $\text{base}(\mu_m)$ and thus $\iota(\partial Y, \mu_m) \leq 2 \cdot |\partial Y|$.

15.5. The distance estimate in Teichmüller space. We end this section by quoting another result of Rafi:

Theorem 15.3. [33, Theorem 2.4] *Fix a surface S and a constant $\epsilon > 0$. There is a constant $C_0 = C_0(S, \epsilon)$ so that for any $c > C_0$ there is a constant A with the following property. Suppose that Σ and Σ' lie in the ϵ -thick part of $\mathcal{T}(S)$. Then*

$$d_{\mathcal{T}}(\Sigma, \Sigma') =_A \sum_X [d_X(\mu, \mu')]_c + \sum_{\alpha} [\log d_{\alpha}(\mu, \mu')]_c$$

where μ and μ' are Bers markings on Σ and Σ' , $Y \subset S$ ranges over non-annular surfaces and α ranges over vertices of $\mathcal{C}(S)$. \square

16. PATHS FOR THE NON-ORIENTABLE SURFACE

Fix F a compact, connected, and non-orientable surface. Let S be the orientation double cover with covering map $\rho_F: S \rightarrow F$. Let $\tau: S \rightarrow S$ be the associated involution. Note that $\mathcal{C}(F) = \mathcal{C}^\tau(S)$. Let $\mathcal{C}^\tau(S) \rightarrow \mathcal{C}(S)$ be the relation sending a symmetric multicurve to its components.

Our goal for this section is to prove Lemma 16.4, the classification of holes for $\mathcal{C}(F)$. As remarked above, Lemma 6.3 and Corollary 6.4 follow, proving the hyperbolicity of $\mathcal{C}(F)$.

16.1. The marking path. We will use the extreme rigidity of Teichmüller geodesics to find τ -invariant marking paths. We first show that τ -invariant Bers pants decompositions exist.

Lemma 16.1. *Fix a τ -invariant hyperbolic metric σ . Then there is a Bers pants decomposition $P = P(\sigma)$ which is τ -invariant.*

Proof. Let $P_0 = \emptyset$. Suppose that $0 \leq k < \xi(S)$ curves have been chosen to form P_k . By induction we may assume that P_k is τ -invariant. Let Y be a component of $S \setminus P_k$ with $\xi(Y) \geq 1$. Note that since τ is orientation reversing, τ does not fix any boundary component of Y .

Pick any systole α for Y .

Claim. Either $\tau(\alpha) = \alpha$ or $\alpha \cap \tau(\alpha) = \emptyset$.

Proof. Suppose not and take $p \in \alpha \cap \tau(\alpha)$. Then $\tau(p) \in \alpha \cap \tau(\alpha)$ as well, and, since τ has no fixed points, $p \neq \tau(p)$. The points p and $\tau(p)$ divide α into segments β and γ . Since τ is an isometry, we have

$$\ell_\sigma(\tau(\alpha)) = \ell_\sigma(\alpha) \quad \text{and} \quad \ell_\sigma(\tau(\beta)) = \ell_\sigma(\beta).$$

Now concatenate to obtain (possibly immersed) loops

$$\beta' = \beta * \tau(\beta) \quad \text{and} \quad \gamma' = \gamma * \tau(\gamma).$$

If β' is null-homotopic then $\alpha \cup \tau(\alpha)$ cuts a monogon or a bigon out of S , contradicting our assumption that α was a geodesic. Suppose, by way of contradiction, that β' is homotopic to some boundary curve $b \subset \partial Y$. Since $\tau(\beta') = \beta'$, it follows that $\tau(b)$ and β' are also homotopic. Thus b and $\tau(b)$ cobound an annulus, implying that Y is an annulus, a contradiction. The same holds for γ' .

Let β'' and γ'' be the geodesic representatives of β' and γ' . Since β and $\tau(\beta)$ meet transversely, β'' has length in σ strictly smaller than $2\ell_\sigma(\beta)$. Similarly the length of γ'' is strictly smaller than $2\ell_\sigma(\gamma)$. Suppose that β'' is shorter than γ'' . It follows that β'' is strictly shorter than α . If β'' is embedded then this contradicts the assumption that α was shortest.

If β'' is not embedded then there is an embedded curve β''' inside of a regular neighborhood of β'' which is again essential, non-peripheral, and has geodesic representative shorter than β'' . This is our final contradiction and the claim is proved. \square

Thus, if $\tau(\alpha) = \alpha$ we let $P_{k+1} = P_k \cup \{\alpha\}$ and we are done. If $\tau(\alpha) \neq \alpha$ then by the above claim $\tau(\alpha) \cap \alpha = \emptyset$. In this case let $P_{k+2} = P_k \cup \{\alpha, \tau(\alpha)\}$ and Lemma 16.1 is proved. \square

Transversals are chosen with respect to a quadratic differential metric. Suppose that $\alpha, \beta \in \mathcal{C}^\tau(S)$. If α and β do not fill S then we may replace S by the support of their union. Following Thurston [38] there exists a square-tiled quadratic differential q with squares associated to the points of $\alpha \cap \beta$. (See [6] for analysis of how the square-tiled surface relates to paths in the curve complex.) Let q_t be image of q under the Teichmüller geodesic flow. We have:

Lemma 16.2. $\tau^* q_t = q_t$.

Proof. Note that τ preserves α and also β . Since τ permutes the points of $\alpha \cap \beta$ it permutes the rectangles of the singular Euclidean metric q_t while preserving their vertical and horizontal foliations. Thus τ is an isometry of the metric and the conclusion follows. \square

We now choose the Teichmüller geodesic $\{\Sigma_t \mid t \in [-M, M]\}$ so that the hyperbolic length of α is less than the Margulis constant in σ_{-M} and the same holds for β in σ_M . Also, α is the shortest curve in σ_{-M} and similarly for β in σ_M .

Lemma 16.3. *Fix t . There are transversals for P_t which are close to being quadratic perpendicular in q_t and which are τ -invariant.*

Proof. Let $P = P_t$ and fix $\alpha \in P$. Let $X = S \setminus (P \setminus \alpha)$. There are two cases: either $\tau(X) \cap X = \emptyset$ or $\tau(X) = X$. Suppose the former. So we choose any transversal $\beta \subset X$ close to being q_t -perpendicular and take $\tau(\beta)$ to be the transversal to $\tau(\alpha)$.

Suppose now that $\tau(X) = X$. It follows that X is a four-holed sphere. The quotient X/τ is homeomorphic to a twice-holed \mathbb{RP}^2 . Therefore there are only four essential non-peripheral curves in X/τ . Two of these are cores of Mobius bands and the other two are their doubles. The cores meet in a single point. Perforce α is the double cover of one core and we take β the double cover of the other.

It remains only to show that β is close to being q_t -perpendicular. Let S^α be the annular cover of S and lift q_t to S^α . Let \perp be the set of q_t^α -perpendiculars. This is a τ -invariant diameter one subset of $\mathcal{C}(S^\alpha)$.

If $d_\alpha(\perp, \beta)$ is large then it follows that $d_\alpha(\perp, \tau(\beta))$ is also large. Also, $\tau(\beta)$ twists in the opposite direction from β . Thus

$$d_\alpha(\beta, \tau(\beta)) - 2d_\alpha(\perp, \beta) = O(1)$$

and so $d_\alpha(\beta, \tau(\beta))$ is large, contradicting the fact that β is τ -invariant. \square

Thus τ -invariant markings exist; these have bounded intersection with the markings constructed in Section 15. It follows that the resulting marking path satisfies the marking path and accessibility requirements, Axioms 13.3 and 13.4.

16.2. The combinatorial path. As in Section 15 break the interval $[-M, M]$ into short subintervals and produce a sequence of τ -invariant markings $\{\mu_n\}_{n=0}^N$. To choose the combinatorial path, pick $\gamma_n \in \text{base}(\mu_n)$ so that γ_n is a τ -invariant curve or pair of curves and so that γ_n is shortest in $\text{base}(\mu_n)$.

We now check the combinatorial path requirements given in Axiom 13.5. Note that $\gamma_0 = \alpha$, $\gamma_N = \beta$; also the reindexing map is the identity. Since

$$\iota(\gamma_n, \mu_{r(n)}) = \iota(\gamma_n, \mu_n) = 2$$

the first requirement is satisfied. Since μ_n and μ_{n+1} have bounded intersection, the same holds for γ_n and γ_{n+1} . Projection to F , surgery, and Lemma 2.2 imply that $d_{\mathcal{C}^\tau}(\gamma_n, \gamma_{n+1})$ is uniformly bounded. This verifies Axiom 13.5.

16.3. The classification of holes. We now finish the classification of large holes for $\mathcal{C}^\tau(S)$. Fix $L_0 > 3C_3 + 2C_2 + 2C_1$. Note that these constants are available because we have verified the axioms that give them.

Lemma 16.4. *Suppose that $\alpha, \beta \in \mathcal{C}^\tau(S)$. Suppose that $X \subset S$ has $d_X(\alpha, \beta) > L_0$. Then X is symmetric.*

Proof. Let (Σ_t, q_t) be the Teichmüller geodesic defined above and let σ_t be the uniformizing hyperbolic metric. Since $L_0 > C_3 + 2C_2$ it follows from the accessibility requirement that $J_X = [m, n]$ is non-empty. Now for all t in the interval of isolation I_X

$$\ell_{\sigma_t}(\delta) < \epsilon,$$

where δ is any component of ∂X and ϵ is the Margulis constant. Let $Y = \tau(X)$. Since τ is an isometry (Lemma 16.2) and since the interval of isolation is metrically defined we have $I_Y = I_X$ and thus $J_Y = J_X$.

Deduce that ∂Y is also short in σ_t . This implies that $\partial X \cap \partial Y = \emptyset$. If X and Y overlap then by (iii) of Lemma 13.13 we have

$$d_X(\mu_m, \mu_n) < C_3$$

and so by the triangle inequality, two applications of (2) of Axiom 13.4, we have

$$d_X(\mu_0, \mu_N) < 3C_3.$$

By the combinatorial axiom it follows that

$$d_X(\alpha, \beta) < 3C_3 + 2C_2$$

a contradiction. Deduce that either $X = Y$ or $X \cap Y = \emptyset$ as desired. \square

As noted in Section 6 this shows that the only hole for $\mathcal{C}^\tau(S)$ is S itself. Thus all holes trivially interfere, verifying Axiom 13.2.

16.4. The replacement axiom. We now verify Axiom 13.6 for subsurfaces $Y \subset S$ with $d_Y(\alpha, \beta) \geq L_0$. (We may ignore all subsurfaces with smaller projection by taking $L_1(Y) > L_0$.)

By Lemma 16.4 the subsurface Y is symmetric. If Y is a hole then $Y = S$ and the first requirement is vacuous. Suppose that Y is not a hole. Suppose that γ_n is such that $n \in J_Y$. Thus $\gamma_n \in \text{base}(\mu_n)$. All components of ∂Y are also pants curves in μ_n . It follows that we may take any symmetric curve in ∂Y to be γ' and we are done.

16.5. On straight intervals. Lastly we verify Axiom 13.11. Suppose that $[p, q]$ is a straight interval. We must show that $d_{\mathcal{C}^\tau}(\gamma_p, \gamma_q) \leq d_S(\gamma_p, \gamma_q)$. Suppose that $\mu_p = \nu_s$ and $\mu_q = \nu_t$; that is, s and t are the times when μ_p, μ_q are short markings. Thus $d_X(\mu_p, \mu_q) \leq L_2$ for every $X \subsetneq S$. This implies that the Teichmüller geodesic, along the straight interval, lies in the thick part of Teichmüller space.

Notice that $d_{\mathcal{C}^\tau}(\gamma_p, \gamma_q) \leq C_2|p - q|$, since for all $i \in [p, q - 1]$, $d_{\mathcal{C}^\tau}(\gamma_i, \gamma_{i+1}) \leq C_2$. So it suffices to bound $|p - q|$. By our choice of B_0 and because the Teichmüller geodesic lies in the thick part we find that $|p - q| \leq d_{\mathcal{T}}(\Sigma_s, \Sigma_t)$. Rafi's distance estimate (Theorem 15.3) gives:

$$d_{\mathcal{T}}(\Sigma_s, \Sigma_t) =_A d_S(\nu_s, \nu_t).$$

Since $\nu_s = \mu_p$, $\nu_t = \mu_q$, and since $\gamma_p \in \text{base}(\mu_p)$, $\gamma_q \in \text{base}(\mu_q)$ deduce that

$$d_S(\mu_p, \mu_q) \leq d_S(\gamma_p, \gamma_q) + 4.$$

This verifies Axiom 13.11. Thus the distance estimate holds for $\mathcal{C}^\tau(S) = \mathcal{C}(F)$. Since there is only one hole for $\mathcal{C}(F)$ we deduce that the map $\mathcal{C}(F) \rightarrow \mathcal{C}(S)$ is a quasi-isometric embedding. As a corollary we have:

Theorem 16.5. *The curve complex $\mathcal{C}(F)$ is Gromov hyperbolic.* \square

17. PATHS FOR THE ARC COMPLEX

Here we verify that our axioms hold for the arc complex $\mathcal{A}(S, \Delta)$. It is worth pointing out that the axioms may be verified using Teichmüller geodesics, train track splitting sequences, or resolutions of hierarchies. Here we use the former because it also generalizes to the non-orientable case; this is discussed at the end of this section.

First note that Axiom 13.2 follows from Lemma 7.3.

17.1. The marking path. We are given a pair of arcs $\alpha, \beta \in \mathcal{A}(X, \Delta)$. Recall that $\sigma_S: \mathcal{A}(X) \rightarrow \mathcal{C}(X)$ is the surgery map, defined in Definition 4.2. Let $\alpha' = \sigma_S(\alpha)$ and define β' similarly. Note that α' cuts a pants off of S . As usual, we may assume that α' and β' fill X . If not we pass to the subsurface they do fill.

As in the previous sections let q be the quadratic differential determined by α' and β' . Exactly as above, fix a marking path $\{\mu_n\}_{n=0}^N$. This path satisfies the marking and accessibility axioms (13.3, 13.4).

17.2. The combinatorial path. Let $Y_n \subset X$ be any component of $X \setminus \text{base}(\mu_n)$ meeting Δ . So Y_n is a pair of pants. Let γ_n be any essential arc in Y_n with both endpoints in Δ . Since $\alpha' \subset \text{base}(\mu_0)$ and $\beta' \subset \text{base}(\mu_N)$ we may choose $\gamma_0 = \alpha$ and $\gamma_N = \beta$.

As in the previous section the reindexing map is the identity. It follows immediately that $\iota(\gamma_n, \mu_n) \leq 4$. This bound, the bound on $\iota(\mu_n, \mu_{n+1})$, and Lemma 4.7 imply that $\iota(\gamma_n, \gamma_{n+1})$ is likewise bounded. The usual surgery argument shows that if two arcs have bounded intersection then they have bounded distance. This verifies Axiom 13.5.

17.3. The replacement and the straight axioms. Suppose that $Y \subset X$ is a subsurface and γ_n has $n \in J_Y$. Let $\mu_n = \nu_t$; that is t is the time when μ_n is a short marking. Thus $\partial Y \subset \text{base}(\mu_n)$ and so $\gamma_n \cap \partial Y = \emptyset$. So regardless of the hole-nature of Y we may take $\gamma' = \gamma_n$ and the axiom is verified.

Axiom 13.11 is verified exactly as in Section 16.

17.4. Non-orientable surfaces. Suppose that F is non-orientable and Δ_F is a collection of boundary components. Let S be the orientation double cover and $\tau: S \rightarrow S$ the involution so that $S/\tau = F$. Let Δ be the preimage of Δ_F . Then $\mathcal{A}^\tau(S, \Delta)$ is the invariant arc complex.

Suppose that α_F, β_F are vertices in $\mathcal{A}(F, \Delta')$. Let α, β be their preimages. As above, without loss of generality, we may assume that $\sigma_F(\alpha_F)$ and $\sigma_F(\beta_F)$ fill F . Note that $\sigma_F(\alpha_F)$ cuts a surface X off of F . The surface X is either a pants or a twice-holed \mathbb{RP}^2 . When X is a pants we define $\alpha' \subset S$ to be the preimage of $\sigma_F(\alpha_F)$. When X is a

twice-holed \mathbb{RP}^2 we take γ_F to be a core of one of the two Möbius bands contained in X and we define α' to be the preimage of $\gamma_F \cup \sigma_F(\alpha_F)$. We define β' similarly. Notice that α and α' meet in at most four points.

We now use α' and β' to build a τ -invariant Teichmüller geodesic. The construction of the marking and combinatorial paths for $\mathcal{A}^\tau(S, \Delta)$ is unchanged. Notice that we may choose combinatorial vertices because $\text{base}(\mu_n)$ is τ -invariant. There is a small annoyance: when X is a twice-holed \mathbb{RP}^2 the first vertex, γ_0 , is disjoint from but not equal to α . Strictly speaking, the first and last vertices are γ_0 and γ_N ; our constants are stated in terms of their subsurface projection distances. However, since $\alpha \cap \gamma_0 = \emptyset$, and the same holds for β, γ_N , their subsurface projection distances are all bounded.

18. BACKGROUND ON TRAIN TRACKS

Here we give the necessary definitions and theorems regarding train tracks. The standard reference is [31]. See also [30]. We follow closely the discussion found in [27].

18.1. On tracks. A *generic train track* $\tau \subset S$ is a smooth, embedded trivalent graph. As usual we call the vertices *switches* and the edges *branches*. At every switch the tangents of the three branches agree. Also, there are exactly two *incoming* branches and one *outgoing* branch at each switch. See Figure 7 for the local model of a switch.




FIGURE 7. The local model of a train track.

Let $\mathcal{B}(\tau)$ be the set of branches. A transverse measure on τ is function $w: \mathcal{B} \rightarrow \mathbb{R}_{\geq 0}$ satisfying the switch conditions: at every switch the sum of the incoming measures equals the outgoing measure. Let $P(\tau)$ be the projectivization of the cone of transverse measures. Let $V(\tau)$ be the vertices of $P(\tau)$. As discussed in the references, each vertex measure gives a simple closed curve carried by τ .

For every track τ we refer to $V(\tau)$ as the marking corresponding to τ (see Section 2.4). Note that there are only finitely many tracks up to the action of the mapping class group. It follows that $\iota(V(\tau))$ is uniformly bounded, depending only on the topological type of S .

If τ and σ are train tracks, and $Y \subset S$ is an essential surface, then define

$$d_Y(\tau, \sigma) = d_Y(V(\tau), V(\sigma)).$$

We also adopt the notation $\pi_Y(\tau) = \pi_Y(V(\tau))$.

A train track σ is obtained from τ by *sliding* if σ and τ are related as in Figure 8. We say that a train track σ is obtained from τ by *splitting* if σ and τ are related as in Figure 9.




FIGURE 8. All slides take place in a small regular neighborhood of the affected branch.




FIGURE 9. There are three kinds of splitting: right, left, and central.

Again, since the number of tracks is bounded (up to the action of the mapping class group) if σ is obtained from τ by either a slide or a split we find that $\iota(V(\tau), V(\sigma))$ is uniformly bounded.

18.2. The marking path. We will use sequences of train tracks to define our marking path.

Definition 18.1. A *sliding and splitting sequence* is a collection $\{\tau_n\}_{n=0}^N$ of train tracks so that τ_{n+1} is obtained from τ_n by a slide or a split.

The sequence $\{\tau_n\}$ gives a sequence of markings via the map $\tau_n \mapsto V_n = V(\tau_n)$. Note that the support of V_{n+1} is contained within the support of V_n because every vertex of τ_{n+1} is carried by τ_n . Theorem 5.5 of [27] verifies the remaining half of Axiom 13.3.

Theorem 18.2. *Fix a surface S . There is a constant A with the following property. Suppose that $\{\tau_n\}_{n=0}^N$ is a sliding and splitting sequence in S of birecurrent tracks. Suppose that $Y \subset S$ is an essential surface. Then the map $n \mapsto \pi_Y(\tau_n)$, as parametrized by splittings, is an A -unparametrized quasi-geodesic. \square*

Note that, when $Y = S$, Theorem 18.2 is essentially due to the first author and Minsky; see Theorem 1.3 of [26].

In Section 5.2 of [27], for every sliding and splitting sequence $\{\tau_n\}_{n=0}^N$ and for any essential subsurface $X \subsetneq S$ an accessible interval $I_X \subset [0, N]$ is defined. Axiom 13.4 is now verified by Theorem 5.3 of [27].

18.3. Quasi-geodesics in the marking graph. We will also need Theorem 6.1 from [27]. (See [16] for closely related work.)

Theorem 18.3. *Fix a surface S . There is a constant A with the following property. Suppose that $\{\tau_n\}_{n=0}^N$ is a sliding and splitting sequence of birecurrent tracks, injective on slide subsequences, where V_N fills S . Then $\{V(\tau_n)\}$ is an A -quasi-geodesic in the marking graph. \square*

19. PATHS FOR THE DISK COMPLEX

Suppose that $V = V_g$ is a genus g handlebody. The goal of this section is to verify the axioms of Section 13 for the disk complex $\mathcal{D}(V)$ and so complete the proof of the distance estimate.

Theorem 19.1. *There is a constant $C_0 = C_0(V)$ so that, for any $c \geq C_0$ there is a constant A with*

$$d_{\mathcal{D}}(D, E) =_A \sum [d_X(D, E)]_c$$

independent of the choice of D and E . Here the sum ranges over the set of holes $X \subset \partial V$ for the disk complex.

19.1. Holes. The fact that all large holes interfere is recorded above as Lemma 12.13. This verifies Axiom 13.2.

19.2. The combinatorial path. Suppose that $D, E \in \mathcal{D}(V)$ are disks contained in a compressible hole $X \subset S = \partial V$. As usual we may assume that D and E fill X . Recall that $V(\tau)$ is the set of vertices for the track $\tau \subset X$. We now appeal to a result of the first author and Minsky, found in [26].

Theorem 19.2. *There exists a surgery sequence of disks $\{D_i\}_{i=0}^K$, a sliding and splitting sequence of birecurrent tracks $\{\tau_n\}_{n=0}^N$, and a reindexing function $r: [0, K] \rightarrow [0, N]$ so that*

- $D_0 = D$,

- $E \in V_N$,
- $D_i \cap D_{i+1} = \emptyset$ for all i , and
- $\iota(\partial D_i, V_{r(i)})$ is uniformly bounded for all i .

□

Remark 19.3. For the details of the proof we refer to [26]. Note that the double-wave curve replacements of that paper are not needed here; as X is a hole, no curve of ∂X compresses in V . It follows that consecutive disks in the surgery sequence are disjoint (as opposed to meeting at most four times). Also, in the terminology of [27], the disk D_i is a *wide dual* for the track $\tau_{r(i)}$. Finally, recurrence of τ_n follows because E is fully carried by τ_N . Transverse recurrence follows because D is fully dual to τ_0 .

Thus V_n will be our marking path and D_i will be our combinatorial path. The requirements of Axiom 13.5 are now verified by Theorem 19.2.

19.3. The replacement axiom. We turn to Axiom 13.6. Suppose that $Y \subset X$ is an essential subsurface and D_i has $r(i) \in J_Y$. Let $n = r(i)$. From Theorem 19.2 we have that $\iota(\partial D_i, V_n)$ is uniformly bounded. By Axiom 13.4 we have $Y \subset \text{supp}(V_n)$ and $\iota(\partial Y, \mu_n)$ is bounded. It follows that there is a constant K depending only on $\xi(S)$ so that

$$\iota(\partial D_i, \partial Y) < K.$$

Isotope D_i to have minimal intersection with ∂Y . As in Section 11.1 boundary compress D_i as much as possible into the components of $X \setminus \partial Y$ to obtain a disk D' so that either

- D' cannot be boundary compressed any more into $X \setminus \partial Y$ or
- D' is disjoint from ∂Y .

We may arrange matters so that every boundary compression reduces the intersection with ∂Y by at least a factor of two. Thus:

$$d_{\mathcal{D}}(D_i, D') \leq \log_2(K).$$

Suppose now that Y is a compressible hole. By Lemma 8.4 we find that $\partial D' \subset Y$ and we are done.

Suppose now that Y is an incompressible hole. Since Y is large there is an I -bundle $T \rightarrow F$, contained in the handlebody V , so that Y is a component of $\partial_h T$. Isotope D' to minimize intersection with $\partial_v T$. Let Δ be the union of components of $\partial_v T$ which are contained in ∂V . Let $\Gamma = \partial_v T \setminus \Delta$. Notice that all intersections $D' \cap \Gamma$ are essential arcs in Γ : simple closed curves are ruled out by minimal intersection and inessential arcs are ruled out by the fact that D' cannot be boundary compressed in the complement of ∂Y . Let D'' be a outermost component

of $D' \setminus \Gamma$. Then Lemma 8.5 implies that D'' is isotopic in T to a vertical disk.

If $D'' = D'$ then we may replace D_i by the arc $\rho_F(D')$. The inductive argument now occurs inside of the arc complex $\mathcal{A}(F, \rho_F(\Delta))$.

Suppose that $D'' \neq D'$. Let $A \in \Gamma$ be the vertical annulus meeting D'' . Let N be a regular neighborhood of $D'' \cup A$, taken in T . Then the frontier of N in T is again a vertical disk, call it D''' . Note that $\iota(D''', D') < K - 1$. Finally, replace D_i by the arc $\rho_F(D''')$.

Suppose now that Y is not a hole. Then some component $S \setminus Y$ is compressible. Applying Lemma 8.4 again, we find that either D' lies in $Z = X \setminus Y$ or in Y . This completes the verification of Axiom 13.6.

19.4. Straight intervals. We end by checking Axiom 13.11. Suppose that $[p, q] \subset [0, K]$ is a straight interval. Recall that $d_Y(\mu_{r(p)}, \mu_{r(q)}) < L_2$ for all strict subsurfaces Y . We must check that $d_{\mathcal{D}}(D_p, D_q) \leq_A d_S(D_p, D_q)$. Since $d_{\mathcal{D}}(D_p, D_q) \leq C_2|p - q|$ it is enough to bound $|p - q|$. Note that $|p - q| \leq |r(p) - r(q)|$ because the reindexing map is increasing. Now, $|r(p) - r(q)| \leq_A d_{\mathcal{M}}(\mu_{r(p)}, \mu_{r(q)})$ because the sequence $\{\mu_n\}$ is a quasi-geodesic in $\mathcal{M}(S)$ (Theorem 18.3). Increasing A as needed and applying Theorem 4.10 we have

$$d_{\mathcal{M}}(\mu_{r(p)}, \mu_{r(q)}) \leq_A \sum_Y [d_Y(\mu_{r(p)}, \mu_{r(q)})]_{L_2}$$

and the right hand side is thus less than $d_S(\mu_{r(p)}, \mu_{r(q)})$ which in turn is less than $d_S(D_p, D_q) + 2C_2$. This completes our discussion of Axiom 13.11 and finishes the proof of Theorem 19.1.

20. HYPERBOLICITY

The ideas in this section are related to the notion of “time-ordered domains” and to the hierarchy machine of [25] (see also Chapters 4 and 5 of Behrstock’s thesis [1]). As remarked above, we cannot use those tools directly as the hierarchy machine is too rigid to deal with the disk complex.

20.1. Hyperbolicity. We prove:

Theorem 20.1. *Fix $\mathcal{G} = \mathcal{G}(S)$, a combinatorial complex. Suppose that \mathcal{G} satisfies the axioms of Section 13. Then \mathcal{G} is Gromov hyperbolic.*

As corollaries we have

Theorem 20.2. *The arc complex is Gromov hyperbolic.* □

Theorem 20.3. *The disk complex is Gromov hyperbolic.* □

In fact, Theorem 20.1 follows quickly from:

Theorem 20.4. *Fix \mathcal{G} , a combinatorial complex. Suppose that \mathcal{G} satisfies the axioms of Section 13. Then for all $A \geq 1$ there exists $\delta \geq 0$ with the following property: Suppose that $T \subset \mathcal{G}$ is a triangle of paths where the projection of any side of T into any hole is an A -unparametrized quasi-geodesic. Then T is δ -slim.*

Proof of Theorem 20.1. As laid out in Section 14 there is a uniform constant A so that for any pair $\alpha, \beta \in \mathcal{G}$ there is a recursively constructed path $\mathcal{P} = \{\gamma_i\} \subset \mathcal{G}$ so that

- for any hole X for \mathcal{G} , the projection $\pi_X(\mathcal{P})$ is an A -unparametrized quasi-geodesic and
- $|\mathcal{P}| =_A d_{\mathcal{G}}(\alpha, \beta)$.

So if $\alpha \cap \beta = \emptyset$ then $|\mathcal{P}|$ is uniformly short. Also, by Theorem 20.4, triangles made of such paths are uniformly slim. Thus, by Theorem 3.11, \mathcal{G} is Gromov hyperbolic. \square

The rest of this section is devoted to proving Theorem 20.4.

20.2. Index in a hole. For the following definitions, we assume that α and β are fixed vertices of \mathcal{G} .

For any hole X and for any geodesic $k \in \mathcal{C}(X)$ connecting a point of $\pi_X(\alpha)$ to a point of $\pi_X(\beta)$ we also define $\rho_k: \mathcal{G} \rightarrow k$ to be the relation $\pi_X|_{\mathcal{G}}: \mathcal{G} \rightarrow \mathcal{C}(X)$ followed by taking closest points in k . Since the diameter of $\rho_k(\gamma)$ is uniformly bounded, we may simplify our formulas by treating ρ_k as a function. Define $\text{index}_X: \mathcal{G} \rightarrow \mathbb{N}$ to be the *index* in X :

$$\text{index}_X(\sigma) = d_X(\alpha, \rho_k(\sigma)).$$

Remark 20.5. Suppose that k' is a different geodesic connecting $\pi_X(\alpha)$ to $\pi_X(\beta)$ and index'_X is defined with respect to k' . Then

$$|\text{index}_X(\sigma) - \text{index}'_X(\sigma)| \leq 17\delta + 4$$

by Lemma 3.7 and Lemma 3.8. After permitting a small additive error, the index depends only on α, β, σ and not on the choice of geodesic k .

20.3. Back and sidetracking. Fix $\sigma, \tau \in \mathcal{G}$. We say σ *precedes* τ by at least K in X if

$$\text{index}_X(\sigma) + K \leq \text{index}_X(\tau).$$

We say σ *precedes* τ by at most K if the inequality is reversed. If σ precedes τ then we say τ *succeeds* σ .

Now take $\mathcal{P} = \{\sigma_i\}$ to be a path in \mathcal{G} connecting α to β . Recall that we have made the simplifying assumption that σ_i and σ_{i+1} are disjoint.

We formalize a pair of properties enjoyed by unparametrized quasi-geodesics. The path \mathcal{P} *backtracks* at most K if for every hole X and all indices $i < j$ we find that σ_j precedes σ_i by at most K . The path \mathcal{P} *sidetracks* at most K if for every hole X and every index i we find that

$$d_X(\sigma_i, \rho_k(\sigma_i)) \leq K,$$

for some geodesic k connecting a point of $\pi_X(\alpha)$ to a point of $\pi_X(\beta)$.

Remark 20.6. As in Remark 20.5, allowing a small additive error makes irrelevant the choice of geodesic in the definition of sidetracking. We note that, if \mathcal{P} has bounded sidetracking, one may freely use in calculation whichever of σ_i or $\rho_k(\sigma_i)$ is more convenient.

20.4. Projection control. We say domains $X, Y \subset S$ *overlap* if ∂X cuts Y and ∂Y cuts X . The following lemma, due to Behrstock [1, 4.2.1], is closely related to the notion of *time ordered* domains [25]. An elementary proof is given in [23, Lemma 2.5].

Lemma 20.7. *There is a constant $M_1 = M_1(S)$ with the following property. Suppose that X, Y are overlapping non-simple domains. If $\gamma \in \mathcal{AC}(S)$ cuts both X and Y then either $d_X(\gamma, \partial Y) < M_1$ or $d_Y(\partial X, \gamma) < M_1$. \square*

We also require a more specialized version for the case where X and Y are nested.

Lemma 20.8. *There is a constant $M_2 = M_2(S)$ with the following property. Suppose that $X \subset Y$ are nested non-simple domains. Fix $\alpha, \beta, \gamma \in \mathcal{AC}(S)$ that cut X . Fix $k = [\alpha', \beta'] \subset \mathcal{C}(Y)$, a geodesic connecting a point of $\pi_Y(\alpha)$ to a point of $\pi_Y(\beta)$. Assume that $d_X(\alpha, \beta) \geq M_0$, the constant given by Theorem 4.6. If $d_X(\alpha, \gamma) \geq M_2$ then*

$$\text{index}_Y(\partial X) - 4 \leq \text{index}_Y(\gamma).$$

Symmetrically, we have

$$\text{index}_Y(\gamma) \leq \text{index}_Y(\partial X) + 4$$

if $d_X(\gamma, \beta) \geq M_2$. \square

20.5. Finding the midpoint of a side. Fix $A \geq 1$. Let $\mathcal{P}, \mathcal{Q}, \mathcal{R}$ be the sides of a triangle in \mathcal{G} with vertices at α, β, γ . We assume that each of \mathcal{P}, \mathcal{Q} , and \mathcal{R} are A -unparametrized quasi-geodesics when projected to any hole.

Recall that $M_0 = M_0(S)$, $M_1 = M_1(S)$, and $M_2 = M_2(S)$ are functions depending only on the topology of S . We may assume that if $T \subset S$ is an essential subsurface, then $M_0(S) > M_0(T)$.

Now choose $K_1 \geq \max\{M_0, 4M_1, M_2, 8\} + 8\delta$ sufficiently large so that any A -unparametrized quasi-geodesic in any hole back and side tracks at most K_1 .

Claim 20.9. If σ_i precedes γ in X and σ_j succeeds γ in Y , both by at least $2K_1$, then $i < j$.

Proof. To begin, as X and Y are holes and all holes interfere, we need not consider the possibility that $X \cap Y = \emptyset$. If $X = Y$ we immediately deduce that

$$\text{index}_X(\sigma_i) + 2K_1 \leq \text{index}_X(\gamma) \leq \text{index}_X(\sigma_j) - 2K_1.$$

Thus $\text{index}_X(\sigma_i) + 4K_1 \leq \text{index}_X(\sigma_j)$. Since \mathcal{P} backtracks at most K_1 we have $i < j$, as desired.

Suppose instead that $X \subset Y$. Since σ_i precedes γ in X we immediately find $d_X(\alpha, \beta) \geq 2K_1 \geq M_0$ and $d_X(\alpha, \gamma) \geq 2K_1 - 2\delta \geq M_2$. Apply Lemma 20.8 to deduce $\text{index}_Y(\partial X) - 4 \leq \text{index}_Y(\gamma)$. Since σ_j succeeds γ in Y it follows that $\text{index}_Y(\partial X) - 4 + 2K_1 \leq \text{index}_Y(\sigma_j)$. Again using the fact that σ_i precedes γ in X we have that $d_X(\sigma_i, \beta) \geq M_2$. We deduce from Lemma 20.8 that $\text{index}_Y(\sigma_i) \leq \text{index}_Y(\partial X) + 4$. Thus

$$\text{index}_Y(\sigma_i) - 8 + 2K_1 \leq \text{index}_Y(\sigma_j).$$

Since \mathcal{P} backtracks at most K_1 in Y we again deduce that $i < j$. The case where $Y \subset X$ is similar.

Suppose now that X and Y overlap. Applying Lemma 20.7 and breaking symmetry, we may assume that $d_X(\gamma, \partial Y) < M_1$. Since σ_i precedes γ we have $\text{index}_X(\gamma) \geq 2K_1$. Lemma 3.7 now implies that $\text{index}_X(\partial Y) \geq 2K_1 - M_1 - 6\delta$. Thus,

$$d_X(\alpha, \partial Y) \geq 2K_1 - M_1 - 8\delta \geq M_1$$

where the first inequality follows from Lemma 3.4.

Applying Lemma 20.7 again, we find that $d_Y(\alpha, \partial X) < M_1$. Now, since σ_j succeeds γ in Y , we deduce that $\text{index}_Y(\sigma_j) \geq 2K_1$. So Lemma 3.4 implies that $d_Y(\alpha, \sigma_j) \geq 2K_1 - 2\delta$. The triangle inequality now gives

$$d_Y(\partial X, \sigma_j) \geq 2K_1 - M_1 - 2\delta \geq M_1.$$

Applying Lemma 20.7 one last time, we find that $d_X(\partial Y, \sigma_j) < M_1$. Thus $d_X(\gamma, \sigma_j) \leq 2M_1$. Finally, Lemma 3.7 implies that the difference in index (in X) between σ_i and σ_j is at least $2K_1 - 2M_1 - 6\delta$. Since this is greater than the backtracking constant, K_1 , it follows that $i < j$. \square

Let $\sigma_\alpha \in \mathcal{P}$ be the *last* vertex of \mathcal{P} preceding γ by at least $2K_1$ in some hole. If no such vertex of \mathcal{P} exists then take $\sigma_\alpha = \alpha$.

Claim 20.10. For every hole X and geodesic h connecting $\pi_X(\alpha)$ to $\pi_X(\beta)$:

$$d_X(\sigma_\alpha, \rho_h(\gamma)) \leq 3K_1 + 6\delta + 1$$

Proof. Since σ_i and σ_{i+1} are disjoint we have $d_X(\sigma_i, \sigma_{i+1}) \geq 3$ and so Lemma 3.7 implies that

$$|\text{index}_X(\sigma_{i+1}) - \text{index}_X(\sigma_i)| \leq 6\delta + 3.$$

Since \mathcal{P} is a path connecting α to β the image $\rho_h(\mathcal{P})$ is $6\delta + 3$ -dense in h . Thus, if $\text{index}_X(\sigma_\alpha) + 2K_1 + 6\delta + 3 < \text{index}_X(\gamma)$ then we have a contradiction to the definition of σ_α .

On the other hand, if $\text{index}_X(\sigma_\alpha) \geq \text{index}_X(\gamma) + 2K_1$ then σ_α precedes and succeeds γ in X . This directly contradicts Claim 20.9.

We deduce that the difference in index between σ_α and γ in X is at most $2K_1 + 6\delta + 3$. Finally, as \mathcal{P} sidetracks by at most K_1 we have

$$d_X(\sigma_\alpha, \rho_h(\gamma)) \leq 3K_1 + 6\delta + 3$$

as desired. \square

We define σ_β to be the first σ_i to succeed γ by at least $2K_1$ — if no such vertex of \mathcal{P} exists take $\sigma_\beta = \beta$. If $\alpha = \beta$ then $\sigma_\alpha = \sigma_\beta$. Otherwise, from Claim 20.9, we immediately deduce that σ_α comes before σ_β in \mathcal{P} . A symmetric version of Claim 20.10 applies to σ_β : for every hole X

$$d_X(\rho_h(\gamma), \sigma_\beta) \leq 3K_1 + 6\delta + 3.$$

20.6. Another side of the triangle. Recall now that we are also given a path $\mathcal{R} = \{\tau_i\}$ connecting α to γ in \mathcal{G} . As before, \mathcal{R} has bounded back and sidetracking. Thus we again find vertices τ_α and τ_γ the last/first to precede/succeed β by at least $2K_1$. Again, this is defined in terms of the closest points projection of β to a geodesic of the form $h = [\pi_X(\alpha), \pi_X(\gamma)]$. By Claim 20.10, for every hole X , τ_α and τ_γ are close to $\rho_h(\beta)$.

By Lemma 3.6, if $k = [\pi_X(\alpha), \pi_X(\beta)]$, then $d_X(\rho_k(\gamma), \rho_h(\beta)) \leq 6\delta$. We deduce:

Claim 20.11. $d_X(\sigma_\alpha, \tau_\alpha) \leq 6K_1 + 18\delta + 2$. \square

This claim and Claim 20.10 imply that the body of the triangle \mathcal{PQR} is bounded in size. We now show that the legs are narrow.

Claim 20.12. There is a constant $N_2 = N_2(S)$ with the following property. For every $\sigma_i \leq \sigma_\alpha$ in \mathcal{P} there is a $\tau_j \leq \tau_\alpha$ in \mathcal{R} so that

$$d_X(\sigma_i, \tau_j) \leq N_2$$

for every hole X .

Proof. We only sketch the proof, as the details are similar to our previous discussion. Fix $\sigma_i \leq \sigma_\alpha$.

Suppose first that no vertex of \mathcal{R} precedes σ_i by more than $2K_1$ in any hole. So fix a hole X and geodesics $k = [\pi_X(\alpha), \pi_X(\beta)]$ and $h = [\pi_X(\alpha), \pi_X(\gamma)]$. Then $\rho_h(\sigma_i)$ is within distance $2K_1$ of $\pi_X(\alpha)$. Appealing to Claim 20.11, bounded sidetracking, and hyperbolicity of $\mathcal{C}(X)$ we find that the initial segments

$$[\pi_X(\alpha), \rho_k(\sigma_\alpha)], \quad [\pi_X(\alpha), \rho_h(\tau_\alpha)]$$

of k and h respectively must fellow travel. Because of bounded backtracking along \mathcal{P} , $\rho_k(\sigma_i)$ lies on, or at least near, this initial segment of k . Thus by Lemma 3.8 $\rho_h(\sigma_i)$ is close to $\rho_k(\sigma_i)$ which in turn is close to $\pi_X(\sigma_i)$, because \mathcal{P} has bounded sidetracking. In short, $d_X(\alpha, \sigma_i)$ is bounded for all holes X . Thus we may take $\tau_j = \tau_0 = \alpha$ and we are done.

Now suppose that some vertex of \mathcal{R} precedes σ_i by at least $2K_1$ in some hole X . Take τ_j to be the last such vertex in \mathcal{R} . Following the proof of Claim 20.9 shows that τ_j comes before τ_α in \mathcal{R} . The argument now required to bound $d_X(\sigma_i, \tau_j)$ is essentially identical to the proof of Claim 20.10. \square

By the distance estimate, we find that there is a uniform neighborhood of $[\sigma_0, \sigma_\alpha] \subset \mathcal{P}$, taken in \mathcal{G} , which contains $[\tau_0, \tau_\alpha] \subset \mathcal{P}$. The slimness of \mathcal{PQR} follows directly. This completes the proof of Theorem 20.4. \square

21. COARSELY COMPUTING HEMPEL DISTANCE

We now turn to our topological application. Recall that a *Heegaard splitting* is a triple (S, V, W) consisting of a surface and two handlebodies where $V \cap W = \partial V = \partial W = S$. Hempel [20] defines the quantity

$$d_S(V, W) = \min \{d_S(D, E) \mid D \in \mathcal{D}(V), E \in \mathcal{D}(W)\}$$

and calls it the *distance* of the splitting. Note that a splitting can be completely determined by giving a pair of cut systems: simplices $\mathbb{D} \subset \mathcal{D}(V)$, $\mathbb{E} \subset \mathcal{D}(W)$ where the corresponding disks cut the containing handlebody into a single three-ball. The triple $(S, \mathbb{D}, \mathbb{E})$ is a *Heegaard diagram*. The goal of this section is to prove:

Theorem 21.1. *There is a constant $R_1 = R_1(S)$ and an algorithm that, given a Heegaard diagram $(S, \mathbb{D}, \mathbb{E})$, computes a number N so that*

$$|d_S(V, W) - N| \leq R_1.$$

Let $\rho_V: \mathcal{C}(S) \rightarrow \mathcal{D}(V)$ be the closest points relation:

$$\rho_V(\alpha) = \{D \in \mathcal{D}(V) \mid \text{for all } E \in \mathcal{D}(V), d_S(\alpha, D) \leq d_S(\alpha, E)\}.$$

Theorem 21.1 follows from:

Theorem 21.2. *There is a constant $R_0 = R_0(V)$ and an algorithm that, given an essential curve $\alpha \subset S$ and a cut system $\mathbb{D} \subset \mathcal{D}(V)$, finds a disk $C \in \mathcal{D}(V)$ so that*

$$d_S(C, \rho_V(\alpha)) \leq R_0.$$

Proof of Theorem 21.1. Suppose that $(S, \mathbb{D}, \mathbb{E})$ is a Heegaard diagram. Using Theorem 21.2 we find a disk D within distance R_0 of $\rho_V(\mathbb{E})$. Again using Theorem 21.2 we find a disk E within distance R_0 of $\rho_W(D)$. Notice that E is defined using D and not the cut system \mathbb{D} .

Since computing distance between fixed vertices in the curve complex is algorithmic [22, 37] we may compute $d_S(D, E)$. By the hyperbolicity of $\mathcal{C}(S)$ (Theorem 3.2) and by the quasi-convexity of the disk set (Theorem 4.9) this is the desired estimate. \square

Very briefly, the algorithm asked for in Theorem 21.2 searches an R_2 -neighborhood in $\mathcal{M}(S)$ about a splitting sequence from \mathbb{D} to α . Here are the details.

Algorithm 21.3. We are given $\alpha \in \mathcal{C}(S)$ and a cut system $\mathbb{D} \subset \mathcal{D}(V)$. Build a train track τ in $S = \partial V$ as follows: make \mathbb{D} and α tight. Place one switch on every disk $D \in \mathbb{D}$. Homotope all intersections of α with D to run through the switch. Collapse bigons of α inside of $S \setminus \mathbb{D}$ to create the branches. Now make τ a generic track by combing away from \mathbb{D} [31, Proposition 1.4.1]. Note that α is carried by τ and so gives a transverse measure w .

Build a splitting sequence of measured tracks $\{\tau_n\}_{n=0}^N$ where $\tau_0 = \tau$, $\tau_N = \alpha$, and τ_{n+1} is obtained by splitting the largest switch of τ_n (as determined by the measure imposed by α).

Let $\mu_n = V(\tau_n)$ be the vertices of τ_n . For each filling marking μ_n list all markings in the ball $B(\mu_n, R_2) \subset \mathcal{M}(S)$, where R_2 is given by Lemma 21.5 below. (If μ_0 does not fill S then output \mathbb{D} and halt.)

For every marking ν so produced we use Whitehead's algorithm (see Lemma 21.4) to try and find a disk meeting some curve $\gamma \in \nu$ at most twice. For every disk C found compute $d_S(\alpha, C)$ [22, 37]. Finally, output any disk which minimizes this distance, among all disks considered, and halt.

We use the following form of Whitehead's algorithm [3]:

Lemma 21.4. *There is an algorithm that, given a cut system $\mathbb{D} \subset V$ and a curve $\gamma \subset S$, outputs a disk $C \subset V$ so that $\iota(\gamma, \partial C) = \min\{\iota(\gamma, \partial E) \mid E \in \mathcal{D}(V)\}$. \square*

We now discuss the constant R_2 . We begin by noticing that the track τ_n is transversely recurrent because α is fully carried and \mathbb{D} is fully dual. Thus by Theorem 18.2 and by Morse stability, for any essential $Y \subset S$ there is a stability constant M_3 for the path $\pi_Y(\mu_n)$. Let δ be the hyperbolicity constant for $\mathcal{C}(S)$ (Theorem 3.2) and let Q be the quasi-convexity constant for $\mathcal{D}(V) \subset \mathcal{C}(S)$ (Theorem 4.9).

Since $\iota(\mathbb{D}, \mu_0)$ is bounded we will, at the cost of an additive error, identify their images in $\mathcal{C}(S)$. Now, for every n pick some $E_n \in \rho_V(\mu_n)$.

Lemma 21.5. *There is a constant R_2 with the following property. Suppose that $n < m$, $d_S(\mu_n, E_n), d_S(\mu_m, E_m) \leq M_3 + \delta + Q$, and $d_S(\mu_n, \mu_m) \geq 2(M_3 + \delta + Q) + 5$. Then there is a marking $\nu \in B(\mu_n, R_2)$ and a curve $\gamma \in \nu$ so that either:*

- γ bounds a disk in V ,
- $\gamma \subset \partial Z$, where Z is a non-hole or
- $\gamma \subset \partial Z$, where Z is a large hole.

Proof of Lemma 21.5. Choose points σ, σ' in the thick part of $\mathcal{T}(S)$ so that all curves of μ_n have bounded length in σ and so that E_n has length less than the Margulis constant in σ' . As in Section 15 there is a Teichmüller geodesic and associated markings $\{\nu_k\}_{k=0}^K$ so that $d_{\mathcal{M}}(\nu_0, \mu_n)$ is bounded and $E_n \in \text{base}(\nu_K)$.

We say a hole $X \subset S$ is *small* if $\text{diam}_X(\mathcal{D}(V)) < 61$.

Claim. There is a constant R_3 so that for any small hole X we have $d_X(\mu_n, \nu_K) < R_3$.

Proof. If $d_X(\mu_n, \nu_K) \leq M_0$ then we are done. If the distance is greater than M_0 then Theorem 4.6 gives a vertex of the $\mathcal{C}(S)$ -geodesic connecting μ_n to E_n with distance at most one from ∂X . It follows from the triangle inequality that every vertex of the $\mathcal{C}(S)$ -geodesic connecting μ_m to E_m cuts X . Another application of Theorem 4.6 gives

$$d_X(\mu_m, E_m) < M_0.$$

Since X is small $d_X(E_m, \mathbb{D}), d_X(E_n, \mathbb{D}) \leq 60$. Since $\iota(\nu_K, E_n) = 2$ the distance $d_X(\nu_K, E_n)$ is bounded.

Finally, because $p \mapsto \pi_X(\mu_p)$ is an A -unparameterized quasi-geodesic in $\mathcal{C}(X)$ it follows that $d_X(\mathbb{D}, \mu_n)$ is also bounded and the claim is proved. \square

Now consider all strict subsurfaces Y so that

$$d_Y(\mu_n, \nu_M) \geq R_3.$$

None of these are small holes, by the claim above. If there are no such surfaces then Theorem 4.10 bounds $d_{\mathcal{M}}(\mu_n, \nu_M)$: taking the cutoff

constant larger than

$$\max\{R_3, C_0, M_3 + \delta + Q\}$$

ensures that all terms on the right-hand side vanish. In this case the additive error in Theorem 4.10 is the desired constant R_2 and the lemma is proved.

If there are such surfaces then choose one, say Z , that minimizes $\ell = \min J_Z$. Thus $d_Y(\mu_n, \nu_\ell) < C_3$ for all strict non-holes and all strict large holes. Since $d_S(\mu_n, E_n) \leq M_3 + \delta + Q$ and $\{\nu_m\}$ is an unparametrized quasi-geodesic [33, Theorem 6.1] we find that $d_S(\mu_n, \nu_\ell)$ is uniformly bounded. The claim above bounds distances in small holes. As before we find a sufficiently large cutoff so that all terms on the right-hand side of Theorem 4.10 vanish. Again the additive error of Theorem 4.10 provides the constant R_2 . Since $\partial Z \subset \text{base}(\nu_\ell)$ the lemma is proved. \square

To prove the correctness of Algorithm 21.3 it suffices to show that the disk produced is close to $\rho_V(\alpha)$. Let m be the largest index so that for all $n \leq m$ we have

$$d_S(\mu_n, E_n) \leq M_3 + \delta + Q.$$

It follows that μ_{m+1} lies within distance $M_3 + \delta$ of the geodesic $[\alpha, \rho_V(\alpha)]$. Recall that $d_S(\mu_n, \mu_{n+1}) \leq C_1$ for any value of n . A shortcut argument shows that

$$d_S(\mu_m, \rho_V(\alpha)) \leq 2C_1 + 3M_3 + 3\delta + Q.$$

Let $n \leq m$ be the largest index so that

$$2(M_3 + \delta + Q) + 5 \leq d_S(\mu_n, \mu_m).$$

If no such n exists then take $n = 0$. Now, Lemma 21.5 implies that there is a disk C with $d_S(C, \mu_n) \leq 4R_2$ and this disk is found during the running of Algorithm 21.3. It follows from the above inequalities that

$$d_S(C, \alpha) \leq 4R_2 + 5M_3 + 5\delta + 3Q + 5 + 2C_1 + d_S(\alpha, \rho_V(\alpha)).$$

So the disk C' , output by the algorithm, is at least this close to α in $\mathcal{C}(S)$. Examining the triangle with vertices $\alpha, \rho_V(\alpha), C'$ and using a final short-cut argument gives

$$d_S(C', \rho_V(\alpha)) \leq 4R_2 + 5M_3 + 9\delta + 5Q + 5 + 2C_1.$$

This completes the proof of Theorem 21.2. \square

REFERENCES

- [1] Jason Behrstock. *Asymptotic geometry of the mapping class group and Teichmüller space*. PhD thesis, SUNY Stony Brook, 2004. <http://www.math.columbia.edu/~jason/thesis.pdf>. [62, 64]
- [2] Jason Behrstock, Cornelia Drutu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. arXiv:math/0512592. [13]
- [3] John Berge. Heegaard documentation. documentation of computer program available at computop.org. [68]
- [4] Mladen Bestvina and Koji Fujiwara. Quasi-homomorphisms on mapping class groups. *Glas. Mat. Ser. III*, 42(62)(1):213–236, 2007. arXiv:math/0702273. [3, 16]
- [5] Joan S. Birman. The topology of 3-manifolds, Heegaard distance and the mapping class group of a 2-manifold. In *Problems on mapping class groups and related topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 133–149. Amer. Math. Soc., Providence, RI, 2006. <http://www.math.columbia.edu/~jb/papers.html>. [2]
- [6] Brian H. Bowditch. Intersection numbers and the hyperbolicity of the curve complex. *J. Reine Angew. Math.*, 598:105–129, 2006. bhb-curvecomplex.pdf. [3, 54]
- [7] Tara E. Brendle and Dan Margalit. Commensurations of the Johnson kernel. *Geom. Topol.*, 8:1361–1384 (electronic), 2004. arXiv:math/0404445. [13]
- [8] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999. [7]
- [9] Jeffrey F. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.*, 16(3):495–535 (electronic), 2003. arXiv:math/0109048. [13]
- [10] Alberto Cavicchioli and Fulvia Spaggiari. A note on irreducible Heegaard diagrams. *Int. J. Math. Math. Sci.*, pages Art. ID 53135, 11, 2006. [2]
- [11] Young-Eun Choi and Kasra Rafi. Comparison between Teichmüller and Lipschitz metrics. *J. Lond. Math. Soc. (2)*, 76(3):739–756, 2007. math.GT/0510136. [11]
- [12] M. Coornaert, T. Delzant, and A. Papadopoulos. *Géométrie et théorie des groupes*. Springer-Verlag, Berlin, 1990. Les groupes hyperboliques de Gromov. [7]
- [13] Robert H. Gilman. The geometry of cycles in the Cayley diagram of a group. In *The mathematical legacy of Wilhelm Magnus: groups, geometry and special functions (Brooklyn, NY, 1992)*, volume 169 of *Contemp. Math.*, pages 331–340. Amer. Math. Soc., Providence, RI, 1994. [9]
- [14] Robert H. Gilman. On the definition of word hyperbolic groups. *Math. Z.*, 242(3):529–541, 2002. [9]
- [15] Mikhael Gromov. Hyperbolic groups. In *Essays in group theory*, pages 75–263. Springer, New York, 1987. [7]
- [16] Ursula Hamenstädt. Geometry of the mapping class groups. I. Boundary amenability. *Invent. Math.*, 175(3):545–609, 2009. [60]
- [17] Kevin Hartshorn. Heegaard splittings of Haken manifolds have bounded distance. *Pacific J. Math.*, 204(1):61–75, 2002. <http://nyjm.albany.edu:8000/PacJ/2002/204-1-5nf.htm>. [27]

- [18] Willam J. Harvey. Boundary structure of the modular group. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, pages 245–251, Princeton, N.J., 1981. Princeton Univ. Press. [5]
- [19] A. Hatcher and W. Thurston. A presentation for the mapping class group of a closed orientable surface. *Topology*, 19(3):221–237, 1980. [13]
- [20] John Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001. arXiv:math/9712220. [2, 5, 67]
- [21] Tsuyoshi Kobayashi. Heights of simple loops and pseudo-Anosov homeomorphisms. In *Braids (Santa Cruz, CA, 1986)*, pages 327–338. Amer. Math. Soc., Providence, RI, 1988. [6, 12, 27]
- [22] Jason Leasure. Geodesics in the complex of curves of a surface. Ph.D. thesis. [68]
- [23] Johanna Mangahas. Uniform uniform exponential growth of subgroups of the mapping class group. *Geom. Funct. Anal.*, 19(5):1468–1480, 2010. arXiv:0805.0133. [64]
- [24] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999. arXiv:math/9804098. [3, 6, 8]
- [25] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000. arXiv:math/9807150. [3, 4, 5, 7, 11, 12, 13, 15, 38, 51, 62, 64]
- [26] Howard A. Masur and Yair N. Minsky. Quasiconvexity in the curve complex. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemp. Math.*, pages 309–320. Amer. Math. Soc., Providence, RI, 2004. arXiv:math/0307083. [3, 4, 12, 60, 61]
- [27] Howard A. Masur, Lee Mosher, and Saul Schleimer. On train track splitting sequences. arXiv:1004.4564. [3, 4, 58, 59, 60, 61]
- [28] Darryl McCullough. Virtually geometrically finite mapping class groups of 3-manifolds. *J. Differential Geom.*, 33(1):1–65, 1991. [7]
- [29] Yair N. Minsky. The classification of Kleinian surface groups, I: Models and bounds. arXiv:math/0302208. [11]
- [30] Lee Mosher. Train track expansions of measured foliations. 2003. <http://newark.rutgers.edu/~mosher/>. [58]
- [31] R. C. Penner and J. L. Harer. *Combinatorics of train tracks*, volume 125 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1992. [58, 68]
- [32] Robert C. Penner. A construction of pseudo-Anosov homeomorphisms. *Trans. Amer. Math. Soc.*, 310(1):179–197, 1988. [6]
- [33] Kasra Rafi. Relative hyperbolicity in teichmüller space, 2010. In preparation. [3, 4, 51, 52, 70]
- [34] Kasra Rafi and Saul Schleimer. Covers and the curve complex. *Geom. Topol.*, 13(4):2141–2162, 2009. arXiv:math/0701719. [38, 42]
- [35] Martin Scharlemann. The complex of curves on nonorientable surfaces. *J. London Math. Soc. (2)*, 25(1):171–184, 1982. [6]
- [36] Saul Schleimer. Notes on the curve complex. <http://math.rutgers.edu/~saulsch/Maths/notes.pdf>. [5]

- [37] Kenneth J. Shackleton. Tightness and computing distances in the curve complex. arXiv:math/0412078. [68]
- [38] William P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 19(2):417–431, 1988. [6, 21, 54]

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, CHICAGO, ILLINOIS
60607

E-mail address: masur@math.uic.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF WARWICK, COVENTRY, CV4
7AL, UK

E-mail address: s.schleimer@warwick.ac.uk