THE GEOMETRY OF THE DISK COMPLEX

HOWARD MASUR AND SAUL SCHLEIMER

ABSTRACT. *This is a work-in-progress and should not be distributed.* We prove that the disk complex is Gromov hyperbolic. Remarkably, a necessary step in the proof is a similar analysis of the geometry of the arc complex.

As an application, we find an algorithm which computes the Hempel distance of a Heegaard splitting, up to an error bounded by a function of the genus.

Contents

1. Introduction	2
1.1. The curve complex	3
1.2. The disk complex	3
1.3. Heegaard distance	5
Acknowledgements	5
2. Background on coarse geometry	5
2.1. Quasi-isometry	5
2.2. Geodesics	6
2.3. Hyperbolicity	6
2.4. A hyperbolicity criterion	7
3. Background on subsurface projection	7
3.1. More complexes	7
3.2. Surfaces and subsurfaces	8
3.3. Natural maps	10
4. Holes in general and the lower bound on distance	11
4.1. Holes, defined	12
4.2. Projection to holes is Lipschitz	13
4.3. Infinite diameter holes	13
4.4. A lower bound on distance	14
5. Holes for the nonorientable surface	14
6. Holes for the arc complex	15
7. Background on three-manifolds	16

Date: January 23, 2007.

This work is in the public domain.

HOWARD MASUR AND SAUL SCHLEIMER

7.1. Compressions	16
7.2. Band sums	18
7.3. Handlebodies and I-bundles	18
8. Holes for the disk complex	19
9. Holes for the disk complex – annuli	20
10. Holes for the disk complex – compressible	23
10.1. Compression sequences of essential disks	23
10.2. Proving the theorem	26
11. Holes for the disk complex – incompressible	26
11.1. Diagonals of polygons	26
11.2. Improving disks	27
11.3. Building the I-bundle	30
11.4. A vertical annulus parallel into the boundary	31
11.5. Finding a pseudo-Anosov map	32
11.6. Corollaries	32
12. Paths in general and the upper bound on distance	34
12.1. The requirements	34
12.2. Basic, inductive, and shortcut intervals	35
12.3. The partition	36
12.4. The upper bound	38
13. Background on Teichmüller space	41
14. Paths for the nonorientable surface	43
15. Paths for the arc complex	44
16. Background on train tracks	44
17. Paths for the disk complex	47
18. The distance estimate	47
19. Hyperbolicity	47
19.1. Index in a hole	48
19.2. Back and sidetracking	49
19.3. Projection control	49
19.4. Finding the midpoint of a side	50
19.5. Another side of the triangle	52
20. An algorithm to coarsely compute Hempel distance	53
Appendix A. Teichmüller geodesics do not backtrack	55
References	59

1. INTRODUCTION

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

 $\mathbf{2}$

we omit the subscript c and write $S_{g,b}$. If the surface is closed and orientable we simply write S_g .

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 .

1.1. The curve complex.

Definition 1.1 (Harvey [Har81]). The curve complex C(S) is the simplical complex with vertex set being isotopy classes of essential, nonperipheral curves in S. There is a k-simplex for every collection of k + 1 distinct isotopy classes having pairwise disjoint representatives.

Definition 1.2. Fix α and β , vertices of $\mathcal{C}(S)$. The distance $d_S(\alpha, \beta)$ is 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 .

For example, if $d_S(\alpha, \beta) \geq 3$ then α and β fill S: every essential non-peripheral curve in S intersects one of α or β . An important tool for this paper is the following theorem of the first author and Yair Minsky [MM99]:

Theorem 1.3. The curve complex of an orientable surface is Gromov hyperbolic. \Box

We will also need this result for nonorientable surfaces; as an application of our techniques we give a proof in Corollary 5.4.

1.2. The disk complex. 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.

Definition 1.4 (McCullough [McC91]). The disk complex $\mathcal{D}(V)$ is the simplicial complex with vertex set being proper isotopy classes of essential disks in V. We have a k-simplex for every collection of k + 1 distinct isotopy classes having pairwise disjoint representatives.

As with the curve complex, define $d_{\mathcal{D}}(D, E)$ to be the distance in the one-skeleton of $\mathcal{D}(V)$ between the disks D and E. Our driving interest is in the intrinsic metric on $\mathcal{D}(V)$ and how it differs from that of $\mathcal{C}(\partial V)$. In fact, we will necessarily treat much more general complexes. In this introduction we restrict to statements about the disk complex. An early motivation for our investigation was:

Theorem 8.3. The natural simplicial inclusion $\nu : \mathcal{D}(V) \to \mathcal{C}(\partial V)$ is not a quasi-isometric embedding.

We develop a theory of the metric obstructions, or *holes*, arising in Theorem 8.3. The complete classification of holes for $\mathcal{D}(V)$ is given by Theorems 9.1, 10.7, and 11.1. These theorems require an delicate analysis of JSJ decompositions of V caused by widely separated disks in $\mathcal{D}(V)$.

Next, inspired by the distance estimates for the mapping class group and the pants complex given in [MM00], we obtain an similar estimate for the disk complex:

Theorem 18.1. Distance in the disk complex $\mathcal{D}(V)$ between two disks D and E is quasi-equal to the sum of distances between D and E, as projected to holes $X \subset \partial V$ for $\mathcal{D}(V)$. To be precise, for any handlebody V there is a constant $c_0 = c_0(V)$ so that, for any $c \ge c_0$ there are constants $a \ge 1$ and $b \ge 0$ so that

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

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.

The paper [MM00] verifies its distance estimates by constructing and then using the *hierarchy* machinary. However, hierarchies do not appear to be flexible enough to deal with the disk complex. Instead we turn to *surgery sequences* of essential disks, as developed in [MM], and to Lee Mosher's version of train track splitting sequences (Section 16). In fact, surgery and splitting sequences are too flexible; an inductive straightening procedure is developed is Section 12.

With the distance estimate in hand we prove:

Theorem 19.3. The intrinsic metric $d_{\mathcal{D}}$ on $\mathcal{D}(V)$ is Gromov hyperbolic.

The closest precursor to Theorem 19.3 is the paper by Brock and Farb [BF] proving the hyperbolicity of the pants complex of the five-holed sphere and twice-holed torus. Another proof may be found in Behrstock's thesis [Beh04]. The former uses an area criterion (see Bowditch [Bow91]) while the latter, following [MM00], proves that certain projection maps are coarsely contracting. We instead use a version of quasi-geodesic stability, see Theorem 2.8.

We end this overview with a remark: to understand the geometry of the disk complex, one is unavoidably forced to consider the corresponding questions in the arc complex of a bounded surface as well as the curve complex of a nonorientable surface. Thus our preliminary material is unavoidably somewhat general.

1.3. Heegaard distance. We now turn to an application. Fix a *Heegaard splitting*: a triple (S, V, W) consisting of a surface and two handlebodies where $V \cap W = \partial V = \partial W = S$. Hempel [Hem01] defines the quantity $d_S(V, W) = d_S(\mathcal{D}(V), \mathcal{D}(W))$ to be the minimal distance in $\mathcal{C}(S)$ between a essential disk of V and of W. Note that a splitting can be completely determined by giving a pair of pants decompositions (maximal simplices), say \mathbb{D} in $\mathcal{D}(V)$ and \mathbb{E} in $\mathcal{D}(W)$. The resulting triple $(S, \mathbb{D}, \mathbb{E})$ is called a *Heegaard diagram*.

We prove:

Theorem 20.8. There is a constant $R_2 = R_2(S)$ and an algorithm which, given a Heegaard diagram $(S, \mathbb{D}, \mathbb{E})$, computes a number N so that

$$|d_S(V,W) - N| \le R_2.$$

Here is the context for Theorem 20.8. A long outstanding problem is to find an algorithm which, given a Heegaard diagram, decides whether or not the underlying splitting surface is *reducible*: has distance zero. Several early ascents of the Poincaré Conjecture fell at essentially this point.

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

2. Background on coarse geometry

Here we review a few ideas from coarse geometry. See [BH99], [CDP90], or even [Gro87] for a fuller discussion.

2.1. Quasi-isometry. As a bit of notation, suppose r, s, a, b are nonnegative real numbers, with $a \ge 1$. If $s \le a \cdot r + b$ then we write $s \le_{a,b} r$. If $s \le_{a,b} r$ and $r \le_{a,b} s$ then we write $s \stackrel{a,b}{=} r$ and call r and s quasi-equal with constants (a, b). We also define the *cut-off function* $[r]_c$ where $[r]_c = 0$ if r < c and $[r]_c = r$ if $r \ge c$.

Suppose that $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ are metric spaces. A map $f : \mathcal{X} \to \mathcal{Y}$ is an (a, b) quasi-isometric embedding for $a \ge 1, b \ge 0$ if, for every $x, y \in \mathcal{X}$,

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

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

We also allow set-valued maps: $f: \mathcal{X} \to \mathcal{P}(\mathcal{Y})$ where $\mathcal{P}(\mathcal{Y})$ is the power set of \mathcal{Y} . If $Z \subset \mathcal{X}$ we adopt the convention $f(Z) = \bigcup_{z \in Z} f(z)$.

The map f is an (a, b) quasi-isometric embedding if $d_{\mathcal{X}}(x, y) \stackrel{a,b}{=} d_{\mathcal{Y}}(x', y')$ for all $x' \in f(x)$ and all $y' \in f(y)$. It necessarily follows that the diameter of f(x) is at most b for any $x \in \mathcal{X}$. Again, f is a quasiisometry if $f(\mathcal{X})$ is b-dense in \mathcal{Y} .

2.2. Geodesics. Fix an interval $[u, v] \subset \mathbb{R}$. A geodesic, connecting x to y in \mathcal{X} , is a (1,0) quasi-isometric embedding $f: [u, v] \to \mathcal{X}$ 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, b) quasi-isometric embedding $g: [m, n] \to \mathcal{X}$ is called an (a, b) quasi-geodesic in \mathcal{X} . A function $g: [m, n] \to \mathcal{X}$ is an (a, b, c) unparametrized quasi-geodesic in \mathcal{X} if

- there is an increasing function $h: [p,q] \to [m,n]$ so that $g \circ h: [p,q] \to \mathcal{X}$ is an (a,b) quasi-geodesic in \mathcal{X} and
- for all $i \in [p, q-1]$, diam_{\mathcal{X}} $(g[h(i), h(i+1)]) \leq c$.

(Compare to the definition of (K, δ, s) quasi-geodesics found in [MM99].)

2.3. Hyperbolicity. Suppose that \mathcal{X} is a geodesic metric space. For convenience we now assume that \mathcal{X} is a graph with metric induced by giving all edges length one.

Definition 2.1. 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 $h \cup k$ contains g.

For the remainder of this section we assume that \mathcal{X} is hyperbolic, $x, y, z \in \mathcal{X}$ are points, and k, g, h are geodesics as in Definition 2.1.

Definition 2.2. We take $\rho_k(z)$ to be the *closest points projection* of z to k:

$$\rho_k(z) = \{ w \in k \mid d_{\mathcal{X}}(z, w) = d_{\mathcal{X}}(z, k) \}.$$

We now list several lemmas useful in the sequel.

Lemma 2.3. The closest points projection $\rho_k(z)$ has diameter at most 4δ .

Lemma 2.4. There is a point on g within distance 2δ of $\rho_k(z)$. The same holds for h.

Lemma 2.5. The diameter of $\rho_g(x) \cup \rho_h(y) \cup \rho_k(z)$ is at most 6 δ . **Lemma 2.6.** Suppose that z' is another point in \mathcal{X} so that $d_{\mathcal{X}}(z, z') \leq R$. Then $d_{\mathcal{X}}(\rho_k(z), \rho_k(z')) \leq 2R + 4\delta$.

Lemma 2.7. 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 $d_X(\rho_k(z), \rho_{k'}(z)) \leq R + 12\delta$.

2.4. A hyperbolicity criterion. Here we give a criterion for a metric space to be Gromov hyperbolic. Again, [BH99] discusses in depth related issues and definitions.

Theorem 2.8. The graph \mathcal{X} is Gromov hyperbolic if and only if there are constants $M \ge 0, \delta \ge 0$ and, for all $x, y \in \mathcal{X}^0$, a collection G(x, y) of edge paths from x to y with the following properties:

- For all $x, y \in \mathcal{X}^0$ and for all $k \in G(x, y)$ if $d_X(x, y) \leq 1$ then the path k has length at most M.
- For all $x, y, z \in \mathcal{X}^0$ and for all $k \in G(x, y), g \in G(y, z), h \in G(z, x)$ the triangle ghk is δ -slim.

Proof. It is enough to prove *stability* of geodesics: There is a constant $R = R(M, \delta)$ so that any geodesic c in \mathcal{X} , connecting x to y, and any path $k \in G(x, y)$ satisfy $d_{\text{Haus}}(c, k) \leq R$. It then follows that geodesic triangles are $R + \delta$ slim.

To prove stability, simply follow the plan of Proposition III.H.1.6 and Theorem III.H.1.7 in [BH99] replacing geodesic triangles by triangles of paths from G, as necessary. The details are left as an exercise for the interested reader.

3. BACKGROUND ON SUBSURFACE PROJECTION

This section recalls some of the theory of *subsurface projections* following [MM00].

3.1. More complexes. Suppose that S is a compact connected surface with boundary. A properly embedded arc $\beta \subset S$ is *essential* if β is not properly isotopic into ∂S .

Definition 3.1. The arc complex $\mathcal{A}(S)$ is the simplical complex with vertex set being proper isotopy classes of essential arcs in S. We have a k-simplex for every collection of k + 1 distinct classes which have pairwise disjoint representatives.

In identical fashion define $\mathcal{AC}(S)$, the arc and curve complex of S, which contains all essential arcs and all essential non-peripheral curves.

Given essential curves or arcs α , β in a surface X the geometric intersection number, $\iota(\alpha, \beta)$, is the minimum intersection possible between α and any β' properly isotopic to β . 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 (properly isotope) β until they do.

Note that the curve complexes $\mathcal{C}(S_{1,1})$ and $\mathcal{C}(S_{0,4})$, with the current definition, are nonempty but have no edges. It is useful to alter the definition in these cases as follows:

Definition 3.2. Suppose that X is one of $S_{1,1}$ or $S_{0,4}$. Add edges to $\mathcal{C}(X)$ between all vertices with geometric intersection exactly one if X $S_{1,1}$ or two if $X = S_{0,4}$. In both cases the result is the *Farey graph*.

Notice also that, with current definitions, the curve complex of an annulus $S_{0,2}$ is empty while the arc complex is finite. Following [MM00]:

Definition 3.3. Fix $X \cong S_{0,2}$. We redefine $\mathcal{C}(X)$ to be the complex with vertices being proper isotopy classes of essential arcs where all isotopies fix the boundary of X pointwise. Also, place an edge between two distinct classes if they have representatives with disjoint interiors.

As above, given α and β vertices in $\mathcal{C}(S)$ define $d_S(\alpha, \beta)$ to be the minimal number of edges required for any path from α to β in the one-skeleton of $\mathcal{C}(S)$. Recall that the geometric intersection of a pair of curves gives an upper bound for their distance, as follows:

Lemma 3.4. Suppose that S is a compact connected surface which is not an annulus. If $\alpha, \beta \subset S$ are essential and non-peripheral curves with $\iota(\alpha, \beta) > 0$ then $d_S(\alpha, \beta) \leq 2\log_2(\iota(\alpha, \beta)) + 2$.

This form of the inequality, stated for closed orientable surfaces, may be found in [Hem01]. A careful proof in the bounded orientable case is given in [Sch]. The nonorientable case is then an exercise. In the case of an annulus, no bound is required: If $X \cong S_{0,2}$ then an easy induction proves that

(3.5)
$$d_X(\alpha,\beta) = 1 + \iota(\alpha,\beta)$$

(see [MM00, Equation 2.3]) for distinct vertices $\alpha, \beta \in \mathcal{C}(X)$.

3.2. Surfaces and subsurfaces. If $X \subset S$ is a connected compact subsurface we call X essential exactly when all boundary components of X are essential in Z.

Definition 3.6. An essential subsurface $X \subset S$ is *cleanly embedded* when a component $\delta \subset \partial X$ is isotopic into ∂S if and only if δ is equal to a component of ∂S .

As a particular case: if $X \subset S$ is a cleanly embedded annulus and S is not an annulus then ∂X is not parallel into ∂S .

Definition 3.7. 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 reversely then we say that X and Y overlap.

Definition 3.8. A compact connected surface S is *simple* if $\mathcal{AC}(S)$ has finite diameter.

Here we address the question of exactly which surfaces are simple.

PSfr**Lempla**: **3**: **9**: **n5**: *uppose 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 a filling lamination.
- S admits a pseudo-Anosov map or S is an annulus.

Lemma 4.6 of [MM99] shows that pseudo-Anosov maps have quasigeodesic orbits, when acting on the associated curve complex. A Dehn twist acting on $\mathcal{C}(S_{0,2})$ has geodesic orbits.

Note that Lemma 3.9 is only used in this paper when ∂S is nonempty. The closed case is included for completeness. We find Figure 1 to be a highly convenient reference.



FIGURE 1. The notation gbc stands for $S_{g,b,c}$. Surfaces in boxes are simple. Euler characteristic is constant along lines of slope 1/2.

Proof of Lemma 3.9. We only sketch the proof. If S admits a pseudo-Anosov map then the stable lamination is filling. If S admits a filling lamination then, by an argument of Kobayashi [Kob88], $\mathcal{AC}(S)$ has infinite diameter. (This argument is also sketched in [MM99], page 124, after the statement of Proposition 4.6.)

Clearly, if the diameter of \mathcal{AC} is infinite then the diameter is at least equal to five. To finish, one may check directly that all the boxed surfaces in Figure 1 have $\mathcal{AC}(S)$ with diameter at most four. (The difficult cases, 012 and 3P, are discussed by Scharlemann [Sch82].) All surfaces not in boxes, other than \mathbb{A} , admit pseudo-Anosov maps. The orientable cases follow from Thurston's construction [Thu88]. Penner's generalization [Pen88] covers the nonorientable cases. $\hfill \Box$

3.3. Natural maps. We now define several maps between our complexes. Remember that our main concern is with coarse geometric properties. Thus our definitions are allowed to take certain liberties; in particular we are not interested in the exact location of image values but only in their location up to bounded error. To avoid making choices we often use set-valued maps.

Begin by fixing attention on a non-annular, non-simple surface S. Choose a hyperbolic metric on S so that all components of ∂S are totally geodesic. Fix also a cleanly embedded $X \subset S$. Let \widehat{S} be the cover of S where X lifts homeomorphically and where $\widehat{S} \cong \operatorname{interior}(X)$. Compactify \widehat{S} via points in its Gromov boundary. In particular, any simple closed curve in S now lifts to a simple closed curve or to a collection of properly embedded arcs. Identify $\mathcal{AC}(X)$ with $\mathcal{AC}(\widehat{S})$.

Definition 3.10. Fix $\alpha \in \mathcal{AC}(S)$. We define the *cutting map* $\kappa_X \colon \mathcal{AC}(S) \to \mathcal{P}(\mathcal{AC}(X))$ as follows: the set $\kappa_X(\alpha)$ contains all isotopy classes of lifts of α to \widehat{S} which are essential and non-peripheral.

Note that κ_S is the identity map. If α can be isotoped out of X then $\kappa_X(\alpha) = \emptyset$.

Definition 3.11. Suppose that X is not an annulus. Fix $\alpha \in \mathcal{AC}(S)$. We define the surgery map $\sigma_X \colon \mathcal{AC}(X) \to \mathcal{P}(\mathcal{C}(X))$ as follows: Let U be a closed regular neighborhood of $\alpha \cup \partial X$. Let $\sigma_X(\alpha)$ be the set of isotopy classes of components of ∂U which are essential and non-peripheral in X.

Recall that $f(A) = \bigcup_{a \in A} f(a)$.

Definition 3.12. If X is not an annulus define the subsurface projection $\pi_X \colon \mathcal{AC}(S) \to \mathcal{P}(\mathcal{C}(X))$ to be the composition $\pi_X = \sigma_X \circ \kappa_X$. If X is an annulus then set $\pi_X = \kappa_X$

Fix $\alpha \in \mathcal{AC}(S)$. If $\pi_X(\alpha) = \emptyset$ then α misses the surface X. If $\pi_X(\alpha)$ is nonempty then α cuts the surface X. If $\alpha, \beta \in \mathcal{AC}(S)$ both cut X we write $d_X(\alpha, \beta) = \operatorname{diam}_X(\pi_X(\alpha) \cup \pi_X(\beta))$. This is the projection distance between α and β in X. As a simple observation we have:

Lemma 3.13. Suppose $\alpha, \beta \in \mathcal{AC}(S)$ are disjoint and cut X. Then $\operatorname{diam}_X(\pi_X(\alpha))$ and $d_X(\alpha, \beta)$ are at most two.

This follows because for any $\delta, \epsilon \in \pi_X(\alpha) \cup \pi_X(\beta)$ we find $\iota(\delta, \epsilon) \leq 4$. See Lemma 2.3 of [MM00] for a very similar statement. As a consequence we have: **Corollary 3.14.** 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 $d_X(\beta_0, \beta_n) \leq 2n$.

If some vertex of $\{\beta_i\}$ does not cut X then the conclusion need not hold: the projection distance $d_X(\beta_0, \beta_n)$ may be arbitrarily large compared to n.

A kind of converse to Lemma 3.13 is:

Lemma 3.15. 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 [CR05] gives a more precise relation between projection distances and intersection number.

Proof of Lemma 3.15. 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 $\lambda \in \mathcal{PML}(S)$. Let Y be any connected component of the subsurface filled by λ , choosen so that α_0 cuts Y. Note that $\pi_Y(\beta_n)$ converges to $\lambda|Y$. Again applying Kobayashi's argument [Kob88], the distance $d_Y(\alpha_0, \beta_n)$ tends to infinity.

4. Holes in general and the lower bound on distance

Suppose that S is a compact connected surface. A *multi-curve* is a collection of disjoint non-parallel essential non-peripheral curves or arcs in S.

In this paper a combinatorial complex $\mathcal{G}(S)$ will always have isotopy classes of certain multi-curves is S as vertices. Vertices will be connected by edges only if there are representatives which are disjoint. We also always assume that \mathcal{G} is connected. There is a natural map $\nu: \mathcal{G} \to \mathcal{AC}(S)$ taking a vertex of \mathcal{G} to the isotopy classes of the components.

One may also entertain complexes where edges are placed between multi-curves with bounded intersection, perhaps of a specified kind. Examples in the literature include the pants complex [Bro03] [BDM], the Hatcher-Thurston complex [HT80], and the complex of separating curves [BM04].

It almost always will suffice to study subcomplexes $\mathcal{G} \subset \mathcal{AC}(S)$. However, the more general case arises when dealing with the curve complex or the arc complex of a nonorientable surface. This allows us to avoid dealing with the Teichmüller space of the nonorientable surface. It also illustrates the generality of our techniques.

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 $d_{S}(\cdot, \cdot)$.

4.1. Holes, defined. We now develop a technique to measure the failure of the natural map $\nu \colon \mathcal{G} \to \mathcal{AC}(S)$ to be a quasi-isometric embedding. Suppose that S is non-simple. Suppose that $\mathcal{G}(S)$ is connected combinatorial complex of multi-curves in S. Suppose that $X \subset S$ is an cleanly embedded subsurface. A vertex $\alpha \in \mathcal{G}$ cuts X if some component of α cuts X.

Definition 4.1. 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 $\pi_X \colon \mathcal{G}(S) \to \mathcal{C}(X)$ 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} . Boundary parallel annuli cannot be cleanly embedded (unless S is also an annulus) and so cannot be holes. A hole $X \subset S$ is a *strict hole* if X is not homeomorphic to S.

Example 4.2. 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 all of ∂S . So, for example, if S is planar then only S is a hole for $\mathcal{A}(S)$. The same holds if $S = S_{1,1}$. In all other cases the arc complex admits infinitely many holes.

Definition 4.3. 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 4.4. 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 $X, Y \subset S$ are disjoint holes for \mathcal{G} . In the presence of symmetry there can be a relationship between $\pi_X | \mathcal{G}$ and $\pi_Y | \mathcal{G}$ as follows:

Definition 4.5. Suppose that X, Y are holes for \mathcal{G} of infinite diameter. Then X and Y are *paired* if there is a homeomorphism $\tau: X \to Y$ and a constant K so that

$$d_Y(\pi_Y(\gamma), \tau(\pi_X(\gamma))) \le K$$

for every $\gamma \in \mathcal{G}$.

4.2. **Projection to holes is Lipschitz.** The following lemma is used repeatedly throughout the paper:

Lemma 4.6. 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

$$d_X(\alpha,\beta) \le 2 + 2 \cdot d_{\mathcal{G}}(\alpha,\beta).$$

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

Proof. This follows directly from Corollary 3.14 and our assumption that vertices of \mathcal{G} connected by an edge represent disjoint multi-curves.

We deduce the well-known:

Lemma 4.7. The inclusion $\nu \colon \mathcal{C}(S) \to \mathcal{AC}(S)$ is a quasi-isometry. The surgery map $\sigma_S \colon \mathcal{AC}(S) \to \mathcal{P}(\mathcal{C}(S))$ is a quasi-inverse for ν .

Proof. Fix a pair of essential, non-peripheral curves α and β inside of S. Since ν is an inclusion we have $d_{\mathcal{AC}}(\alpha,\beta) \leq d_S(\alpha,\beta)$. As S is a hole for $\mathcal{AC}(S)$ by Lemma 4.6 we have $d_S(\alpha,\beta) \leq 2 + 2 \cdot d_{\mathcal{AC}}(\alpha,\beta)$.

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 $\sigma_S(\gamma)$.

4.3. Infinite diameter holes. Brian Bowditch posed several questions (Newton Institute, August 2003) regarding the geometry of $\mathcal{A}(S)$: in particular, is the inclusion $\nu : \mathcal{A}(S) \to \mathcal{AC}(S)$ a quasi-isometry? The absence of a natural quasi-inverse is one hint; the presence of holes for \mathcal{A} is another. In fact, we have:

Lemma 4.8. Suppose that $\mathcal{G}(S)$ is a combinatorial complex, and $X \subset S$ is a strict hole of infinite diameter. Then $\nu : \mathcal{G} \to \mathcal{AC}(S)$ is not a quasi-isometric embedding.

This lemma and Example 4.2 completely determines when the inclusion of $\mathcal{A}(S)$ into $\mathcal{AC}(S)$ is a quasi-isometric embedding and so answers Bowditch's initial question. It quickly becomes clear that the set of holes tightly constrains the intrinsic geometry of a combinatorial complex.

Lemma 4.9. Suppose that $\mathcal{G}(S)$ is a combinatorial complex invariant under the natural action of $\mathcal{MCG}(S)$. Then every non-simple hole for \mathcal{G} has infinite diameter. Furthermore, if $X, Y \subset S$ are disjoint nonsimple holes for \mathcal{G} then there is a quasi-isometric embedding of \mathbb{Z}^2 into \mathcal{G} . \Box We will not use Lemmas 4.8 or 4.9 in an essential way and accordingly omit proofs. Instead our interest lies in proving the far more powerful distance estimate for $\mathcal{G}(S)$.

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

Theorem 4.10. 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 are constants $a \geq 1, b \geq 0$ satisfying

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

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

The proof follows the proof of Theorems 6.10 and 6.12 of [MM00], practially 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 [MM00], which records how markings differing by an elementary move project to an essential subsurface, by Lemma 4.6 of this paper, which records how \mathcal{G} projects to a hole.

One major goal of this paper is to give criteria sufficient obtain the reverse inequality; Theorem 12.1.

5. Holes for the nonorientable surface

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

Definition 5.1. A multi-curve or multi-arc $\gamma \subset S$ is symmetric if $\tau(\gamma) \cap \gamma = \emptyset$ or $\tau(\gamma) = \gamma$. A multi-curve or multi-arc 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 5.2. The *invariant complex* $C^{\tau}(S)$ is the simplical complex with vertex set being isotopy classes of invariant multi-curves. 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 $\nu : \mathcal{C}^{\tau}(S) \to \mathcal{C}(S)$. We will prove:

Lemma 5.3. $\nu: \mathcal{C}^{\tau}(S) \to \mathcal{C}(S)$ is a quasi-isometric embedding.

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

Corollary 5.4. C(F) is Gromov hyperbolic.

This brief discussion is not really satisfying. Perhaps we should sketch the proof... One half of the proof of Lemma 5.3 is straight-forward: since ν sends adjacent vertices to adjacent edges we have

(5.5)
$$d_S(\alpha,\beta) \le d_{\mathcal{C}^\tau}(\alpha,\beta)$$

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 4.6 or indeed from Theorem 4.10.

The other half of the proof 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 14.1. There is a constant K with the following property: Suppose that α, β are invariant multi-curves or arcs in S and $X \subset S$ so that $d_X(\alpha, \beta) > K$. Then X is symmetric.

From this it follows that:

Corollary 5.6. With K as in Lemma 14.1: If $X \subset S$ is a hole for $C^{\tau}(S)$ with diameter greater than K then X = S.

Proof. Suppose that $X \subset S$ is a strict subsurface, cleanly embedded. Suppose that $\operatorname{diam}_X(\mathcal{C}^{\tau}(S)) > K$. 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 $\mathcal{C}^{\tau}(S)$.

This, together with the upper bound (Theorem 12.1), proves Lemma 5.3.

6. Holes for the ARC Complex

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

Definition 6.1. Suppose that S is a non-simple surface with boundary. Let Δ be a nonempty collection of components of ∂S . The *arc complex* $\mathcal{A}(S, \Delta)$ is the full subcomplex of $\mathcal{A}(S)$ containing all essential arcs $\alpha \subset S$ with $\partial \alpha \subset \Delta$.

Note that $\mathcal{A}(S, \partial S)$ and $\mathcal{A}(S)$ are identical. If X is non-simple, cleanly embedded in S, and $\Delta \subset X$ then the meaning of $\mathcal{A}(X, \Delta)$ is clear.

Lemma 6.2. Suppose $X \subset S$ is cleanly embedded. If $\Delta \subset \partial X$ then X is a hole for $\mathcal{A}(S, \Delta)$.

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

Lemma 6.3. If $X, Y \subset S$ are holes for $\mathcal{A}(S, \Delta)$ then $X \cap Y \neq \emptyset$. \Box

The proof follows immediately from Lemma 6.2. Lemma 4.9 indicates that Lemma 6.3 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 nonorientable case we must carefully lift to the orientation cover. We require the following definitions.

Suppose that F is non-simple, nonorientable, and has nonempty boundary. Let $\rho_F \colon S \to F$ be the orientation double cover and let $\tau \colon S \to S$ be the induced involution. Fix $\Delta' \subset \partial F$ and let $\Delta = \rho_F^{-1}(\Delta')$. **Definition 6.4.** 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 simplically isomorphic to $\mathcal{A}(F, \Delta')$. Now, it follows from Lemma 14.1 that all holes for $\mathcal{A}^{\tau}(S, \Delta)$, with sufficiently large diameter, are symmetric subsurfaces $X \subset S$ so that $\Delta \subset X \cup \tau(X)$. If $X \cap \tau(X) = \emptyset$ then the subsurfaces X and $\tau(X)$ are paired holes, as in Definition 4.5. 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, all infinite diameter holes for the invariant arc complex *interfere*: any such hole Z intersects any other such hole X or intersects a hole paired with X.

Lemma 6.5. Suppose that $X, Z \subset S$ are holes for $\mathcal{A}^{\tau}(S, \Delta)$, both of infinite diameter. Then Z intersects X or $\tau(X)$.

Proof. Notice that $\Delta \subset \partial X \cup \partial Y$. Since Δ is also contained in $Z \cup \tau(Z)$, the subsurface Z meets either X or Y.

7. Background on three-manifolds

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

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

7.1. Compressions. Suppose that F is a surface embedded in M. Then F is *compressible* if there is a disk B embedded in M with interior $(B) \cap \partial M = \emptyset$, $B \cap F = \partial B$, 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 7.1. An embedded surface F is boundary compressible if there is a disk B embedded in M with

- 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 $\partial M \smallsetminus \partial F$.

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 others 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 fr $N \setminus F$. Again, F' is obtained by boundary compressing F along B. If no such bigon exists then F is boundary incompressible.

Remark 7.2. Recall that any surface F properly embedded in a handlebody V_g , $g \ge 2$, is either compressible or boundary compressible.

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

Remark 7.3. 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, after tightening, each F_i meets 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 7.4. Suppose that M is irreducible. Fix X an 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 nonempty then an outermost bigon of E gives the desired boundary compression lying in X.

Suppose instead that $F \cap E = \emptyset$ but F does meet 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 $fr(N) \setminus F$ 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 $\partial D = \gamma$, F is incompressible, and M is irreducible deduce that F is isotopic to E.

7.2. Band sums. 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 $fr(N) \setminus F'$. The surface F obtained is the result of band summing F_1 to F_2 along δ . It is straight-forward to show 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'.

Remark 7.5. Note that, if F is a band sum of F' then every component of ∂F in ∂M is disjoint from, or parallel to, every curve of $\partial F'$.

7.3. 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 Fis 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 $\rho_F \colon T \to F$ 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 7.6. Suppose that F is a compact, connected surface with $\partial F \neq \emptyset$. Let $\rho_F \colon T \to F$ 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. \Box

8. 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 $\nu: \mathcal{D}(V) \to \mathcal{C}(S)$.

Remark 8.1. 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,
- $\overline{S \setminus X} = \overline{\partial V \setminus X}$ 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 9.1. Suppose X is a hole for $\mathcal{D}(V)$ and X is an annulus. Then the diameter of X is at most 3.

Theorem 10.7. 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.

Thus X supports a pseudo-Anosov map.

Theorem 11.1. Suppose X is an incompressible hole for $\mathcal{D}(V)$ with diameter at least $4\log_2(g) + 60$. Then there is an I-bundle $\rho_F \colon T \to F$ embedded in V so that

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

As a corollary of these theorems we have:

Corollary 8.2. If X is hole for $\mathcal{D}(V)$ with diameter at least $4 \log_2(g) + 60$ then X has infinite diameter.

Proof. We only give a sketch. If X is a hole with diameter at least $4 \log_2(g) + 60$ then either Theorem 10.7 or 11.1 applies.

If X is compressible then Dehn twists, in opposite directions, about the given disks D and E yields a automorphisms $f: V \to V$ so that f|Xis pseudo-Anosov. This follows from Thurston's construction [Thu88]. By Lemma 3.9 the hole X has infinite diameter.

If X is incompressible then $X \subset \partial_h T$ where $\rho_F \colon T \to F$ is the given *I*-bundle. Let $f \colon F \to F$ 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. \Box

Applying Lemma 4.8 we find another corollary:

Theorem 8.3. If ∂V contains a strict hole with diameter at least $4 \log_2(g) + 60$ then the inclusion $\nu \colon \mathcal{D}(V) \to \mathcal{C}(\partial V)$ is not a quasiisometric embedding.

9. Holes for the disk complex – annuli

The proof of Theorem 9.1 occupies the rest of this section. This proof shares many features with the proofs of Theorems 10.7 and 11.1. However, the exceptional definition of $\mathcal{C}(S_{0,2})$ prevents a unified approach. Fix V, a handlebody.

Theorem 9.1. Suppose X is a hole for $\mathcal{D}(V)$ and X is an annulus. Then the diameter of X is at most 3.

Assume, to obtain a contradiction, that X has diameter at least 4. Suppose that $D \in \mathcal{D}(V)$ is a disk choosen to minimize $D \cap X$. Among all disks $E \in \mathcal{D}(V)$ with $d_X(D, E) \geq 2$ 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 2.



FIGURE 2. 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. This follows from two observations. First, after the tightening we find that the arcs $D \cap X$ and $E \cap X$ lie inside of X just as $\pi_X(D)$ and $\pi_X(E)$ lie inside of \hat{X} , the annular cover of ∂V . Second, by Equation 3.5, $d_X(D \cap X, E \cap X) = 1 + \iota(D \cap X, E \cap X)$.

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 3 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 $B \subset E \setminus D$: a subdisk where $B \cap D$ is a single arc lying in the boundary of B. 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 Bis outermost or the fact that every arc of $E \cap X$ meets some arc of $D \cap X$. Using this observation, Figure 3 lists the possible ways for Bto lie inside of E.



FIGURE 3. The arc α cuts 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 3. 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 3.

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 $|D \cap X| = 2$ and the claim holds.

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 meets X at most once after tightening. If $C \cup B$ is disjoint from X then X is not a hole, a contradiction. If $C \cup B$ meets X exactly once then we may form a new disk C' by boundary compressing X along $C \cup B$. Then this disk C' is disjoint from X, a contradiction. \Box

We finish the proof of Theorem 9.1 by noting that $D \cup B$ is a *tripod*: $D \cup B$ is homeomorphic to $\Upsilon \times I$ where Υ is the simplicial graph with three edges and four vertices, three of valence one. 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. As in the previous paragraph we find that X is not a hole, a contradiction. \Box

10. Holes for the disk complex – compressible

The proof of Theorem 10.7 occupies the second half of this section.

10.1. Compression sequences of essential disks. Fix a nonempty, pairwise disjoint collection of essential curves $\Gamma \subset \partial V$. Fix also a essential disk $D \subset V$. Properly isotope D to make ∂D tight with respect to Γ .

Assume for the moment that D meets some component of Γ .

Definition 10.1. Define a compression sequence of essential disks as follows: $\{(\Delta_k, B_k)\}_{k=1}^n$ where $\Delta_1 = \{D\}$, B_k is a boundary compression of Δ_k , B_k is disjoint from Γ , and Δ_{k+1} is obtained by boundary compressing Δ_k along B_k and tightening with respect to Γ . 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 every 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 $S \smallsetminus \Gamma$ 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 D' with $\partial D' \subset Z$.

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

Note that all compression sequences must end by Remark 7.3. Given a maximal sequence we may relate the various disks in the sequence as follows:

Definition 10.2. 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.

We require the following somewhat technical lemma:

Lemma 10.3. Fix $\Gamma \subset \partial V$ a nonempty collection of pairwise disjoint essential curves. Suppose that D meets Γ and choose a maximal sequence starting at D. Fix any component $X \subset \partial V \setminus \Gamma$ so that Δ_n cuts X. Fix any disk $D_k \in \Delta_k$ which also cuts X. Then either $D_k \in \Delta_n$ or there are four distinct disjointness pairs $\{(\alpha_i, \beta_i)\}_{i=1}^4$ for D_k 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. If not, then D_k is the unique disk of Δ_k which is boundary compressed at level 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_k .

Let δ be the band sum arc dual to B_k . We may assume that $|\Gamma \cap \delta|$ is minimal over all arcs dual to B_k . 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_k 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: $|X \cap \partial \delta|$ equals zero, one, or two. We deal with each in turn.

First suppose that $X \cap \partial \delta = \emptyset$. Then $D_k \cap X$ contains $(D_{k+1} \cup D'_{k+1}) \cap X$. If D_{k+1} and D'_{k+1} are both components of Δ_n then choose any arc $\beta \subset D_{k+1} \cap X$ and any arc $\beta' \subset D'_{k+1} \cap X$. The four disjointness pairs are then all ordered combinations of β and β' . Suppose instead that D_{k+1} , say, is not a component of D_n . Then as $D_{k+1} \cap X \subset D_k \cap X$ the disk D_k inherits four disjointness pairs from D_{k+1} .

Second suppose that $X \cap \partial \delta$ is a single point of D_{k+1} , contained in $\gamma \subset D_{k+1} \cap X$. Let α, α' be the arcs of $D_k \cap X$ meeting a regular neighborhood of $\delta \cup \gamma$. Both of these are disjoint from γ in X, by Remark 7.5. 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 (β, 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 (β, β) and (α, β) , say, give the desired conclusion.

Third suppose that $X \cap \partial \delta$ meets $\gamma \subset D_{k+1}$ and $\gamma' \subset D'_{k+1}$. Let α and α' be two arcs of $D_k \cap X$ meeting the regular neighborhood of $(\delta \cup \gamma \cup \gamma') \cap X$. These are disjoint from all arcs of $(D_{k+1} \cup D'_{k+1}) \cap X$ again by Remark 7.5. Suppose both D_{k+1} or 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 and the proof is complete. \Box

Lemma 10.3 allows us to relate a disk D to another disk D' obtained from D by repeated boundary compressions in the complement of Γ . Lemma 10.4. Fix $X \subset \partial V$, a hole 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 8.1, that D cuts ∂X . Choose a maximal sequence starting at D.

Assume first that the sequence is non-trivial (n > 1). By Lemma 10.3 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. \Box

Remark 10.5. At first sight Lemma 10.4 appears surprising; after all, any pair of curves in C(X) can be connected by a sequence of band sums. However, the sequences of band sums arising in Lemma 10.4 are quite special due to the fact that D has bounded genus.

When V has the structure of an I bundle an important special case of Lemma 10.4 appears: Let $T \cong F \times I$, $X \cup Y = \partial_h T$, and $\rho_F \colon T \to F$ be the natural projection. If α, β are arcs in $\partial_h T$ define $d_{\mathcal{A}(F)}(\alpha, \beta) = d_{\mathcal{A}(F)}(\rho_F(\alpha), \rho_F(\beta))$.

Lemma 10.6. Suppose that D is an essential disk in $T \cong F \times I$. Tighten ∂D with respect to $\partial \partial_h T$. Then $d_{\mathcal{A}(F)}(D \cap X, D \cap Y) \leq 6$.

Proof. As in Lemma 10.4 obtain a disk D' which cannot be boundary compressed into X or Y. As D' is ∂_X incompressible Lemma 7.6 proves D' may be isotoped to be vertical with respect to T (that is, D' is a union of fibres). Thus $D' \cap X$ is a single arc as is $D' \cap Y$. We find that with $d_{\mathcal{A}(X)}(D, D') \leq 3$ and $d_{\mathcal{A}(Y)}(D, D') \leq 3$ and the lemma follows. \Box 10.2. **Proving the theorem.** It is now trivial to prove Theorem 10.7. We first recall the statement:

Theorem 10.7. 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.

Thus X supports a pseudo-Anosov map.

Proof. Choose disks D' and E' in $\mathcal{D}(V)$ so that $d_X(D', E') \ge 15$. By Lemma 10.4 there are disks D and E so that $\partial D, \partial E \subset X, d_X(D', D) \le 6$, and $d_X(E', E) \le 6$. It follows from the triangle inequality that $d_X(D, E) \ge 3$ and the theorem is proved. \Box

11. Holes for the disk complex – incompressible

This section classifies incompressible holes for the disk complex.

Theorem 11.1. Suppose X is an incompressible hole for $\mathcal{D}(V)$ with diameter at least $4\log_2(g) + 60$. Then there is an I-bundle $\rho_F \colon T \to F$ 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 10.4 we may assume that D, E are essential disks, without boundary compressions into X or $S \setminus X$, with $d_X(D, E) > 4 \log_2(g) + 48$. Examine the intersection pattern of Dand E to find rectangles R and Q. The intersection of these rectangle in V will determine the desired I-bundle T. The third and fourth conclusions of the theorem follow from standard facts about primitive annuli. The fifth requires another application of Lemma 10.4 as well as Lemma 3.9.

11.1. **Diagonals of polygons.** Before beginning the proof of Theorem 11.1 we must briefly discuss diagonals of polygons.

Let D be a 2n sided regular polygon. Label the sides of D with the letters X and Y in an alternating fashion. Any side labeled X (or Y) will be called an X side (or Y side).

Definition 11.2. 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 we say that γ and γ' are *non-parallel*.

Lemma 11.3. 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. First, an easy counting argument shows that $|\Gamma| \leq 4n - 3$. Now, if every X side meets at least nine non-parallel diagonals then $|\Gamma| \geq \frac{9}{2}n > 4n - 3$, a contradiction. \Box

11.2. **Improving disks.** Suppose now that X is an incompressible hole for $\mathcal{D}(V)$ with diameter at least $4 \log_2(g) + 60$. (From Theorem 9.1 it follows that X is not an annulus.) Let $Y = \overline{S \setminus X}$.

Choose disks D' and E' in V so that $d_X(D', E') \ge 4 \log_2(g) + 60$. By Lemma 10.4 there are a pair of disks D and E so that both are essential in V, cannot be boundary compressed into X or into $\overline{S \setminus X}$, and so that $d_{\mathcal{A}(X)}(D', D) \le 3$ and $d_{\mathcal{A}(X)}(E', E) \le 3$. Thus $d_X(D', D) \le 6$ and $d_X(E', E) \le 6$ (Lemma 4.6). By the triangle inequality $d_X(D, E) \ge$ $4 \log_2(g) + 60 - 12 = 4 \log_2(g) + 48$.

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. Now consider D and E to be evensided polygons, with vertices being the points $\partial D \cap \partial X$ and $\partial E \cap \partial X$ respectively. Now let $\Gamma = D \cap E$. See Figure 4 for one *a priori* possible collection $\Gamma \subset D$.



FIGURE 4. Can $\Gamma \subset D$ contain simple closed curves and non-diagonal arcs?

From our assumptions and the irreducibility of V it follows that Γ contains no simple closed curves. Suppose now that there is an $\gamma \subset \Gamma$ so that, in D, both endpoints of γ lie in the same side of D – so γ is not a diagonal for 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 $d_X(D, E) \geq 4 \log_2(g) + 48$ we have that $d_X(\alpha, \beta) \geq 4 \log_2(g) + 48 - 4 = 4 \log_2(g) + 44$. As $d_X(\alpha, \beta) = d_X(\sigma_X(\alpha), \sigma_X(\beta))$ and $\iota(\sigma_X(\alpha), \sigma_X(\beta)) \leq 4\iota(\alpha, \beta) + 4$ it follows from Lemma 3.4 that $\iota(\alpha, \beta) \geq 2^{\frac{4 \log_2(g) + 44 - 6}{2}} - 1 > 2^{13}$

Now break α and β into subarcs, each subarc meeting all of the diagonals of fixed type and only meeting the diagonals of that type. There is a pair of such subarcs $\alpha' \subset \alpha$ and $\beta' \subset \beta$ with the following property:

$$|\alpha' \cap \beta'| \ge \frac{\iota(\alpha, \beta)}{64}.$$

Note that it follows that $|\alpha' \cap \beta'| \ge 128$.

Claim 11.4. The graph $\Theta' = \alpha' \cup \beta'$ fills X.

Proof. Suppose not. Let $W \subset X$ be the surface obtained by taking a regular neighborhood of Θ' union with all disks and boundary parallel annuli (in X) of $X \setminus \Theta'$. Thus W is an essential subsurface strictly contained in X. We now assume, breaking symmetry, that $d_X(\alpha, \partial W) \geq \frac{4\log_2(g)+44}{2} = 2\log_2(g) + 22$. As α and β are tight in X it follows that both α and β are tight with

As α and β are tight in X it follows that both α and β are tight with respect to ∂W . Let $A = \{\alpha_i\}_{i=0}^n = \alpha \cap W$ be the collection of subarcs of α meeting W, with $\alpha' \subset \alpha_0$. Partition A into subsets A_k so that

- $\alpha_0 \in A_0$ and
- $\alpha_i \in A_h, \alpha_j \in A_k$ are properly isotopic if and only if h = k.

Let \mathcal{A} be the collection of partitions A_k so that $|A_k| \ge 65$.

Subclaim. Cutting the surface W along \mathcal{A} gives a collection of disks.

Suppose, for a contradiction, that $W' \subset W \setminus \mathcal{A}$ is not a disk. Deduce that if W is not an annulus there is an essential curve $\gamma \subset W'$ with $|\gamma \cap \alpha| \leq 64 \cdot (-3\chi(W))$. This is because $\alpha \cap W$ yields at most $-3\chi(W)$ different proper isotopy classes of arcs. If W is an annulus (and so W' = W) then $|\gamma \cap \alpha| \leq 64$.

In either case $|\gamma \cap \alpha| \leq 64 \cdot 6g \leq 512g$ where g is the genus of the handlebody V. By Lemma 3.4 deduce that $d_X(\gamma, \alpha) \leq 2(9 + \log_2(g)) + 2 = 2\log_2(g) + 20$. As $\gamma \subset W' \subset W$ it follows that $d_X(\gamma, \partial W) \leq 1$ and so $d_X(\alpha, \partial W) \leq 2\log_2(g) + 21$. This contradicts our assumption that $d_X(\alpha, \partial W) \geq 2\log_2(g) + 22$.

On the other hand:

Subclaim. A_0 contains at most 64 arcs.

If A_0 contains 65 arcs or more then $|\alpha \cap \beta| > 64 \cdot |\alpha' \cap \beta'| \ge \iota(\alpha, \beta)$, a contradiction.

So let W_0 be the disk component of $W \setminus \mathcal{A}$ containing $\alpha_0 = \alpha'$. Note that α_0 separates W_0 . Thus every time β' crosses α_0 , except for the last, β' must also cross some partition $A_k \in \mathcal{A}$. We deduce that:

$$\iota(\alpha,\beta) \ge |\alpha \cap \beta'| \ge 65(|\alpha' \cap \beta'| - 1) > 64 \cdot |\alpha' \cap \beta'| \ge \iota(\alpha,\beta).$$

(The strict inequality follows from the fact that $|\alpha' \cap \beta'| \ge 128$.) As this is a contradiction, Claim 11.4 is proven.

Let $R \subset D$ be the rectangle with top side equal to α' , left and right sides equal to diagonals $D \cap E$ of D, and bottom side in $\partial D \setminus \partial X$ parallel (along the diagonals) to α' . Denote this bottom side of R by α'' . See Figure 5. We define $Q \subset E$ and $\beta'' \subset \partial Q$ similarly.



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

As above we have $\Theta' = \alpha' \cup \beta'$. Let $\Theta'' = \alpha'' \cup \beta''$. There are now two possibilities: either Θ' and Θ'' are disjoint or they are not. In the first case let $\Theta = \Theta'$. In the second case let $\Theta = \Theta' \cup \Theta''$. In either case notice that, by Claim 11.4, the graph Θ is connected and fills X.

The following claim will be useful in the next section:

Claim 11.5. The graph Θ' is not contained in any disk C embedded in ∂V . The same holds for Θ'' .

Proof. This is clear for Θ' because Θ' fills X, an essential subsurface of ∂V . Now suppose that $\Theta'' \subset C$, a disk. Note that $|\alpha'' \cap \beta''| = |\alpha' \cap \beta'| > 1$. We are supposing that all intersections of α'' and β'' occur in C. It follows that ∂D and ∂E are not tight on ∂V , a contradiction. \Box

11.3. Building the I-bundle. As in the previous section we are given a pair of rectangles $R \subset D$ and $Q \subset E$ so that $R \cap \partial V = \alpha' \cup \alpha''$ and $Q \cap \partial V = \beta' \cup \beta''$. We also have a connected graph Θ which fills X. Note that $R \cup Q$ is an *I*-bundle and Θ is the component of its horizontal boundary meeting X. See Figure 6.



FIGURE 6. $R \cup Q$ is an *I*-bundle because all arcs of intersection are parallel.

Let T_0 be a regular neighborhood of $R \cup Q$, taken in V. This 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 11.4. We wish to 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 11.5). 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 $\partial V \setminus X$, 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 $\Theta \subset \partial_h T$, which fills *X*, by induction. It follows that the component of $\partial_h T$ containing Θ is isotopic to *X*.

This finishes the construction of the promised I-bundle T and demonstrates the first two conclusions of Theorem 11.1. For future use we record:

Remark 11.6. Every curve of $\partial \partial_v T = \partial \partial_h T$ is essential in $S = \partial V$.

11.4. A vertical annulus parallel into the boundary. Here we obtain the third conclusion of Theorem 11.1: 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 11.7. All components of $\partial_v T$ are incompressible in V.

Proof. Suppose that $A \subset \partial_v T$ was compressible. By Remark 11.6 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 8.1.

Claim 11.8. Some component of $\partial_v T$ is boundary parallel.

Proof. Since $\partial_v T$ is incompressible (Claim 11.7) by Remark 7.2, 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 $V \smallsetminus A$. Let A' be the closure of $A \searrow N$ (so A' is a rectangle). Let $B' \cup B''$ be the closure of $fr(N) \searrow A$. 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 $D \subset \overline{V \searrow T}$. 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 11.6. 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. \Box

Remark 11.9. 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.

11.5. 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 11.2, pick essential disks D' and E' essential disks in Vso that $d_X(D', E') \ge 4 \log_2(g) + 60$. Lemma 10.4 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, $d_X(D, E) \ge 4 \log_2(g) + 60 - 12 = 4 \log_2(g) + 48$.

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 7.6 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. Note that we still have $d_X(A, B) \ge 4 \log_2(g) + 48$.

We now begin to work in the base surface F. Recall that $\rho_F: T \to F$ is an *I*-bundle. Take $\alpha = \rho_F(A)$ and $\beta = \rho_F(B)$. Note that the natural map $\mathcal{C}(F) \to \mathcal{C}(X)$, defined by taking a curve to its lift, is distance non-increasing (see Equation 5.5). Thus $d_F(\alpha, \beta) \ge 4 \log_2(g) + 48$. By Theorem 9.1 the surface F cannot be an annulus. Thus, by Lemma 3.9 the subsurface F supports a pseudo-Anosov map and we are done.

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

Lemma 11.10. Suppose that X is incompressible hole for $\mathcal{D}(V)$ supported by the *I*-bundle $\rho_F \colon T \to F$. Suppose that $D \subset V$ is an essential disk.

- If F is orientable then set $X \cup Y = \partial_h T$. Then $d_{\mathcal{A}(F)}(D \cap X, D \cap Y) \leq 5$.
- If F is nonorientable then set $X = \partial_h T$. Let $\tau \colon X \to X$ be the associated involution. Then $d_X(D, \mathcal{C}^{\tau}(X)) \leq 6$.

Proof. By Lemma 10.4 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 $X \subset \partial_h T$ we have $d_{\mathcal{A}(X)}(D, D') \leq 3$.

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 $E = D' \cap T$. Deduce that the components of E are, after a proper isotopy in T, vertical with respect to T.

Pick any component E' of E and let $\alpha \subset F$ be the arc of intersection between E' and F. Let $\alpha' = E' \cap \partial_h T$. This is, after a proper isotopy of D', a subarc of $D' \cap \partial_h T$. Note that α' is symmetric under the involution $\tau : \partial_h T \to \partial_h T$. The conclusion follows. \Box

Recall Lemma 6.3: all holes for the arc complex intersect. This cannot hold for the disk complex. For example if $\rho_F: T \to F$ 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 11.10, X and Y are paired holes, in the sense of Definition 4.5. So, as with the invariant arc complex (Lemma 6.5), all holes for the disk complex interfere:

Lemma 11.11. Suppose that $X, Z \subset \partial V$ are holes for $\mathcal{D}(V)$, both of infinite diameter. 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 8.1 that both X and Z are incompressible. Let $\rho_F \colon T \to F$ be the *I*-bundle in V with $X \subset \partial_h T$, as provided by Theorem 11.1. We also have a component $A \subset \partial_v T$ so that A is boundary parallel. Let U be the solid torus component of $V \setminus A$.

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.

Now, if F is nonorientable then $\partial D \subset X \cup (\partial U \setminus A)$. Also, Z cannot meet U (because $Z \cap X = \emptyset$) and it follows that $D \cap Z = \emptyset$, a contradiction.

Deduce that F is orientable. Let $X \cup Y = \partial_h T$. So $\partial D \subset (X \cup (\partial U \setminus A) \cup Y)$. Now, if Z misses Y then $Z \subset \partial U \setminus A$. In this case, Z is itself an annulus, contradicting Theorem 9.1. It follows that Z cuts Y and we are done.

12. Paths in general and the upper bound on distance

Suppose that $\mathcal{G} = \mathcal{G}(S)$ is a combinatorial complex. The goal of the section is to give conditions on \mathcal{G} , the *path requirements*, which imply that the distance $d_{\mathcal{G}}$ is bounded above by the sum of projections to holes. We begin by stating the target theorem.

Theorem 12.1. Fix S a compact connected non-simple surface. Suppose that $\mathcal{G}(S)$ is a combinatorial complex satisfying the path requirements. Then there is a constant $c_0 = c_0(S)$ so that for all $c \ge c_0$ there are constants $a \ge 1, b \ge 0$ satisfying

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

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

The proof is more difficult than that of the lower bound, Theorem 4.10. Essentially, we must build a path in \mathcal{G} which is not too long. We will be given a path $\Lambda = \{\mu_n\}$ of markings (defined below) which is *locally* an unparametrized quasi-geodesic: when X is a hole for \mathcal{G} the path $\pi_X(\Lambda)$ will be an unparametrized quasi-geodesic.

We will also be given a combinatorial path $\Gamma = \{\gamma_i\}$ lying in \mathcal{G} . A reindexing function ties Γ to Λ . It will follow that Γ is also locally an unparametrized quasi-geodesic, with respect to the holes for \mathcal{G} . Unfortunately Γ is almost surely too long to give the upper bound; it may spend time traveling through $\mathcal{C}(Y)$ where Y is not a hole.

Our goal is to partition Γ into *basic* intervals where the path Γ makes definite progress through $\mathcal{C}(S)$, *shortcut* intervals where Γ wastes time in $\mathcal{C}(Y)$ for some non-hole $Y \subset S$, and *inductive* intervals where Γ enters and exits $\mathcal{C}(X)$ for some strict hole $X \subset S$. We do not alter Γ inside of the basic intervals, we shorten Γ so that it only spends a constant amount of time in non-holes, and we call on induction to cope with strict holes. A similar scheme then gives an upper bound on the length of the shortened Γ . Theorem 12.1 then follows. We begin by setting out the path requirements.

12.1. The requirements. We are given the following data: $\mathcal{G} = \mathcal{G}(S)$, a combinatorial complex, as well as vertices $\alpha, \beta \in \mathcal{G}$. We are also given $\Lambda = {\mu_n}_{n=0}^N$ a collection of markings. There is also a sufficiently large constant c_0 , depending only on the topology of S.

Definition 12.2. Recall that a marking $\mu = \{\nu_i\}$ is a set of curves in S with the following properties:

- $\iota(\nu_i, \nu_j) \leq c_0$ and
- the ν_i fill S.

Requirements for the marking path.

- (1) $\iota(\alpha, \mu_0), \iota(\beta, \mu_N) \leq c_0.$
- (2) For any hole X for \mathcal{G} , the map $n \mapsto \pi_X(\mu_n)$ is an unparametrized quasi-geodesic with uniform constants.
- (3) Λ locally satisfies the reverse triangle inequality: For any subsurface X, and for any m < n < p, we have

 $d_X(\mu_m, \mu_p) \ge d_X(\mu_m, \mu_n) + d_X(\mu_n, \mu_p) - c_0.$

From the marking path $\{\mu_n\}$ we will extract a *combinatorial* path $\Gamma = \{\gamma_i\}_{i=0}^K$ consisting of vertices in \mathcal{G} . We also will have a *reindexing* function $r: [0, K] \to [0, N]$. There is a uniform constant $c_1 \in \mathbb{R}$ with the following properties:

Requirements for the combinatorial path.

- (1) $\alpha = \gamma_0$ and $\beta = \gamma_K$.
- (2) The reindexing map is strictly increasing.
- (3) $\iota(\gamma_i, \mu_{r(i)}) \leq c_1$, for every $i \in [0, K]$.
- (4) $\iota(\gamma_i, \gamma_{i+1}) \leq c_1$, for every $i \in [0, K]$.

Finally we will be given, for every essential subsurface $X \subset S$, a (possibly empty) interval $J_X \subset [0, N]$. For every $n \in J_X$ we say that X is *accessible* from the marking μ_n . There are uniform constants $c_2, c_3 \in \mathbb{R}$ with the following properties:

Requirements for accessibility.

- (1) If $d_X(\alpha,\beta) > c_2$ then J_X is nonempty.
- (2) If $m, n \leq \min J_X$ or if $m, n \geq \max J_X$ then $d_X(\mu_m, \mu_n) \leq c_3$.
- (3) If X is nested in Y then for any $m, n \in J_X$ we have $d_Y(\mu_m, \mu_n) \le c_3$.
- (4) If X and Y overlap then for any $m, n \in J_X \cap J_Y$ we have $d_X(\mu_m, \mu_n), d_Y(\mu_m, \mu_n) \leq c_3.$

12.2. **Basic, inductive, and shortcut intervals.** Our goal is to partition the indices of $\Gamma = \{\gamma_i\}_{i=0}^{K}$ into basic, inductive, and shortcut intervals. We first identify the relevant surfaces. Take $L_0 \geq c_3$ and sufficiently large. We are given $\alpha, \beta \in \mathcal{G}$ and a combinatorial path $\Gamma = \{\gamma_i\}_{i=0}^{K}$. Set $\alpha_S = \alpha$ and $\beta_S = \beta$. Define

$$B_S = \{ Y \subsetneq S \mid d_Y(\alpha_S, \beta_S) \ge 3L_0 \}.$$

In general we are given a hole $X \subset S$, an interval $I_X = [h, k] \subset [0, K]$, and a set of subsurfaces B_X . Set $\alpha_X = \gamma_h$ and $\beta_X = \gamma_k$. We are told that:

• If $X \neq S$, $d_X(\alpha_X, \beta_X) \geq L_0$.

- For all $Y \in B_X$, $Y \subsetneq X$ and $d_Y(\alpha_X, \beta_X) \ge 3L_0$.
- For all $Y \notin B_X$, if $Y \subsetneq X$ and Y is a hole then $d_Y(\alpha_X, \beta_X) \leq 5L_0$.

For any subinterval $[i, j] \subset I_X$ we define

 $B_X(i,j) = \{ Y \in B_X \mid d_Y(\gamma_i, \gamma_j) \ge 3L_0 \}.$

We may now define basic, shortcut, and inductive subintervals of I_X . **Definition 12.3.** If $B_X(i, j)$ is empty then [i, j] is a *basic* interval. We require:

(12.4)
$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_{a,b} d_X(\gamma_i, \gamma_j).$$

Definition 12.5. Suppose $Y \in B_X$ is a non-hole. If $r([i, j]) \subset J_Y$ then [i, j] is a *shortcut* interval associated to Y. We require a constant $L_1 = L_1(X)$ so that:

(12.6)
$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_{a,b} L_1$$

Definition 12.7. Suppose $Y \in B_X$ is a hole. If $r([i, j]) \subset J_Y$ and $d_Y(\gamma_i, \gamma_j) \geq L_0$ then [i, j] is an *inductive* interval associated to Y. The induction hypothesis is:

(12.8)
$$d_{\mathcal{G}}(\gamma_i, \gamma_j) \leq_{a,b} \sum [d_Z(\gamma_i, \gamma_j)]_c$$

where the sum is over all holes Z for $\mathcal{G}(S)$.

If [i, j] is inductive and associated to Y then we may write $I_Y = [i, j]$ and set $\alpha_Y = \gamma_i$, $\beta_Y = \gamma_j$. In addition we may define the set of subsurfaces:

$$B_Y = \{ Z \in B_X(i,j) \mid Z \subsetneq Y \}.$$

Lemma 12.9. For L_0 sufficiently large, the following holds. Suppose $I_Y = [i, j]$ is inductive. If $Z \subsetneq Y$, Z is a hole, and $Z \notin B_Y$ then $d_Z(\alpha_Y, \beta_Y) \leq 5L_0$.

Proof. If $d_Z(\alpha_Y, \beta_Y) < 3L_0$ then there is nothing to prove. Supposing the contrary, since $Z \subsetneq Y \subsetneq X$ and $Z \notin B_Y$, we have $d_Z(\alpha_X, \beta_X) < 3L_0$. As the path $\pi_Z(\mu_n)$ is an unparametrized quasi-geodesic in $\mathcal{C}(Z)$, the same holds for the path $\pi_Z(\gamma_k)$, because Z is a hole. Thus for L_0 sufficiently large, $\pi_Z(\alpha_Y)$ and $\pi_Z(\beta_Y)$ lie within L_0 of the geodesic between $\pi_Z(\alpha_X)$ and $\pi_Z(\beta_X)$. The conclusion follows from the triangle inequality.

12.3. The partition. We now partition the interval I_X . In general, we are given a subinterval $[i, j] \subset I_X$. (To begin with, $[i, j] = I_X$.)

If $B_X(i, j)$ is empty then [i, j] is basic and we are done. Suppose then that $B_X(i, j)$ is nonempty. If there is a hole $Y \in B_X(i, j)$ then choose that hole $Y \in B_X(i, j)$ so that

- Y is maximal with respect to the partial order (by inclusion) on $B_X(i, j)$
- $J_Y \cap r([i, j])$ is as large as possible. (Subject to the first condition.)

If there are no holes in $B_X(i, j)$ then instead take Y to be a non-hole satisfying these same two conditions as above.

In either case let $i', j' \in [i, j]$ be the first and last indices, respectively, so that $r(i'), r(j') \in J_Y$. Partition $[i, j] = [i, i' - 1] \cup [i', j'] \cup [j' + 1, j]$. (The first or third parts, or both, may be empty). Take $I_Y = [i', j']$ and set $\alpha_Y = \gamma_{i'}, \beta_Y = \gamma_{j'}$.

We will show later that if Y is a non-hole then I_Y is a shortcut interval. If Y is a hole, then set

$$B_Y = \{ Z \in B_X(i', j') \mid Z \subsetneq Y \}.$$

We must check that I_Y is an inductive interval – the only nontrivial point is that $d_Y(\gamma_{i'}, \gamma_{j'}) \geq L_0$. But this follows immediately from (2) of accessibility requirements and the triangle inequality.

Now partition [i, i' - 1] and [j' + 1, j] in similar fashion. This completes our recursive description to the partition of I_X .

We establish preliminary results about the partition.

Lemma 12.10. If Y, Z are holes that either overlap or are nested, then $I_Z \subset J_Y$ implies $Z \subset Y$.

Proof. This follows immediately from the definition of the intervals I_Z , the assumption $M_0 \ge C_3$, and the accessibility requirement (2) for nested subsurfaces and requirement (3) for overlapping surfaces. \Box

Proposition 12.11. For L_0 large enough, the partition of [m, n] has the properties that

- (1) For any hole Y there is at most 1 inductive interval I_Y associated to Y.
- (2) For any hole Y there are at most 2 inductive intervals I_Z such that $J_Y \cap I_Z \neq \emptyset$, such that $Y \subsetneq Z$, or Z and Y overlap. There is at most one interval I_Z for Z paired with Y.
- (3) Suppose there exists $l \ge N_1$, intervals I_{Y_i} that are shortcut intervals associated to a nonholes Y_i or inductive intervals associated to holes Y_i . Then $\{\partial Y_i\}$ fill X, and $d_X(\partial Y_1, \partial Y_l) \ge l/N_1 2$. In particular, there are at most N_1 shortcut intervals I_Y associated to any fixed Y.

Proof. We prove the first conclusion. If there were two intervals $I_1 = [p_1, q_1]$ and $I_2 = [p_2, q_2]$ associated to the same hole Y with $q_1 \leq p_2$,

then we claim that

$$d_Y(\mu_{q_2},\mu_{p_1}) \ge L_0.$$

If that is the case, then we would have a contradiction, for then an interval, either $[p_1, q_2]$ or an interval containing it, would have been chosen as the inductive interval rather than either $[p_i, q_i]$. We need to prove the claim. Let $\alpha_Y = \pi_Y(\gamma_{p_1})$ and $\beta_Y = \pi_Y(\gamma_{q_2})$. We notice that since the path $\pi_Y(\mu_n)$ is an unparametrized quasi-geodesic, and $\mathcal{C}(Y)$ is a hyperbolic space, there is a constant D depending on these constants such that for any two points on the path, any point between them is within D of the geodesic joining the endpoints. If $q_1 = p_2$ we apply it to the endpoints α_Y and β_Y with $\pi_Y(\gamma_{q_1})$ between them to conclude that

$$d_Y(\beta_Y, \alpha_Y)) \ge 2L_0 - 2D \ge L_0,$$

for L_0 large enough, proving the claim. If $p_2 > q_1$, then dividing the interval $[p_1, q_2]$ into $[p_1, p_2]$ with interior point q_1 , and $[p_2, q_2]$, we conclude that

$$d_Y(\mu_{q_2}, \mu_{p_1}) \ge 2L_0 - 3D \ge L_0,$$

again for L_0 big enough, proving the claim.

We prove the second conclusion. The interval J_Y may be contained in a single I_W for W paired with Y. If J_Y intersected three or more I_Z such that either Y and Z overlap, or $Y \subset Z$, then J_Y would contain one of them. This is impossible by Lemma 12.10. Thus it intersects at most 2 such I_Z .

We prove the third statement. If there were more than N_1 such I_{Y_i} , then Lemma ?? says that some hole or nonhole Z containing a collection of these Y_i would satisfy $d_Z(\mu_{m_1}, \mu_{n_k}) \ge 3M_0$, where μ_{m_1} is the furthest point to the left, and μ_{n_k} the furthest to the right. Suppose Z is a proper subsurface of X. If Z is a hole, then either Z or some W containing Z would have been chosen to define an inductive interval I rather than any of the I_{Y_i} . This is a contradiction. If Z is a nonhole, then since it contains no holes, either it or some W containing it would have been chosen rather than I_{Y_i} , and again we have a contradiction.

Thus Z = X and so the $\{\partial Y_i\}$ fill X. Again let g_X be a geodesic in $\mathcal{C}(X)$ joining γ_i and γ_j . For each Y_i there is a curve $\omega_i \in m$ such that $d_X(\omega_i, \partial Y_i) \leq 1$. For any ω , the set of ∂Y_i within distance 1 of ω do not fill X. There are at most N_1 of these Y_i . This forces the total number of $\omega \in g_X$ that are within distance 1 of some ∂Y_i to be at least l/N_1 , giving the desired estimate.

12.4. The upper bound. We are now equipped to prove Theorem 12.1. Recall we are given curves γ_h, γ_k and a hole X. We construct a path l

in \mathcal{G} joining γ_h and γ_k . For each basic interval [r(i), r(j)] with corresponding curves γ_i, γ_j we take the given subpath joining γ_i and γ_j . For each shortcut interval [r(i), r(j)] we take the shortcut path of length at most N_0 joining γ_i, γ_j , guaranteed to exist by assumption (??). For each inductive interval we take the path defined inductively through that interval, guaranteed to exist by (??).

Proposition 12.12. For some fixed a, b, c the path l satisfies

$$|l| \leq_{a,b} \sum [d_Y(\alpha,\beta)]_c,$$

the sum over holes $Y \subset X$.

Proof. Let N_S be the number of shortcut intervals. The path though each shortcut interval is bounded by N_0 . Therefore we need to prove that N_S is bounded in terms of $d_X(\gamma_h, \gamma_k)$. By Proposition 12.11 1), and 2) any hole Y is associated to at most 1 inductive interval $I_Y =$ [r(i), r(j)] and any J_Y overlaps at most 2 other $I_{Z_1} = [r(p_1), r(q_1)], I_{Z_2} =$ $[r(p_2), r(q_2)]$ and we can assume $r(q_1) \leq r(i) < r(j) \leq r(p_2)$. The path in $\mathcal{C}(Y)$ is an unparametrized quasi-geodesic. Therefore

$$d_Y(\gamma_{p_1}, \gamma_{q_1}) + d_Y(\gamma_i, \gamma_j) + d_Y(\gamma_{p_2}, \gamma_{q_2}) \leq_{a,b} d_Y(\gamma_{p_1}, \gamma_{q_2}).$$

Thus the path length of l through each hole Y inductively is bounded by $K_1 d_Y(\gamma_h, \gamma_l) + K_2$, for fixed K_1, K_2 . Thus if we let N_I denote the number of inductive intervals it follows that it is enough to prove that there are constants a, b such that

(12.13)
$$\sum_{I_B} |I_B| + N_S + N_I \leq_{a,b} d_X(\gamma_h, \gamma_k),$$

where the sum is over all basic intervals I_B . We prove this bound.

Let δ be the hyperbolicity constant for the space $\mathcal{C}(X)$. The fact that the sequence $\{\gamma_i\}$ is an unparametrized quasi-geodesic in $\mathcal{C}(X)$ says that there are constants a, b and $D > \delta$ such that for any partition of the $\{\gamma_i\}$,

(12.14)
$$P = \gamma_i = \beta_1, \beta_2 \dots, \beta_M = \gamma_j$$

with the property that for all $p, d_X(\beta_p, \beta_{p+1}) \ge D$, then

(12.15)
$$\sum_{p=1}^{M-1} d_X(\gamma_p, \gamma_{p+1}) \leq_{a,b} d_X(\gamma_i, \gamma_j)$$

We partition the entire interval [r(h), r(l)] as

$$I_1 \cup J_1 \cup I_2 \cup J_2 \cup \ldots \cup I_k \cup J_k,$$

where the intervals I_i are basic and satisfy

 $|I_i| \geq 4D.$

Each interval J_i is a connected union of basic, shortcut, and inductive intervals with the following properties. Each J_i begins with either a shortcut or inductive interval associated with some proper subsurface Y_i^1 . It is followed by a possibly empty basic interval with length at most 4D, followed by a shortcut or inductive interval associated to some Y_i^2 , then a basic interval again with length at most 4D, and so forth ending with a shortcut or inductive interval associated to some subsurface $Y_i^{l(i)}$. (We may have $I_1 = \emptyset$ or $J_k = \emptyset$.)

We wish to form a partition P of the entire interval [r(h), r(k)]. Suppose the shortcut and inductive intervals in J_i are associated to $Y_i^1, \ldots Y_i^{n_i}$. Let ω_i, ω'_i the initial and terminal curves. Let

$$K_1 = K_0 + C_2$$

By (??)

$$d_X(\omega_i, \omega_i') \ge d_X(\partial Y_i^1, \partial Y_i^{n_i}) - 2K_1$$

The contribution to the left side of (12.13) from J_i is

(12.16)
$$n_i + \sum_{I_B \cap J} |I_B| \le n_i (4D+1)$$

By 3) of Proposition 12.11,

$$d_X(\partial Y_1, \partial Y_i^{n_i})) \ge n_i/N_1 - 2$$

and therefore

(12.17)
$$d_X(\omega_i, \omega'_i) \ge n_i/N_1 - 2 - 2K_1.$$

We group the set of J_i into two types. The first possibility is that

$$d_X(\omega_i, \omega_i') \le D$$

 \mathbf{SO}

$$n_i/N_1 - 2 - 2K_1 \le D.$$

Then

$$n_i + \sum_{I_B \cap J_i} |I_B| \le n_i (4D+1) \le (4D+1)(D+2+2K_1)N_1.$$

This quantity then can by incorporated into the length through a basic interval of length at least 4D bordering J_i by multiplying by a fixed factor (depending on D) and so can be ignored in terms of (12.13). In our partition P we include the midpoints α_i, τ_i of the basic intervals I, I' on each side of J_i . Let $I_{1/2} \subset I, I'_{1/2} \subset I'$ the half intervals bordering J_i . Then a simple thin triangle argument in the hyperbolic space $\mathcal{C}(X)$, and the fact that $D > \delta$, shows that

$$d_X(\alpha_i, \tau_i) \ge D_s$$

and furthermore

$$|I_{1/2}| + |I'_{1/2}| \leq_{a,b} d_X(\alpha_i, \tau_i)$$

The second possibility for J_i is that the endpoints satisfy

$$d_X(\omega_i, \omega_i') \ge D.$$

In this case we include the endpoints ω_i, ω'_i in the partition. Combining (12.16) and (12.17) we get a bound

(12.18)
$$n_i + \sum_{I_B \cap J_i} |I_B| \le (4D+1)N_1(2+2K_1+d_X(\omega_i,\omega_i')).$$

We now have a partition of the sequence of curves such that two consecutive curves are at least D apart. The points of the partition include the endpoints of intervals J_i , if J_i is of the second type, and the midpoints of intervals I surrounding J_i if it is of the first type.

Let $I_p = [r(p), r(p')]$ be a basic interval of length at least 4D or half of a basic interval of length at least 2D with endpoints $\gamma_p, \gamma_{p'}$. By (??) there exists C_4 such that

(12.19)
$$d_X(\gamma_p, \gamma_{p'}) \ge \frac{|I_p|}{C_4}$$

Combining (12.18) and (12.19) and applying (12.15), we conclude that the left side of (12.13) which can be expressed as $\sum_{p=1}^{k} |I_p| + \sum_{p=1}^{k} n_p + \sum_{I_B \cap J_i} |I_B|$ satisfies $\sum_{p=1}^{k} |I_p| + \sum_{p=1}^{k} n_p + \sum_{I_B \cap J_i} |I_B| \le C_4 \sum_{p=1}^{n} d_X(\gamma_p, \gamma'_p) + (4D+1)N_1(2+2K_1+\sum_{p=1}^{n} d_X(\gamma_p, \gamma'_p))$

$$\sum_{p} |I_p| + \sum_{p} n_p + \sum_{I_B \cap J_p} |I_B| \le C_4 \sum_{p} d_X(\gamma_p, \gamma_p) + (4D+1)N_1(2+2K_1 + \sum_{p} d_X(\omega_p, \omega_p))$$
$$\le_{a,b} d_X(\gamma_h, \gamma_k).$$

13. BACKGROUND ON TEICHMÜLLER SPACE

Fix now a surface $S = S_{g,n}$ of genus g with n punctures. Let $\mathcal{M} = \mathcal{M}(S)$ be the *Teichmüller space* of S. This is slightly non-standard notation; however this fits with our general framework where \mathcal{M} is a marking space.

Recall that $\mathcal{M}(S)$ is the set of equivalence classes of conformal structures μ on S. Here $\mu_1 \sim \mu_2$ if there is a conformal map $f: \mu_1 \to \mu_2$

which is isotopic to the identity. We equip Teichmüller space with the Teichmüller metric. That is,

$$d_{\mathcal{M}}(\mu_1, \mu_2) = \inf_f \left\{ \frac{1}{2} \log K(f) \right\}.$$

Here K(f) is the maximal dilatation of f. The infimum ranges over all maps $f: \mu_1 \to \mu_2$ isotopic to the identity and it is realized by a Teichmüller map. This, in turn, may be defined in terms of a quadratic differential.

Definition 13.1. A holomorphic quadratic differential $q(z)dz^2$ on a Riemann surface μ_0 , possibly with punctures, is an assignment of a holomorphic function $q_z(z)$ to each local holomorphic coordinate chart. If z and ζ are overlapping charts then we require the equality

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

to hold in the intersection of the charts.

The sum of the orders of the zeroes and poles of q totals 4g - 4. At any point away from the zeroes and poles 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. For every $t \in \mathbb{R}$ there is a new Riemann surface μ_t and associated quadratic differential $g_t(q)$ on μ_t . We are also given a map $f_t: \mu_0 \to \mu_t$, called a *Teichmüller map*. The Teichmüller map is defined by the property that in the natural coordinates z = x + iy of q and $z_t = x_t + iy_t$ of $g_t(q)$ we have

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

The Teichmüller map is the unique map realizing the infimum in the definition of the Teichmüller distance.

Let σ_t be the unique hyperbolic metric in the conformal class of μ_t . Let P_t be a *Bers pants decomposition* of S;

 P_t of shortest length curves, called a Bers pants decomposition and also a shortest marking, again denoted μ_t . For any subsurface X there is a projection $\pi_X : \mathcal{M} \to \mathcal{C}(X)$ which first associates to $\mu \in \mathcal{M}$ its minimal length marking and then projects that marking into X. This map is Lipschitz. Let ϵ_1 a Margulis constant so that intersecting curves cannot both be of Poincare length smaller than ϵ_1 .

I don't think that is true. Consider an annulus.

14. Paths for the nonorientable surface

Again fix F a compact, connected, and nonorientable surface. Let S be the orientation double cover with covering map $\rho_F \colon S \to F$. Let $\tau \colon S \to S$ be the associated involution. Let $\nu \colon \mathcal{C}^{\tau}(S) \to \mathcal{C}(S)$ be the natural map from the complex of τ -invariant curves to $\mathcal{C}(S)$. As noted above, the inclusion is Lipschitz.

Our goal here is to obtain a reverse inequality:

$$d_{\mathcal{C}^{\tau}}(\alpha,\beta) \geq_{a,b} d_S(\alpha,\beta).$$

We begin by giving the promised proof of:

Lemma 14.1. There is a constant K with the following property: Suppose that α, β are invariant multi-curves or arcs in S and $X \subset S$ so that $d_X(\alpha, \beta) > K$. Then X is symmetric.

Proof. Let $Y = \tau(X)$. Since α and β are invariant deduce $d_Y(\alpha, \beta) = d_X(\alpha, \beta)$. We can assume $d_{\mathcal{C}}(\alpha, \beta) \geq 3$. This implies that α, β fill the surface S, and so there exists a quadratic differential q on some Riemann surface X_0 which has as its horizontal leaves, closed curves homotopic to α and vertical leaves homotopic to β . If α has two components then we require that the moduli of the corresponding cylinders coincide. The same is true for β . Since τ preserves α and β , τ^*q is another quadratic differential on X_0 all of whose horizontal leaves are homotopic to α and vertical leaves homotopic to β . But then $\tau^*q = q$ and so q is symmetric.

Let $g_t(q), -\infty \leq t \leq \infty$ the corresponding path of quadratic differentials. Each is τ symmetric. For each time t, let σ_t be the hyperbolic metric on the surface of $g_t(q)$ and μ_t the shortest marking. The hyperbolic metric is τ invariant. Let ϵ_1 a Margulis constant so that intersecting curves cannot both be of Poincare length smaller than ϵ_1 .

By Theorem () of Rafi there is a constant K such that if $d_X(\alpha, \beta) \ge 3K$, there are $t_1 < t_2$ so that

(14.2)
$$d_X(\mu_{t_1}, \mu_{t_2}) \ge K$$

and such that for all $t_1 \leq t \leq t_2$

$$\sigma_t(\partial X) < \epsilon_1.$$

Since σ_t is symmetric, we have $\sigma_t(\tau(\partial X)) = \sigma_t(\partial X)$. This implies $i(\partial X, \tau(\partial X) = \emptyset$. We cannot have any component of $\tau(\partial X)$ contained in X for this would violate (14.2). Thus either $\tau(X) = X$ or $\tau(X) \cap X = \emptyset$.

As $t \to -\infty$, $\rho_t(\alpha) \to 0$ and no other curve has hyperbolic length going to 0 and as t We can let We have a sequence γ_n of shortest curves on these surfaces. The function

 $i: \mathcal{AC} \times \mathcal{M} \to \mathbb{R}^+$

is given by the length of the shortest geodesic in the homotopy class.

We need to define sequences $\alpha = \gamma_1, \ldots, \gamma_I = \beta$ of curves, a corresponding sequence μ_n of points in Teichmuller space that satisfy the path requirements. We also have to define the accessible intervals so as to satisfy the Accessible requirements. We also need then to check (??), (??) and (??).

We check the marking requirements on our sequences in Teichmuller space and curves. We already have Condition 1). Condition 2) is exactly Theorem () of [?] for the whole surface, and Theorem 1 of [?] in the case of subsurfaces. Condition 3) holds since intersection number is always a bound for distance. shortest curves can intersect a bounded number of times.

We check the combinatorial path requirements. In all cases there was no reindexing so 2) holds. Condition 3) holds since bounded length curves can intersect only a bound number of times. A bounded length surgery path produces curves that connect them. Condition 4) follows for example from the distance formula given in Theorem 1.1 of [?]. We are given

$$d_T(\mu_i, \mu_{i+1}) = 1,$$

where $d_T(\cdot, \cdot)$ is Teichmuller distance. On the other hand

$$d_T(\mu_i, \mu_{i+1}) \stackrel{a,b}{=} d_Y(\mu_i, \mu_{i=1}).$$

This implies there is a bound for any subsurface projection which implies there is a bound for $i(\gamma_i, \gamma_{i+1})$

15. Paths for the arc complex

Here we build the pair of paths required to obtain the upper bound on distance for the arc complex $\mathcal{A}(S, \Delta)$.

16. Background on train tracks

Here we give the necessary definitions and theorems regarding train tracks. This treatment is essentially due to Lee Mosher.

Recall that a generic train track $\tau \subset S$ is a smooth (C^2) 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 $N = N(\tau)$ be an closed *I*-bundle neighborhood of τ . Note that ∂N splits into two collections of arcs: the *vertical boundary* $\partial_v N$ and the *horizontal boundary* $\partial_h N$. (If τ has no switches then $\partial_v N$ is empty.) See Figure 8. Conversely, given a foliated neighborhood N, taking the quotient of the *I*-fibres recovers τ , up to isotopy.



FIGURE 8. A vertically foliated neighborhood of τ .

A simple closed curve α is carried by τ , written $\alpha \prec \tau$, if α lies in the interior of $N(\tau)$, transverse to the vertical fibres. If α is an arc then we additionally require that $\partial \alpha \subset \partial_v N$. The weights of α on τ are then a non-negative integer for each branch of τ : the number of times α meets any fibre in the neighborhood of that branch. The set of all allowable weights on τ , denoted $P(\tau)$ is a cone over \mathbb{Z} . A simple closed curve carried by τ , with weights lying on an extremal ray of $P(\tau)$, is called a vertex of τ . Denote the set of vertices of τ by $\operatorname{Vert}(\tau)$. If τ and σ are tracks, and $Y \subset S$ is an essential surface, then define $d_Y(\tau, \sigma) = d_Y(\operatorname{Vert}(\tau), \operatorname{Vert}(\sigma))$.

If $\alpha \prec \tau$ then we may *split* the track τ along α : Form $N = N(\tau)$, let $n(\alpha)$ be an open neighborhood of α , let $N' = N - n(\alpha)$, and set $\tau' = N'/I$. If all weights of α on τ are two or less then we call τ' a *central splitting* of τ . Notice that there are only finitely many central splittings possible of any given track.

A train track σ is a *subtrack* of τ if σ is obtained by deleting branches of τ . In this situation we write $\sigma \subset \tau$. We now have a crucial definition:

Definition 16.1. An essential surface $Y \subset S$ is *accessible* from τ if Y can be realized as follows: there is a subtrack $\sigma \subset \tau$ and a central splitting σ' of σ so that a regular neighborhood N of σ' , plus any disk components of $S \setminus N$ to N, is isotopic to Y. We will denote σ' by $\tau|Y$.

Actually, I'm pretty sure that $\tau | Y$ is only defined up to some fuzzyness...

There is a bound, depending only on S, on the number of surfaces accessible from any one given train track.

We say that a train track σ is obtained from τ by *sliding* if σ and τ are related as in Figure 9. We say that a train track σ is obtained from τ by *splitting* if σ and τ are related as in Figure 10.



FIGURE 9. All slides take place in a small regular neigborhood of the affected branch.



FIGURE 10. There are three kinds of splitting: right, left, and central.

If σ is a train track obtained from τ by a sequence of slides and exactly one splitting then we say that σ is obtained via a *wide splitting*. In this case we write $\sigma \prec \tau$. We remark that if there is a wide splitting from τ to σ then there is one using an *a priori* bounded number of slides. We shall always use these parsimonious wide splittings. For an in-depth discussion see Section 3.13 of Mosher's monograph [Mos].

Definition 16.2. A splitting sequence is a collection $\{\tau_i\}$ of train tracks so that $\tau_0 \succ \tau_1 \succ \ldots \succ \tau_n$.

We may now state the first of Mosher's results:

Theorem 16.3 (Mosher). Fix a surface S. There are constants (a, b, c)with the following property: Suppose that $\{\tau_i\}$ is a splitting sequence in S and $Y \subset S$ is an essential surface. Then the map $i \mapsto \pi_Y(\operatorname{Vert}(\tau_i))$ is an (a, b, c) unparametrized quasi-geodesic.

Note that, when Y = S, Theorem 16.3 is essentially due to the first author and Minsky; see Theorem 1.3 of [MM]. To see why Theorem 16.3 holds when Y a strict subsurface we define:

Definition 16.4. For any essential subsurface $Y \subset S$ and any splitting sequence $\{\tau_i\}$ define the *accessible region* J_Y to be set of indices so that Y is accessible from τ_i .

Lemma 16.5 (Mosher). The accessible region J_Y is a (possibly empty) interval.

We now have:

Theorem 16.6 (Mosher). Fix a surface S. There is a constant K with the following property: Suppose that $\{\tau_i\}$ is a splitting sequence and Y is an essential subsurface. There is a splitting sequence $\{\sigma_k\}$ in Y, parametrized by $J_Y = [i, j]$, so that

• If $k \leq i$ (respectively, $j \leq k$), then

 $d_Y(\tau_k, \tau_i) \leq K \quad (resp. \ d_Y(\tau_j, \tau_k) \leq K).$

• For $k \in J_Y$, the tracks $\tau_k | Y$ and σ_k have splitting distance at most K.

Note that Theorem 16.3 follows from the cited result of [MM] and Theorem 16.6.

17. Paths for the disk complex18. The distance estimate

Needs work.

Needs work.

Theorem 18.1. For any handlebody V there is a constant $c_0 = c_0(V)$ so that, for any $c \ge c_0$ there are constants $a \ge 1$ and $b \ge 0$ satisfying

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

independent of the choice of D and E. The sum ranges over subsurfaces $X \subset \partial V$ which are holes for $\mathcal{D}(V)$.

19. Hyperbolicity

We remark that some of the ideas in this section are similar to the heirarchy machine of [MM00] (see also Chapters 4 and 5 of Behrstock's thesis [Beh04]). Our discussion has a different focus; we avoid the

heirarchy machine as it appears to be too rigid to deal with the disk complex. As another difference from [MM00] and [Beh04] we rely on Theorem 2.8 instead of proving that projection maps are coarsely contracting.

We wish to prove:

Theorem 19.1. Suppose that $\mathcal{G}(S)$ satisfies the distance estimate and the path requirements. Suppose also that all holes for \mathcal{G} interfere. Then \mathcal{G} is Gromov hyperbolic.

As corollaries we have

Theorem 19.2. The arc complex is Gromov hyperbolic.

Theorem 19.3. The disk complex is Gromov hyperbolic.

In fact, Theorem 19.1 follows quickly from:

Theorem 19.4. Fix \mathcal{G} , a combinatorial complex. Suppose that \mathcal{G} satisfies the distance estimate. Suppose that all holes for \mathcal{G} interfere. For any constants (a, b, c) there is a δ so that for any triangle of paths $T \subset \mathcal{G}$ we have: if the projection of any side of T into into any hole is an (a, b, c) unparametrized quasi-geodesic, then T is δ -slim.

Proof of Theorem 19.1. Suppose that $\mathcal{G}(S)$ satisfies the distance estimate and the path requirements. Suppose also that all holes for \mathcal{G} interfere. Then, by Theorem ???, there are uniform constants (a, b, c) so that for any pair $\alpha, \beta \in \mathcal{G}$ there is a path $\mathcal{P} = \{\gamma_i\} \subset \mathcal{G}$ with

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

So if $\alpha \cap \beta = \emptyset$ then $|\mathcal{P}|$ is uniformly short. Also, by Theorem 19.4, triangles made of such paths are uniformly slim. Thus, by Theorem 2.8, \mathcal{G} is Gromov hyperbolic.

The rest of this section is devoted to proving Theorem 19.4.

19.1. 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 $h \in \mathcal{C}(X)$ connecting a point of $\pi_X(\alpha)$ to a point of $\pi_X(\beta)$ we also define $\rho_h: \mathcal{G} \to h$ to be the map $\pi_X | \mathcal{G}: \mathcal{G} \to \mathcal{C}(X)$ followed by closest points projection to h. Define index_X: $\mathcal{G} \to \mathbb{N}$ to be the *index* in X:

$$\operatorname{index}_X(\sigma) = d_X(\alpha, \rho_h(\sigma)).$$

48

Remark 19.5. Suppose that h' is a different geodesic connecting $\pi_X(\alpha)$ to $\pi_X(\beta)$ and index'_X is defined with respect to h'. Then

$$|\operatorname{index}_X(\sigma) - \operatorname{index}'_X(\sigma)| \le 12\delta + 2$$

by Lemma 2.7. Thus, if we are willing to accept a small additive error, the choice of geodesic h is irrelevant. Accordingly we will supress the superscript whenever possible.

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

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

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 . We assume that σ_i and σ_{i+1} are disjoint.

We now formalize a pair of properties enjoyed by unparametrized quasi-geodesics to the situation at hand. 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

$$d_X(\sigma_i, \rho_h(\sigma_i)) \le K,$$

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

Remark 19.6. As in Remark 19.5, 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 $\rho_h(\sigma_i)$ is more convenient.

19.3. **Projection control.** We say domains $X, Y \subset S$ overlap if X and Y intersect but are not nested. The following theorem (see Theorem 4.2.1 of Behrstock's thesis [Beh04]) follows from Masur and Minsky's idea (see [MM00]) of *time ordered* domains in S:

Theorem 19.7. 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 $d_Y(\partial X, \gamma) < M_1$.

We also require a more specialized version of Theorem 19.7 for the case where X and Y are nested.

Lemma 19.8. There is a constant $M_2 = M_2(S)$ with the following property. Suppose that $X \subset Y$ are nested non-simple domains. Fix $\alpha, \beta, \gamma \in \mathcal{AC}(S)$ which cut both X and Y. 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 $d_X(\alpha, \beta) \ge M_0$, the constant given by the Bounded Geodesic Image Theorem [MM00]. If $d_X(\alpha, \gamma) \ge M_2$ then

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

Symmetrically, we have

$$\operatorname{index}_Y(\gamma) \le \operatorname{index}_Y(\partial X) + 4$$

if $d_X(\gamma, \beta) \ge M_2$.

19.4. Finding the midpoint of a side. Let (a, b, c) be arbitrary. 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, b, c) unparametrized quasigeodesics 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 $K_1 \ge \max\{M_0, 4M_1, M_2, 8\} + 6\delta$ sufficiently large so that any (a, b, c) unparametrized quasi-geodesic in any hole back and side tracks at most K_1 .

Claim 19.9. 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

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

Thus $\operatorname{index}_X(\sigma_i) + 4K_1 \leq \operatorname{index}_X(\sigma_j)$. 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 $d_X(\alpha, \beta) \ge 2K_1 \ge M_0$ and $d_X(\alpha, \gamma) \ge 2K_1 - 2\delta \ge M_2$. Apply Lemma 19.8 to deduce $\operatorname{index}_Y(\partial X) - 4 \le \operatorname{index}_Y(\gamma)$. Since σ_j succeeds γ in Y it follows that $\operatorname{index}_Y(\partial X) - 4 + 2K_1 \le \operatorname{index}_Y(\sigma_j)$. Again using the fact that σ_i precedes γ in X we have that $d_X(\sigma_i, \beta) \ge M_2$. We deduce from Lemma 19.8 that $\operatorname{index}_Y(\sigma_i) \le \operatorname{index}_Y(\partial X) + 4$. Thus

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

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

Suppose now that X and Y overlap. Applying Theorem 19.7 and breaking symmetry, we may assume that $d_X(\gamma, \partial Y) < M_1$. Since σ_i

precedes γ we have $\operatorname{index}_X(\gamma) \ge 2K_1$. Thus, it follows that $\operatorname{index}_X(\partial Y) \ge 2K_1 - 2M_1 - 4\delta$ and so

$$d_X(\alpha, \partial Y) \ge 2K_1 - 2M_1 - 6\delta \ge M_1.$$

Applying Theorem 19.7 again, we find that $d_Y(\alpha, \partial X) < M_1$. Now, since σ_j succeeds γ in Y, we deduce that $\operatorname{index}_Y(\sigma_j) \geq 2K_1$. Similar considerations to the above show that

$$d_Y(\partial X, \sigma_i) \ge 2K_1 - M_1 - 2\delta \ge M_1.$$

Applying Theorem 19.7 one last time, we find that $d_X(\partial Y, \sigma_j) < M_1$. Thus $d_X(\gamma, \sigma_j) \leq 2M_1$. Finally, we deduce that the difference in index (in X) between σ_i and σ_j is at least $2K_1 - 4M_1 - 4\delta$. Since this is again greater than K_1 , it follows that i < j.

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 19.10. There is a constant $N_1 = N_1(S)$ with the following property. For every hole X and geodesic h connecting $\pi_X(\alpha)$ to $\pi_X(\beta)$:

$$d_X(\sigma_\alpha, \rho_h(\gamma)) \le N_1$$

Proof. Since σ_i and σ_{i+1} are disjoint we have

$$|\operatorname{index}_X(\sigma_{i+1}) - \operatorname{index}_X(\sigma_i)| \le 4\delta + 2.$$

Since \mathcal{P} is a path connecting α to β the image $\rho_h(\mathcal{P})$ is $4\delta + 2$ -dense in h. Thus, if $\operatorname{index}_X(\sigma_\alpha) + 2K_1 + 4\delta + 2 < \operatorname{index}_X(\gamma)$ then we have a contradiction to the definition of σ_α .

On the other hand, if $\operatorname{index}_X(\sigma_\alpha) \geq \operatorname{index}_X(\gamma) + K_1$ then σ_α succeeds γ . This directly contradicts Claim 19.9.

We deduce that the difference in index between σ_{α} and γ in X is at most $2K_1 + 4\delta + 2$. Finally, as \mathcal{P} sidetracks by at most K_1 we have

$$d_X(\sigma_\alpha, \rho_h(\gamma)) \le 3K_1 + 4\delta + 2$$

as desired.

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 19.9, we immediately deduce that σ_{α} comes before σ_{β} in \mathcal{P} . A symmetric version of Claim 19.10 applies to σ_{β} : for every hole X

$$d_X(\rho_h(\gamma), \sigma_\beta) \le N_1$$

19.5. 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 the closest points projection of β to geodesics of the form $l = [\pi_X(\alpha), \pi_X(\gamma)]$. By Claim 19.10, for every hole X, τ_{α} and τ_{γ} are close to $\rho_l(\beta)$.

By Lemma 2.5, if $h = [\pi_X(\alpha), \pi_X(\beta)]$, then $d_X(\rho_h(\gamma), \rho_l(\beta)) \leq 6\delta$. We deduce:

Claim 19.11. $d_X(\sigma_{\alpha}, \tau_{\alpha}) \leq 2N_1 + 6\delta$.

We now prove:

Claim 19.12. 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

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

for every hole X.

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

Suppose first that no vertex of \mathcal{R} precedes σ_i by more than $2K_1$. Fix a hole X and geodesics $h = [\pi_X(\alpha), \pi_X(\beta)]$ and $l = [\pi_X(\alpha), \pi_X(\gamma)]$. Then $\rho_l(\sigma_i)$ is within distance $2K_1$ of $\pi_X(\alpha)$. Appealing to Claim 19.11, bounded sidetracking, and hyperbolicity of $\mathcal{C}(X)$ we find that the initial segments

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

of h and l respectively must fellow travel. Because of bounded backtracking along \mathcal{P} , $\rho_h(\sigma_i)$ lies on, or at least near, this initial segment of h. Thus by Lemma 2.7 $\rho_l(\sigma_i)$ is close to $\rho_h(\sigma_i)$ 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 19.9 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 19.10.

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 19.4.

Definition 20.1. Let $\rho_V \colon \mathcal{C}(S) \to \mathcal{P}(\mathcal{D}(V))$ be the closest points projection map: for every $\gamma \in \mathcal{C}(S)$ take $\rho_V(\gamma)$ equal to the set of $D \in \mathcal{D}(V)$ so that $d_S(\gamma, D) = \min\{d_S(\gamma, E) \mid E \in \mathcal{D}(V)\}.$

Recall that a maximal simplex \mathbb{D} in $\mathcal{D}(V)$ is a *pants decomposition* of V. The first application of our ideas is the "projection to handlebodies" theorem:

Theorem 20.6. There is a constant $R_1 = R_1(V)$ and an algorithm which, given an essential curve $\gamma \subset S$ and a pants decomposition $\mathbb{D} \subset \mathcal{D}(V)$, finds a disk $E \in \mathcal{D}(V)$ so that

$$d_S(E, \rho_V(\gamma)) \le R_1.$$

Recall the definition of the closest points projection map $\rho_V \colon \mathcal{C}(S) \to \mathcal{P}(\mathcal{D}(V))$. For every $\gamma \in \mathcal{C}(S)$ we take $\rho_V(\gamma)$ equal to the set of $D \in \mathcal{D}(V)$ so that $d_S(\gamma, D) = \min\{d_S(\gamma, E) \mid E \in \mathcal{D}(V)\}.$

Recall that a maximal simplex \mathbb{D} in $\mathcal{D}(V)$ is a *pants decomposition* of V. The goal of this section is to prove Theorem 20.6: there is an algorithm which finds a point of $\rho_V(\gamma)$, up to bounded error.

Fix a handlebody V with $\partial V = S$. Fix γ an essential curve in S and \mathbb{D} a pants decomposition of V. Let D be any element of $\rho_V(\gamma)$.

Let H be a hierarchy connecting \mathbb{D} to γ . Let σ be the last slice of H preceding D on H. Note that $d_S(\sigma, D)$ is bounded by $2\delta + R$, the latter being the quasi-convexity constant of $\mathcal{D}(V)$.

Claim 20.2. There is a multicurve μ so that $d_X(\sigma, \mu)$ is uniformly bounded for all essential subsurfaces X and either

- μ is the boundary of a disk,
- μ is the boundary of a non-hole, or
- μ is the boundary of a hole of large diameter.

Proof. Connect σ to D by a hierarchy L. Connect D to \mathbb{D} by a hierarchy K. If |L| is small then we may take $\mu = \partial D$. Otherwise, as $d_S(\sigma, D)$ is bounded, we may assume that the hierarchy L contains at least one large projection in a strict subsurface. Let X be the first such, under time-order. We may assume that $d_X(\sigma, D)$ is large.

We now claim that $\mu = \partial X$ satisfies the conclusion of the claim. This is clear if X is not a hole. If X is a hole then, as σ is the last slice preceding D along H, we deduce that $d_X(\mathbb{D}, D)$ is also large. We are done.

We immediately deduce:

53

Clean up.

Claim 20.3. There is a disk $E \in \mathcal{D}(V)$ so that E and μ meet at most twice and so that $d_S(E, \rho_V(\gamma)) \leq R_1$.

Also, by Lemma 3.15 we have:

Claim 20.4. The multi-curve μ and the slice σ have bounded intersection.

Let $R_3 = R_3(S)$ be the given bound on intersection number. We may now describe the algorithm:

Algorithm 20.5. We are given $\gamma \in \mathcal{C}(S)$ and $\mathbb{D} \subset \mathcal{D}(V)$. Build H a hierarchy between γ and \mathbb{D} ; see Leasure [Lea]. Find a resolution into slices $\{\sigma^i\}$ of H. For each slice σ^i list all multicurves μ so that $\iota(\mu, \sigma^i) \leq R_3$.

For every multicurve μ so produced determine whether

- a component of μ bounds a disk in V (using the solution of the word problem in the free group $\pi_1(V)$),
- a component of $S \searrow \mu$ compresses in V (using Haken's algorithm), or
- a component of $S \setminus \mu$ cobounds an *I*-bundle in *V* (using the Jaco-Tollefson algorithm for detecting the JSJ decomposition).

For every μ falling into one of the three categories above the algorithm also produces a disk E_{μ} with $\iota(E_{\mu}, \mu) \leq 2$. For every disk E_{μ} so produced compute $d_S(\gamma, E_{\mu})$ (again, see [Lea]). Finally take E equal to any of the E_{μ} which minimized $d_S(\gamma, E_{\mu})$, among all disks considered.

By Claim 20.3 the disk E satisfies $d_S(E, \rho_V(\gamma)) \leq R_1$. This proves: **Theorem 20.6.** There is a constant $R_1 = R_1(V)$ and an algorithm which, given an essential curve $\gamma \subset S$ and a pants decomposition $\mathbb{D} \subset \mathcal{D}(V)$, finds a disk $E \in \mathcal{D}(V)$ so that

$$d_S(E, \rho_V(\gamma)) \le R_1.$$

Remark 20.7. We obviously make no claims as to the efficiency of this algorithm.

Theorem 20.8. There is a constant $R_2 = R_2(S)$ and an algorithm which, given a Heegaard diagram $(S, \mathbb{D}, \mathbb{E})$, computes a number N so that

$$|d_S(V,W) - N| \le R_2.$$

Proof. Using Theorem 20.6 we find a disk D close to $\pi_V(\mathbb{E})$. Similarly, we find a disk E close to $\pi_W(D)$.

Recall that computing distance between fixed vertices in the curve complex is algorithmic [Lea]. So, we may compute the distance $d_S(D, E)$. By the quasi-convexity of $\mathcal{D}(V)$ and of $\mathcal{D}(W)$ (see [MM]) this is the desired number.

APPENDIX A. TEICHMÜLLER GEODESICS DO NOT BACKTRACK

KASRA RAFI

needs to be fixed

Let $G: \mathbb{R} \to \mathcal{T}(S)$ be a Teichmüller geodesic, and let ν_{\pm} be the corresponding vertical and horizontal foliations. Let Y be a subsurface of S that is not an annulus or a pair of pants. To be more precise, for $t \in \mathbb{R}$, let μ_t be a shortest marking of S in G(t). We call the map $\rho_Y: \mathbb{R} \to \mathcal{C}(Y)$ a shadow of G to $\mathcal{C}(Y)$ if $\rho_Y(t)$ is a curve in $\pi_Y(\mu_t)$. In these notes, we prove the following theorem.

Theorem A.1. Every shadow of a Teichmüller geodesic G to C(Y) is an un-parametrized quasi-geodesic in C(Y).

We remark that this contrasts with the way geodesics behave in the Lipschitz metric on $\mathcal{T}(S)$, studied by Thurston in [Thu98], where the projection of a geodesic to a subsurface can be move back and forth for arbitrarily long periods. (Examples can easily be produced using Thurston's construction of minimal stretch maps [Thu98] and the results in [CR05]).

To begin, note that if Y = S, the above is a theorem of Masur and Minsky [MM00, Theorem 3.3], that is, we already know that the shadow of G to $\mathcal{C}(S)$ does not backtrack. To prove Theorem A.1 in general, in addition to the Masur-Minsky result we require the following theorem: Let $FN_Y: \mathcal{T}(S) \to \mathcal{T}(Y)$ be the map defined by restriction of Fenchel-Nielsen coordinates to Y; that is, choose a pants decomposition of S that contains the boundary of Y. The forgetful map sends Fenchel-Nielsen coordinates on $\mathcal{T}(S)$, with respect to this pants decomposition, to Fenchel-Nielsen coordinates on $\mathcal{T}(Y)$ (see [Min96]).

Theorem A.2. For every Teichmüller geodesic G and every subsurface Y, there exists an interval $J_Y = [a, b]$ and a geodesic $G_Y \colon J_Y \to \mathcal{T}(Y)$ such that

• If t < a (resp. b < t), then

 $d_Y(\mu_t, \mu_a) = O(1)$ (resp. $d_Y(\mu_b, \mu_t) = O(1)$).

• For $t \in J_Y$,

$$d_{\mathcal{T}}\big((\mathrm{FN}_Y \circ G)(t), G_Y(t)\big) = O(1).$$

From the result of Masur and Minsky, we know that any shadow of G_Y is an un-parametrized quasi-geodesic. The first conclusion of the above theorem implies that the image a shadow of $G|(-\infty, a]$ and $G|[b, \infty)$ is a bounded set in $\mathcal{C}(Y)$ and its second conclusion implies that any shadow of G|[a, b] fellow travels any shadow of G_Y . That is, Theorem A.1 follows from Theorem A.2.



FIGURE 11. (FN_Y \circ G) and G_Y are bounded away in $\mathcal{T}(Y)$

Theorem A.2 is essentially follows from results contained in [Min96], [Raf05], [Raf06] and [Raf07]. We sketch the arguments below.

Let q_t be the quadratic differential given by the geodesic G at time t. We call the flat metric $|q_t|$ the q_t -metric. For a curve α in S, let $q_t(\alpha)$ be the q_t -length of α : the length of a geodesic representative of α in the q_t -metric. Let α be a boundary component of Y, and let ω be an arc in Y with both endpoints in α . By the q_t -length of ω , we mean the length of the shortest arc representing ω that starts and ends on a geodesic representative of α – all considered in the q_t -metric. Again, denote this length by $q_t(\omega)$. Define

$$M_t(\alpha, Y) = \min_{\omega} \frac{q_t(\omega)}{q_t(\alpha)},$$

where ω ranges over all arcs in Y with both endpoints on α , as above. Let t_{α} be the time when α is balanced (its intersections with the horizontal and the vertical foliations are equal). The following lemma is contained in the proof of Theorem 3.1 of [Raf06].

Lemma A.3. There is a uniform constant $c \ge 0$ so that

$$M_s(\alpha, Y) \le M_t(\alpha, Y) + c$$

for all $s \leq t \leq t_{\alpha}$ and for all $t_{\alpha} \leq t \leq s$.

Choose a large enough M_0 (see Theorem A.4 below). Define $J_{\alpha,Y} \subset \mathbb{R}$ to be empty when $M_{t_{\alpha}}(\alpha, Y) < M_0$. On the other hand, if $M_{t_{\alpha}}(\alpha, Y) \geq M_0$, then take $J_{\alpha,Y}$ to be the largest interval containing t_{α} so that $M_t(\alpha, Y) \geq M_0$ for all $t \in J_{\alpha,Y}$. We now define

$$J_Y = \bigcap_{\alpha \subset \partial Y} J_{\alpha, Y}.$$

Note that, by Lemma A.3, for every $t \notin J_Y$, there is a boundary component α such that $M_t(\alpha, Y) \leq M_0 + c$. Immediately from the proof of Theorem 5.5 of [Raf05] we now have:

Theorem A.4. There are constants M_0 and C so that, if $M_t(\alpha, Y) \leq M_0 + c$ for some boundary component α , then either

$$d_Y(\mu_t, \nu_-) \le C \quad or \quad d_Y(\mu_t, \nu_+) \le C.$$

This proves the first conclusion of Theorem A.2. In fact, for t < a,

 $d_Y(\mu_t, \nu_-) \leq C$ and $d_Y(\mu_a, \nu_-) \leq C \Longrightarrow d_Y(\mu_t, \mu_a) \leq 2C$,

and, for t > b,

 $d_Y(\mu_t, \nu_+) \leq C$ and $d_Y(\mu_b, \nu_+) \leq C \Longrightarrow d_Y(\mu_t, \mu_b) \leq 2C.$

To obtain the second conclusion of Theorem A.2, we must first construct the candidate geodesic arc G_Y in $\mathcal{T}(Y)$. Our plan, as in [Raf07], is to fill all components of ∂Y with locally flat punctured disks. The quadratic differential $q_t|Y$ will extend over these disks, giving an extension \overline{q}_t . The map $G_Y: J_Y \to \mathcal{T}(Y)$, where $G_Y(t)$ is the conformal structure of \overline{q}_t , will be a Teichmüller geodesic. We will then use a theorem of Minsky [Min96] to show that the distance in $\mathcal{T}(Y)$ between $\operatorname{FN}_Y(t)$ and $G_Y(t)$ is uniformly bounded, depending only on M_0 . We begin by examining the boundary of Y.

For every curve α , the q_t -geodesic representatives of α form a (possibly degenerate) flat annulus $F_t(\alpha)$. Let \mathbf{Y}_t be a representative of the isotopy class of Y that has q_t -geodesic boundaries and that is disjoint from the interior of $F_t(\alpha)$ for every curve $\alpha \subset \partial Y$. For $t \in J_Y$ and $r = M_t(\alpha, Y)/2$, the r-regular neighborhood of a boundary component of Y inside \mathbf{Y}_t is an annulus $A = A_t(\alpha)$. We call the boundary component of A that is a q_t -geodesic representative of α the *inner boundary*. Along the inner boundary we attach a foliated locally flat punctured disk as follows (the construction is identical to the one given in [Raf07] and is repeated here for completeness):

Let A' be the double cover of A, and let q'_t be the lift of q_t . Let x_1, \ldots, x_n be the points on the inner boundary of A which have angle $\theta_i > \pi$ in Y_t . Note that this set is nonempty: if not, Y_t meets the interior of the flat annulus $F_t(\alpha)$, a contradiction. Denote the lifts of x_i by y_i and z_i . We now fill the inner boundary of A' by symmetrically adding 2(n-1) Euclidean triangles to obtain a disk, D', equipped with a singular flat structure such that the total angle at each point is a multiple of π and is at least 2π .

We start by attaching a Euclidean triangle to vertices y_1, y_2, y_3 , which we denote by $\Delta(y_1, y_2, y_3)$ (see Figure 12). We choose the angle at vertex y_2 , $\angle y_2$, so that the total angle at y_2 , $\theta_2 + \angle y_2$, is a multiple of π . Assuming $0 \leq \angle y_2 < \pi$, there is a unique such triangle. Attach an isometric triangle to z_1, z_2, z_3 . Now consider points y_1, y_3, y_4 . Again, there



FIGURE 12. The filling of the annulus A'

exists a Euclidean triangle with one edge equal to the newly introduced segment $[y_1, y_3]$, another edge equal to the segment $[y_3, y_4]$ and an angle at y_3 that makes the total angle at y_3 , including the contribution from the triangle $\triangle(y_1, y_2, y_3)$, a multiple of π . Attach this triangle to vertices y_1, y_3, y_4 and an identical triangle to vertices z_1, z_3, z_4 . Continue in this fashion until finally adding triangles $\triangle(y_1, y_n, z_1)$ and $\triangle(z_1, z_n, y_1)$. Because of the symmetry, the two edges connecting y_1 and z_1 have equal length, and we can glue these together. We call the union of the added triangles D'. Notice that the involution on A' extends to D'. Let Dbe the quotient of D', and note that D is a punctured disk attached to the inner boundary of A.

For $i \neq 1$, the total angles at y_i and at z_i are multiples of π and are larger than $\theta_i > \pi$; therefore, they are at least 2π . We have added 2(n-1) triangles. Hence, the sum of the total angles of all vertices is $2\sum_i \theta_i + 2(n-1)\pi$, which is a multiple of 2π . Therefore, the sum of the angles at y_1 and z_1 is also a multiple of 2π . But they are equal to each other, and each one is larger than π . This implies that they are both at least 2π .

It follows that the quadratic differential q'_t extends over D' symmetrically with quotient \overline{q}_t , an extension of q_t to D. Then \overline{q}_t is a quadratic differential on Y. Denote the Riemann surface obtained after capping off all boundaries of \mathbf{Y}_t by \overline{Y}_t . This completes the construction.

It remains to show that the distance in $\mathcal{T}(Y)$ between $\operatorname{FN}_Y(G((t)))$ and $G_Y(t)$ is uniformly bounded. For this, we examine the extremal lengths of curves in two metrics. We say two quantities P and Qare *comparable* if they are the same up to up a multiplicative error depending on the topology of S only and we write $P \approx Q$. We know from [Min96] that, for any essential curve γ in Y, the extremal lengths of γ in G(t) and in $FN_Y(G(t))$ are comparable:

(A.5)
$$\operatorname{Ext}(\gamma, G(t)) \stackrel{:}{\asymp} \operatorname{Ext}(\gamma, \operatorname{FN}_Y(G(t))),$$

(See the proof of Theorem 6.1 in [Min96], page 283, line 19.) We need to show that the extremal lengths of γ in G(t) and in $G_Y(t)$ are comparable as well. For this, we need the following lemma (see [Min96, Lemma 4.2]).

Lemma A.6 (Minsky). Let Σ be any Riemann surface. There exists a constant m, depending only on the topological type of Σ , such that if $Y \subset \Sigma$ is a subsurface of S with negative Euler characteristic for which each component α of ∂Y bounds an annulus A_{α} in Y, with modulus $\operatorname{Mod}(A_{\alpha}) \geq m$, then for any essential γ in Y,

$$\operatorname{Ext}(\gamma, Y) \stackrel{:}{\asymp} \operatorname{Ext}(\gamma, \Sigma).$$

After choosing M_0 large enough, Y_t satisfies the condition of the above lemma, for every $t \in J_Y$. In fact, for $\alpha \subset \partial Y$ and $r = q_t(\alpha) M_t(\alpha, Y)$, the *r*-neighborhood of α in Y_t is an annulus and has a modulus log $M_t(\alpha, Y)$ [Raf05, Lemma 3.6]. Choosing M_0 large enough, we can assume that, for $t \in J_Y$, log $M_t(\alpha, Y) \ge m_0$.

Now, considering Y_t as a subsurface of G(t) and as a subsurface of $\overline{Y}_t = G_Y(t)$, Lemma A.6 implies that

$$\operatorname{Ext}(\gamma, G(t)) \stackrel{:}{\asymp} \operatorname{Ext}(\gamma, \mathsf{Y}_t) \stackrel{:}{\asymp} \operatorname{Ext}(\gamma, G_Y(t)).$$

This and (A.5) finishes the proof of Theorem A.2.

References

- [BDM] Jason Behrstock, Cornelia Drutu, and Lee Mosher. Thick metric spaces, relative hyperbolicity, and quasi-isometric rigidity. arXiv:math.GT/0512592.
- [Beh04] 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.
- [BF] Jeffrey Brock and Benson Farb. Curvature and rank of Teichmüller space. arXiv:math.GT/0109045.
- [BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.
- [BM04] Tara E. Brendle and Dan Margalit. Commensurations of the Johnson kernel. *Geom. Topol.*, 8:1361–1384 (electronic), 2004. arXiv:math.GT/0404445.
- [Bow91] B. H. Bowditch. Notes on Gromov's hyperbolicity criterion for pathmetric spaces. In Group theory from a geometrical viewpoint (Trieste, 1990), pages 64–167. World Sci. Publ., River Edge, NJ, 1991.

- [Bro03] 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.GT/0109048.
- [CDP90] M. Coornaert, T. Delzant, and A. Papadopoulos. Géométrie et théorie des groupes. Springer-Verlag, Berlin, 1990. Les groupes hyperboliques de Gromov.
- [CR05] Young-Eun Choi and Kasra Rafi. Comparison Between Teichmuller and Lipschitz Metrics, 2005. math.GT/0510136.
- [Gro87] Mikhael Gromov. Hyperbolic groups. In Essays in group theory, pages 75–263. Springer, New York, 1987.
- [Har81] 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.
- [Hem01] John Hempel. 3-manifolds as viewed from the curve complex. *Topology*, 40(3):631–657, 2001. arXiv:math.GT/9712220.
- [HT80] A. Hatcher and W. Thurston. A presentation for the mapping class group of a closed orientable surface. *Topology*, 19(3):221–237, 1980.
- [Kob88] 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.
- [Lea] Jason Leasure. Geodesics in the complex of curves of a surface. Ph.D. thesis.
- [McC91] Darryl McCullough. Virtually geometrically finite mapping class groups of 3-manifolds. J. Differential Geom., 33(1):1–65, 1991.
- [Min96] Yair N. Minsky. Extremal length estimates and product regions in Teichmüller space. *Duke Math. J.*, 83(2):249–286, 1996. arXiv:math.GT/9407222.
- [MM] Howard A. Masur and Yair N. Minsky. Quasiconvexity in the curve complex. arXiv:math.GT/0307083.
- [MM99] Howard A. Masur and Yair N. Minsky. Geometry of the complex of curves. I. Hyperbolicity. Invent. Math., 138(1):103–149, 1999. arXiv:math.GT/9804098.
- [MM00] 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.GT/9807150.
- [Mos] Lee Mosher. Train track expansions of measured foliations. http://newark.rutgers.edu/~mosher/.
- [Pen88] Robert C. Penner. A construction of pseudo-Anosov homeomorphisms. Trans. Amer. Math. Soc., 310(1):179–197, 1988.
- [Raf05] Kasra Rafi. A characterization of short curves of a Teichmüller geodesic. Geom. Topol., 9:179–202 (electronic), 2005. arXiv:math.GT/0404227.
- [Raf06] Kasra Rafi. A combinatorial model for the Teichmüller metric. To appear in Geometric and Functional Analysis, 2006. arXiv:math.GT/0509584.
- [Raf07] Kasra Rafi. Thick-thin decomposition for quadratic differ-MathematicsResearchentials. То appear $_{in}$ Letters, 2007.http://www.math.uconn.edu/~rafi/research/research.html/TT.pdf.

THE GEOMETRY OF THE DISK COMPLEX

- [Sch] Saul Schleimer. Notes on the curve complex. http://math.rutgers.edu/~saulsch/Maths/notes.pdf.
- [Sch82] Martin Scharlemann. The complex of curves on nonorientable surfaces. J. London Math. Soc. (2), 25(1):171–184, 1982.
- [Thu88] William P. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. Bull. Amer. Math. Soc. (N.S.), 19(2):417–431, 1988.
- [Thu98] William P. Thurston. Minimal stretch maps between hyperbolic surfaces, 1998. arXiv:math.GT/9801039.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, CHICAGO, ILLINOIS 60607

E-mail address: masur@math.uic.edu *URL*: http://www.math.uic.edu/~masur

Department of Mathematics, Rutgers University, Piscataway, New Jersey 08854

E-mail address: saulsch@math.rutgers.edu *URL*: http://www.math.rutgers.edu/~saulsch