

Subsurface distances for hyperbolic 3-manifolds fibering over the circle

Yair N. Minsky and Samuel J. Taylor

Abstract. For a hyperbolic fibered 3-manifold M , we prove results that uniformly relate the structure of surface projections as one varies the fibrations of M . This extends our previous work from the fully-punctured to the general case.

1. Introduction

Let M be a hyperbolic 3-manifold, and let S be a fiber in a fibration of M over the circle. The corresponding monodromy is a pseudo-Anosov homeomorphism $\varphi: S \rightarrow S$ and comes equipped with invariant stable and unstable laminations $\lambda^\pm$ on S . Let $N_S \rightarrow M$ be the infinite cyclic covering of M corresponding to S .

As a consequence of the proof of Thurston’s ending lamination conjecture, Minsky [19] and Brock–Canary–Minsky [10] develop combinatorial tools to study the geometry of a hyperbolic manifold homeomorphic to $S \times \mathbb{R}$. Applying their work to the special case of $N_S \cong S \times \mathbb{R}$ explains how the geometry of N_S , and hence that of M , is *coarsely* determined by combinatorial data associated to the pair of laminations $\lambda^\pm$. In particular, using only the pair $\lambda^\pm$, a combinatorial model of N_S (called the *model manifold*) is constructed and this model is shown to be bi-Lipschitz to N_S , where the bi-Lipschitz constant depends only on the complexity of the surface S . For this, one of the main combinatorial tools are the Masur–Minsky subsurface projections [17], which associate to each subsurface $Y \subset S$ a *subsurface projection distance* $d_Y(\lambda^-, \lambda^+)$ measuring the complexity of $\lambda^\pm$ as seen from Y . In fact, subsurface projections have proven to be useful in several settings [6, 16, 22] and have been generalized in many directions [5, 7, 8, 25].

These developments suggest the following outline for studying the geometry of a hyperbolic fibered 3-manifold M : apply the model manifold machinery to an infinite cyclic cover of M associated to a fiber and use this to make conclusions about the structure of M . In fact, because of Agol’s resolution of the virtual fibering conjecture [3], this simple idea generalizes to any hyperbolic manifold by first passing to a finite sheeted cover which fibers over the circle.

Unfortunately, this approach is too naïve for a number of reasons, perhaps the most important of which is that the complexity of a fiber in the appropriate cover is *not* known at the outset. Indeed, even when a fibered manifold M is fixed, if $\dim(H^1(M; \mathbb{R})) \geq 2$, then M fibers in infinitely many ways, and the complexities of the corresponding fibers are necessarily unbounded. Since the bi-Lipschitz constants in the model manifold theorem depend on the complexity of the underlying surface, this approach goes nowhere without a precise understanding of how the constants relating geometry to combinatorics vary as the surface changes.

To salvage this approach, one would like control over how the tools at the center of the construction depend on complexity. The purpose of this paper is to give such uniform, explicit control on subsurface projection distances as one varies the fibers within a fixed fibered manifold. This extends our previous work [20] that handled the special case of *fully-punctured* fibered manifolds (see below for details).

Main results

Recall that the fibrations of a manifold M are organized into finitely many “fibered faces” of the unit ball in $H^1(M, \mathbb{R})$ of the Thurston norm [26], where each fibered face $\mathbf{F}$ has the property that all primitive integral classes in the open cone $\mathbb{R}_+\mathbf{F}$ represent a fiber (see Section 2.3). Associated to each fibered face is a pseudo-Anosov flow which is transverse to every fiber represented in $\mathbb{R}_+\mathbf{F}$ [13].

Our first main result bounds the size and projection distance for all subsurfaces of all fibers over a fixed fibered face $\mathbf{F}$. The constant D in the statement is no more than 15 (see Section 2.1.2) and $|\chi'(Y)| = \max\{|\chi(Y)|, 1\}$.

Theorem 1.1 (Bounding projections for M). *Let M be a hyperbolic fibered 3-manifold with fibered face $\mathbf{F}$. Then for any fiber S contained in $\mathbb{R}_+\mathbf{F}$ and any subsurface Y of S ,*

$$|\chi'(Y)| \cdot (d_Y(\lambda^-, \lambda^+) - 16D) \leq 2D|\mathbf{F}|,$$

where $|\mathbf{F}|$ is a constant depending only on $\mathbf{F}$.

In particular, subsurface projections are uniformly bounded over the fibered face $\mathbf{F}$ as are the complexities of subsurfaces whose projection distances are greater than $16D$. Note that since M has only finitely many fibered faces, this bounds the size of all subsurface projections among all fibers of M .

Second, we relate subsurfaces of different fibers in the same fibered face of M . Note that here the constants involved do not depend on the manifold M .

Theorem 1.2 (Subsurface dichotomy). *Let M be a hyperbolic fibered 3-manifold, and let S and F be fibers of M which are contained in the same fibered cone. If $W \subset F$ is a subsurface of F , then either W is homotopic, through surfaces transverse to the associated flow, to an embedded subsurface W' of S with*

$$d_{W'}(\lambda^-, \lambda^+) = d_W(\lambda^-, \lambda^+)$$

or the fiber S satisfies

$$9D \cdot |\chi(S)| \geq d_W(\lambda^-, \lambda^+) - 16D.$$

Along the way to establishing our main theorems, we prove several results that may be independently interesting. While most of these concern the connection between the manifold M and the veering triangulation of the associated fully-punctured manifold (see the next section for details), we also obtain information about subsurface projections to *immersed* subsurfaces.

For a finitely generated subgroup $\Gamma < \pi_1(S)$, let $S_\Gamma \rightarrow S$ be the corresponding cover. If $Y \subset S_\Gamma$ is a compact core, the covering map restricted to Y is an immersion and we say that $Y \rightarrow S$ *corresponds to* $\Gamma < \pi_1(S)$. Lifting to the cover induces a (partially defined) map of curve and arc graphs which we denote by $\pi_Y: \mathcal{A}(S) \rightarrow \mathcal{A}(Y)$. (When Γ is cyclic, we set $\mathcal{A}(Y)$ to be the annular curve graph $\mathcal{A}(S_\Gamma)$ as usual.) Note that these constructions depend on Γ but not on the choice of Y and that π_Y agrees with the usual subsurface projection when $Y \subset S$ is an embedded subsurface.

Theorem 1.3 (Immersion to covers). *There is a constant $M \leq 38$ satisfying the following: Let S be a surface, and let $Y \rightarrow S$ be an immersion corresponding to a finitely generated $\Gamma < \pi_1(S)$. Then either*

- *there is a subsurface $W \subset S$ so that $Y \rightarrow S$ factors (up to homotopy) through a finite sheeted covering $Y \rightarrow W$, or*
- *the diameter of the entire projection of $\mathcal{A}(S)$ to $\mathcal{A}(Y)$ is bounded by M .*

The novelty of Theorem 1.3 is that the constant $M \leq 38$ is explicit and uniform over all surfaces and immersions. Previously, Rafi and Schleimer proved that for any finite cover $\tilde{S} \rightarrow S$, there is a constant $T \geq 0$ (depending on $\tilde{S}$ and S) such that if $Y \subset \tilde{S}$ is a subsurface with $d_Y(\alpha, \beta) \geq T$ for $\alpha, \beta \in \mathcal{A}(S)$, then Y covers a subsurface of S [23, Lemma 7.2].

Relation to our previous work

Given a fibered face $\mathbf{F}$ of M and its associated pseudo-Anosov flow, the stable/unstable laminations $\Lambda^\pm$ of the flow intersect each fiber to give the laminations associated to its monodromy. Removing the singular orbits of the flow produces the *fully-punctured* manifold $\mathring{M}$ associated to the face $\mathbf{F}$. If $\varphi: S \rightarrow S$ is the monodromy of some fiber S representing a class in $\mathbb{R}_+ \mathbf{F}$, then $\mathring{M}$ is the mapping torus of the surface $\mathring{S}$ obtained from S by puncturing at the singularities of φ . The fibered face of $\mathring{M}$ containing $\mathring{S}$ is denoted by $\mathring{\mathbf{F}}$, and the inclusion $\mathring{M} \subset M$ induces an injective homomorphism $H^1(M; \mathbb{R}) \hookrightarrow H^1(\mathring{M}; \mathbb{R})$ mapping $\mathbb{R}_+ \mathbf{F}$ into $\mathbb{R}_+ \mathring{\mathbf{F}}$.

In our previous work [20], we restricted our study of subsurface projections in fibered manifolds to the fully-punctured settings. When $\mathring{M}$ is fully-punctured, it admits a canonical *veering triangulation* τ [1, 14] associated to the fibered face $\mathring{\mathbf{F}}$. We found that the combinatorial structure of this triangulation encodes the hierarchy of subsurface projections for each fiber $\mathring{F}$ in $\mathbb{R}_+ \mathring{\mathbf{F}}$. As a result, we established versions of Theorems 1.1 and 1.2

in that restricted setting (though with better constants than available in general). In fact, when the fibered manifold is fully-punctured there are additional surprising connections between the veering triangulation and the curve graph. For example, a fiber $\mathring{F}$ of $\mathring{M}$ is necessarily a punctured surface, and edges of the triangulation τ (when lifted to the cover of $\mathring{M}$ corresponding to $\mathring{F}$) form a subset of the arc graph $\mathcal{A}(\mathring{F})$. This subset is *geodesically connected* in the sense that for any pair of arcs of $\mathring{F}$ coming from edges of τ , there is a geodesic in $\mathcal{A}(\mathring{F})$ joining them consisting entirely of veering edges [20, Theorem 1.4]. Such a result cannot have a precise analog if, for example, the manifold M is closed.

In this paper, we extend our study to general (e.g., closed) hyperbolic fibered manifolds. The main difficulty here is that these manifolds do not admit veering triangulations. So our approach is to start with an arbitrary fibered manifold M and consider the veering triangulation of the associated fully-punctured manifold $\mathring{M}$. (For example, the constant $|\mathbf{F}|$ appearing in Theorem 1.1 is precisely the number of tetrahedra of the veering triangulation of $\mathring{M}$ associated to $\mathring{\mathbf{F}}$.) Unfortunately, results about subsurface projections to fibers of M do not directly imply the corresponding statements in $\mathring{M}$. Instead, we develop tools to relate *sections* of the veering triangulation (i.e., ideal triangulations of the fully-punctured fiber by edges of the veering triangulation) to subsurface projections in the fibers of M .

Summary of paper

In Section 2, we present background material. In particular, we summarize the definition of the veering triangulation (Section 2.2) and recall the main constructions from [20] that connect the structure of the veering triangulation on $\mathring{S} \times \mathbb{R}$ to subsurface representatives in $\mathring{S}$ (Section 2.4).

Section 3 introduces the lattice structure of *sections* of the veering triangulation. It concludes with Section 3.2 which details how sections (which are ideal triangulations of the fully-punctured surface $\mathring{S}$) are used to define projections to the curve graph of subsurfaces of the original surface S . This is followed by Section 4 where Theorem 1.3 is proven. This section does not use veering triangulations and can be read independently from the rest of the paper.

In Section 5, we prove two estimates that relate the veering triangulation of the fully-punctured manifold $\mathring{M}$ to fibers of M . The first (Proposition 5.1) shows that for each subsurface Y of S , there are top and bottom sections of $\mathring{S} \times \mathbb{R}$ which project close to the images of $\lambda^\pm$ in $\mathcal{A}(Y)$. The second (Lemma 5.2) shows that these projections to $\mathcal{A}(Y)$ move slowly from $\pi_Y(\lambda^-)$ to $\pi_Y(\lambda^+)$ depending on the size of Y . Both of these estimates are needed for proofs of Theorems 1.1 and 1.2.

Finally, in Sections 6 and 7, Theorems 1.1 and 1.2 are proven. The bulk of the proof of Theorem 1.1 involves building a simplicial *pocket* for the subsurface Y that embeds into the veering triangulation of $\mathring{M}$ whose “width” is approximately $|\chi(Y)|$ and whose “depth” is at least $d_Y(\lambda^-, \lambda^+)$. For Theorem 1.2, we show that if a subsurface W of a fiber F is not homotopic into another fiber S (in the same fibered cone) then, after puncturing along singular orbits of the flow, a section (triangulation) of $\mathring{S}$ contains many edges, proportional

to the “depth” of the pocket for W . For each of these arguments, the difficulty lies in the fact that we are extracting information about the original manifold M and projections to subsurfaces of its fibers by relying on the veering triangulation of $\mathring{M}$, which a priori only records information about projections to its fully-punctured fibers.

2. Background

Here we record basic background that we will need throughout the paper. We begin with some material that follows easily from standard facts about curve graphs and then recall the definition of the veering triangulation. We conclude by reviewing results from [20] which develop connections between the two.

Remark 2.1 (Pseudo-Anosov conventions). For a pseudo-Anosov homeomorphism $\varphi: S \rightarrow S$, we denote its stable and unstable lamination by λ^+ and λ^- , respectively. Here we use the dynamical convention that the leaves of λ^+ are contracted by φ . Hence, in natural coordinates away from singularities for the quadratic differential q whose vertical/horizontal foliations are determined by λ^+/λ^- , respectively, φ has the form

$$\begin{pmatrix} k & 0 \\ 0 & k^{-1} \end{pmatrix},$$

where $k > 1$ is the stretch factor of φ .

The choice to denote the stable lamination by λ^+ comes from the 3-dimensional perspective. The mapping torus of φ is defined by

$$M_\varphi = \frac{S \times \mathbb{R}}{(x, t) \sim (\varphi(x), t - 1)},$$

so that the first return map of the positive flow in the $\mathbb{R}$ direction is φ . This means that the deck translation on the cover $S \times \mathbb{R}$ in the positive $\mathbb{R}$ direction acts by φ^{-1} on S , so that the positive deck translations of any fixed curve in S converge in $\mathcal{PML}$ to λ^+ . This is consistent with the fact from hyperbolic geometry that if one considers the hyperbolic structure on $S \times \mathbb{R}$ that covers the unique hyperbolic structure on M_φ , then the ending lamination for the positive end of $S \times \mathbb{R}$ is λ^+ .

Note that with this convention, we have the slightly counterintuitive fact that the unstable lamination λ^- is the attracting point for φ with respect to its action on the space of projective measured laminations $\mathcal{PML}$.

2.1. Curve graph facts and computations

The arc and curve graph $\mathcal{A}(Y)$ for a compact surface Y is the graph whose vertices are homotopy classes of essential simple closed curves and proper arcs. Edges join vertices precisely when the vertices have disjoint representatives on Y . Here, essential curves/arcs are those which are not homotopic (rel endpoints) to a point or into the boundary.

If Y is not an annulus, homotopies of arcs are assumed to be homotopies through maps sending boundary to boundary. This is equivalent to considering proper embeddings $\mathbb{R} \rightarrow \text{int}(Y)$ into the interior of Y up to proper homotopy, and we often make use of this perspective. When Y is an annulus, the homotopies are also required to fix the endpoints. We consider $\mathcal{A}(Y)$ as a metric space by using the graph metric, although we usually only consider distance between vertices. For additional background, the reader is referred to [17, 19].

If $Y \subset S$ is an essential subsurface (i.e., one that is π_1 -injective and contains an essential curve), we have subsurface projections $\pi_Y(\lambda)$ which are defined for simplices $\lambda \subset \mathcal{A}(S)$ that intersect Y essentially, otherwise the projection is defined to be empty. Namely, after lifting λ to the cover S_Y associated to $\pi_1(Y)$, we obtain a collection of properly embedded disjoint essential arcs and curves, which determine a simplex of $\mathcal{A}(Y) := \mathcal{A}(S_Y)$. We let $\pi_Y(\lambda)$ be the union of these vertices. The same definition applies to a lamination λ that intersects ∂Y essentially.

When Y is an annulus, these arcs have natural endpoints coming from the standard compactification of $\tilde{S} = \mathbb{H}^2$ by a circle at infinity. We remark that π_Y does not depend on any choice of hyperbolic metric on S .

When Y is not an annulus and λ and ∂Y are in minimal position, we can also identify $\pi_Y(\lambda)$ with the isotopy classes of components of $\lambda \cap Y$.

When λ, λ' are two arc/curve systems or laminations, we denote by $d_Y(\lambda, \lambda')$ the diameter of the union of their images in $\mathcal{A}(Y)$, that is,

$$d_Y(\lambda, \lambda') = \text{diam}_{\mathcal{A}(Y)}(\pi_Y(\lambda) \cup \pi_Y(\lambda')).$$

2.1.1. An ordering on subsurface translates. Here we prove a lemma that will be needed in Section 6. It establishes an ordering on translates of a fixed subsurface under a pseudo-Anosov map by appealing to a more general ordering of Behrstock–Kleiner–Minsky–Mosher [6], as refined in Clay–Leininger–Mangahas [11].

Fix a pseudo-Anosov $\varphi: S \rightarrow S$ with stable and unstable laminations λ^+ and λ^- , respectively. Recall from Remark 2.1 that λ^- is the unstable lamination of φ and hence its attracting fixed point on $\mathcal{PM}\mathcal{L}$.

Lemma 2.2. *If $d_Y(\lambda^-, \lambda^+) \geq 20$, then for any $n \geq 1$ with $\pi_Y(\varphi^n(\partial Y)) \neq \emptyset$,*

$$d_Y(\varphi^n(\partial Y), \lambda^-) \leq 4.$$

Proof. First consider the set of subsurfaces $\mathcal{S} = \{Y: d_Y(\lambda^-, \lambda^+) \geq 20\}$. If Y, Z are members of $\mathcal{S}$ that overlap nontrivially then, following [11], we say $Y < Z$ if

$$d_Y(\partial Z, \lambda^+) \leq 4.$$

According to [11, Proposition 3.6], this is equivalent to the condition that $d_Z(\partial Y, \lambda^-) \leq 4$, and any two overlapping $Y, Z \in \mathcal{S}$ are ordered. Moreover, by [11, Corollary 3.7], $<$ is a strict partial order on $\mathcal{S}$.

Returning to our setting, suppose that $Y \in \mathcal{S}$ and that $\varphi^n(Y)$ and Y overlap for some $n \geq 1$. Consider the sequence $Y_i = \varphi^{in}(Y)$ and note that $Y_0 = Y$.

Now, we know that $\partial Y_i \rightarrow \lambda^-$ in $\mathcal{PM}\mathcal{L}$ as $i \rightarrow +\infty$. This implies that for large enough i , we have $d_Y(\partial Y_i, \lambda^-) \leq 1$, and hence $Y_i \prec Y_0$.

On the other hand, if $Y_0 \prec Y_1$ then, since φ preserves $\lambda^\pm$, we have $Y_i \prec Y_{i+1}$ for all i . Since $\prec$ is transitive, this would imply that $Y_0 \prec Y_i$, a contradiction.

Since Y_0 and Y_1 are ordered, we must have $Y_1 \prec Y_0$. Hence, $d_Y(\varphi^n(\partial Y), \lambda^-) \leq 4$, which is what we wanted to prove. ■

2.1.2. Distance and intersection number. For an orientable surface S with genus $g \geq 0$ and $p \geq 0$ punctures, set $\zeta = \zeta(S) = 2g + p - 4 = |\chi(S)| - 2$. The following lemma of Bowditch will be important in making our estimates uniform over complexity. Asymptotically stronger, yet less explicit, bounds were first proven by Aougab [4].

Lemma 2.3 (Bowditch [9]). *For any integer $n \geq 0$ and curves $\alpha, \beta \in \mathcal{A}(S)^{(0)}$ with $\zeta = \zeta(S)$, we have*

$$2^n \cdot i(\alpha, \beta) \leq \zeta^{n+1} \Rightarrow d_S(\alpha, \beta) \leq 2(n+1).$$

If the surface S is punctured, then for any arcs a and b in $\mathcal{A}(S)^{(0)}$ there are a curve $\alpha \in \mathcal{A}(S)^{(0)}$ disjoint from a and a curve $\beta \in \mathcal{A}(S)^{(0)}$ disjoint from b such that $i(\alpha, \beta) \leq 4 \cdot i(a, b) + 4$. These curves are constructed using the standard projection from the arc graph to the curve graph: α is a boundary component of a neighborhood of $a \cup P_1 \cup \dots \cup P_k$, where P_i is a loop circling the i th puncture of S .

Applying Lemma 2.3 together with the above observation, we compute that for curves/arcs $a, b \in \mathcal{A}(S)^{(0)}$ and $\zeta(S) \geq 3$,

$$d_S(a, b) < 6 + 2 \cdot \frac{\log(2i(a, b) + 1)}{\log(\frac{\zeta}{2})},$$

where $\log = \log_2$. We also recall the standard complexity independent inequality (see, e.g., [15, 24])

$$d_S(a, b) \leq 2 \log(i(a, b)) + 2,$$

where $i(a, b) > 0$. Using these inequalities, straightforward computations (which we omit) show that for any curves/arcs $a, b \in \mathcal{A}(S)^{(0)}$,

$$i(a, b) \leq 8|\chi(S)| + 4 \Rightarrow d_S(a, b) \leq 15, \quad (2.1)$$

and

$$i(a, b) \leq 32|\chi(S)| + 8 \Rightarrow d_S(a, b) \leq 18. \quad (2.2)$$

We remark that the above mentioned work of Aougab [4] implies that if $i(a, b) \leq K|\chi(S)|$, then $d_S(a, b) \leq 3$, as long as $|\chi(S)|$ is sufficiently large (depending on K).

2.1.3. Proper graphs. Throughout the paper, we will use curve graph tools to study objects that arise from (partially) ideal triangulations of surfaces. To do this, we introduce the notion of a proper graph.

A *proper graph* in S is a one-complex G minus some subset of the vertices, properly embedded in S . A connected proper graph is *essential* if it is not properly homotopic into an end of S or to a point. In general, G is essential if some component is essential.

A proper arc or curve a is *nearly simple* in a proper graph G if a is properly homotopic to a proper path or curve in G which visits no vertex of G more than twice. Note that a proper graph G in S is essential if and only if it carries an essential arc or curve. Define

$$\mathcal{A}_S(G) = \{a \in \mathcal{A}(S)^{(0)} : a \text{ is nearly simple in } G\}. \quad (2.3)$$

Corollary 2.4. *Suppose that G is an essential proper graph in S with at most $2|\chi(S)| + 1$ vertices. Then $\text{diam}_S(\mathcal{A}_S(G)) \leq D$ for $D = 15$.*

Proof. Let a and b be essential arcs or curves that are nearly simple in G . Realize a and b in a small neighborhood of G so that they intersect only in neighborhoods of the vertices of G . Since a and b each pass through any neighborhood of a vertex at most twice, they intersect at most $4(2|\chi(S)| + 1)$ times. By the computations in equation (2.1), this implies that $d_S(a, b) \leq 15$. ■

2.2. Veering triangulations

Our basic object here is a Riemann surface X with an integrable holomorphic quadratic differential q , which fits into a sequence

$$\overset{\circ}{X} \subset X \subset \bar{X}$$

as follows: $\bar{X}$ is a closed Riemann surface on which q extends to a meromorphic quadratic differential, and $\mathcal{P} = \bar{X} \setminus X$ is a finite set of *punctures* which includes the poles of q , if any. Let $\text{sing}(q)$ be the union of $\mathcal{P}$ with the zeros of q , so that

$$\text{poles}(q) \subseteq \mathcal{P} \subseteq \text{sing}(q)$$

and set $\overset{\circ}{X} = \bar{X} \setminus \text{sing}(q)$. When $X = \overset{\circ}{X}$, we say that X is *fully-punctured*.

Let λ^+ and λ^- be the vertical and horizontal foliations of q , which we assume contain no saddle connections.

The constructions of Agol [1] and Gueritaud [14] yield a fibration

$$\Pi: \mathcal{N} \rightarrow \overset{\circ}{X}$$

whose fibers are oriented lines, so that $\mathcal{N} \cong \overset{\circ}{X} \times \mathbb{R}$, and $\mathcal{N}$ is equipped with an ideal triangulation τ whose tetrahedra, called τ -simplices, are characterized by the following description.

Let $p: \tilde{X} \rightarrow \hat{X}$ be the universal covering map and $\hat{X}$ the metric completion of $\tilde{X}$. The completion points, or *singularities*, of $\hat{X}$ are in bijective correspondence with the preimages of once-punctured disks centered at punctures of $\tilde{X}$. Note that p extends to an infinite branched covering $\hat{X} \rightarrow \bar{X}$. In more detail, let $x \in \hat{X}$ be a puncture with angle $k\pi$. Letting D_1 be the quotient by $z \rightarrow -z$ of a disk in the complex plane, the Euclidean metric descends to a metric with a cone point of angle π at $0 \in D_1$. The pullback of this by a k -fold branched cover of D_1 over 0 gives a model for the metric in $\tilde{X}$ in a neighborhood of x , which we call D_x .

A component of the preimage of $\mathring{D}_x = D_x \setminus \{x\}$ in $\tilde{X}$ can then be identified with its universal cover $\tilde{D}_x$ which is the infinite cyclic cover of $D_1 \setminus \{0\}$. The metric completion $\hat{D}_x$ of $\tilde{D}_x$ is obtained by adding a single point $\hat{x}$, whose neighborhood basis is the set of preimages of open neighborhoods of 0 in D_1 . The covering map extends to an infinite branched cover $\hat{D}_x \rightarrow D_x \subset \hat{X}$ sending $\hat{x}$ to x . Note that in this sense, each completion point of $\hat{X}$ has infinite total angle.

We remark that one can define the completion of the universal cover of X in the same manner, where the singularities are exactly the completion points, which have infinite total angle, along with the singular points of the lifted quadratic differential, which have finite total angle.

A *singularity-free rectangle* in $\hat{X}$ is an embedded rectangle whose edges are leaf segments of the lifts of $\lambda^\pm$ and whose interior contains no singularities of $\hat{X}$. If R is a *maximal* singularity-free rectangle in $\hat{X}$, then it contains exactly one singularity on the interior of each edge. The four singularities span a quadrilateral in R which we can think of as the image of a tetrahedron by a projection map whose fibers are intervals, as pictured in Figure 1.




Figure 1. A maximal singularity-free rectangle R defines a tetrahedron equipped with a map into R .

The tetrahedra of $\mathcal{N}$ are identified with all such tetrahedra, up to the action of $\pi_1(\hat{X})$, where the restriction of Π is exactly this projection to the rectangles, followed by p .

A τ -edge in $\hat{X}$ is the Π -image of an edge of τ , or equivalently a saddle connection of q whose lift to $\tilde{X}$ spans a singularity-free rectangle. A τ -edge in X or $\bar{X}$ is the closure of a τ -edge in $\hat{X}$.

When $\varphi: S \rightarrow S$ is a pseudo-Anosov homeomorphism, let X denote S endowed with a Riemann surface structure and a holomorphic quadratic differential q whose foliations

are the stable and unstable foliations of φ . Then $\lambda^\pm$ have no saddle connections, so we may construct $\mathcal{N}$, on which φ induces a simplicial homeomorphism Φ of $\mathcal{N}$, whose quotient $\mathring{M}_\varphi$ is the mapping torus of $\varphi|_{\mathcal{S}}$. Equivalently, $\mathring{M} = \mathring{M}_\varphi$ is obtained from the mapping torus M_φ by removing the singular orbits of its suspension flow, which we discuss next.

2.3. Fibered faces of the Thurston norm

Let M be a finite-volume hyperbolic 3-manifold. A fibration $\sigma: M \rightarrow S^1$ of M over the circle comes with the following structure: There is a primitive integral cohomology class in $H^1(M; \mathbb{Z})$ represented by $\sigma_*: \pi_1 M \rightarrow \mathbb{Z}$, which is the Poincaré dual of the fiber F . There is also a representation of M as a quotient $(F \times \mathbb{R})/\Phi$, where $\Phi(x, t) = (\varphi(x), t - 1)$ and where $\varphi: F \rightarrow F$ is a pseudo-Anosov homeomorphism called the monodromy map. The map φ has stable and unstable (singular) measured foliations λ^+ and λ^- on F . Finally, there are the suspension flow inherited from the natural $\mathbb{R}$ action on $F \times \mathbb{R}$, and suspensions $\Lambda^\pm$ of $\lambda^\pm$ which are flow-invariant 2-dimensional foliations of M . Note that the deck transformation Φ translates in the *opposite* direction of the lifted flow. This is so that the first return map to the fiber F equals φ .

The fibrations of M are organized by the *Thurston norm* $\|\cdot\|$ on $H^1(M; \mathbb{R})$ [26] (see also [12]). This norm has a polyhedral unit ball B with the following properties:

- (1) Every cohomology class dual to a fiber is in the cone $\mathbb{R}_+ \mathbf{F}$ over a top-dimensional open face $\mathbf{F}$ of B .
- (2) If $\mathbb{R}_+ \mathbf{F}$ contains a cohomology class dual to a fiber, then *every* primitive integral class in $\mathbb{R}_+ \mathbf{F}$ is dual to a fiber. The face $\mathbf{F}$ is called a *fibered face* and its primitive integral classes are called fibered classes.
- (3) For a fibered class ω with associated fiber F , $\|\omega\| = -\chi(F)$.

In particular, if $\dim H^1(M; \mathbb{R}) \geq 2$ and M is fibered, then there are infinitely many fibrations, with fibers of arbitrarily large complexity. We will abuse terminology by saying that a fiber (rather than its Poincaré dual) is in $\mathbb{R}_+ \mathbf{F}$.

The fibered faces also organize the suspension flows and the stable/unstable foliations: If $\mathbf{F}$ is a fibered face, then there are a single flow ψ and a single pair $\Lambda^\pm$ of foliations whose leaves are invariant by ψ , such that *every* fibration associated to $\mathbb{R}_+ \mathbf{F}$ may be isotoped so that its suspension flow is ψ up to a reparameterization, and the foliations $\lambda^\pm$ for the monodromy of its fiber F are $\Lambda^\pm \cap F$. These results were proven by Fried [13]; see also McMullen [18].

Finally, we note that the veering triangulation of $\mathring{M}$, like the flow itself, is an invariant of the fibered face containing the fiber $\mathring{S}$ (see [2] or [20, Proposition 2.7]).

2.4. Subsurfaces, q -compatibility, and τ -compatibility

We conclude this section by reviewing some essential constructions from [20] and direct the reader there for the full details. In short, the idea is that if Y is a subsurface of

(X, q) with $d_Y(\lambda^-, \lambda^+)$ sufficiently large, then Y has particularly nice forms; the first with respect to the q -metric, and the second with respect to τ .

Let $Y \subset X$ be an essential compact subsurface, and let $X_Y = \tilde{X}/\pi_1(Y)$ be the associated cover of X . We say a boundary component of Y is *puncture-parallel* if it bounds a disk in $\tilde{X} \setminus Y$ that contains a single point of $\mathcal{P}$. We denote the corresponding subset of $\mathcal{P}$ by $\mathcal{P}_Y$ and refer to its elements as the *punctures* of Y . Let $\tilde{\mathcal{P}}_Y$ denote the subset of punctures of X_Y which are encircled by the boundary components of the lift of Y to X_Y . In terms of the completed space $\bar{X}_Y$, $\tilde{\mathcal{P}}_Y$ is exactly the set of completion points which have finite total angle. Let $\partial_0 Y$ denote the union of the puncture-parallel components of ∂Y , and let $\partial' Y$ denote the rest. Observe that the components of $\partial_0 Y$ are in natural bijection with $\mathcal{P}_Y$ and set $Y' = Y \setminus \partial_0 Y$.

Identifying $\tilde{X}$ with $\mathbb{H}^2$, let $\Lambda_Y \subset \partial\mathbb{H}^2$ be the limit set of $\pi_1(Y)$, $\Omega_Y = \partial\mathbb{H}^2 \setminus \Lambda_Y$, and $\hat{\mathcal{P}}_Y \subset \Lambda_Y$ the set of parabolic fixed points of $\pi_1(Y)$. Let $C(X_Y)$ denote the compactification of X_Y given by $(\mathbb{H}^2 \cup \Omega_Y \cup \hat{\mathcal{P}}_Y)/\pi_1(Y)$, adding a point for each puncture-parallel end of X_Y , and a circle for each of the other ends.

q -convex hulls. As above, identify $\tilde{X}$ with $\mathbb{H}^2$, and let $\hat{X}$ be its metric completion with respect to the lift of q . Let $\Lambda \subset \partial\mathbb{H}^2$ be a closed set, and let $\text{CH}(\Lambda)$ be the convex hull of Λ in $\mathbb{H}^2$. Using the results of [20, Section 2.3], we define the q -convex hull $\text{CH}_q(\Lambda)$ as follows.

Assume first that Λ has at least 3 points. Each boundary geodesic l of $\text{CH}(\Lambda)$ has the same endpoints as a (biinfinite) q -geodesic l_q in $\hat{X}$ (we note that l_q may meet $\partial\mathbb{H}^2$ at interior points). Further, l_q is unique unless it is part of a parallel family of geodesics, making a Euclidean strip.

The metric completion $\hat{X}$ is divided by l_q into two sides, and one of the sides, which we call $\mathcal{D}_l$, meets $\partial\mathbb{H}^2$ in a subset of the complement of Λ . The side $\mathcal{D}_l$ is either a disk or a string of disks attached along completion points. If l_q is one of a parallel family of geodesics, we include this family in $\mathcal{D}_l$. After deleting from $\hat{X}$ the interiors of $\mathcal{D}_l$ for all l in $\partial\text{CH}(\Lambda)$, we obtain $\text{CH}_q(\Lambda)$, the q -convex hull. See Figure 2.




Figure 2. On the left are a few boundary geodesics of $\text{CH}(\Lambda)$. On the right are those geodesics replaced with the corresponding q -geodesics. Note that if l denotes the lower-right geodesic on the left side, then the interior of l_q (on the right side) meets $\partial\mathbb{H}^2$ in a single completion point. In this case, $\mathcal{D}_l$ is a string of two disks.

If Λ has two points, then $\text{CH}_q(\Lambda)$ is the closed Euclidean strip formed by the union of q -geodesics joining those two points.

Now, fixing a subsurface Y , we can define a q -convex hull for the cover X_Y by taking a quotient of the q -convex hull $\text{CH}_q(\Lambda_Y)$ of the limit set Λ_Y of $\pi_1(Y)$. This quotient, which we will denote by $\text{CH}_q(X_Y)$, lies in the completion $\bar{X}_Y$. We remark that in general $\text{CH}_q(X_Y)$ may be a total mess, e.g., it may have empty interior. The opposite situation, when the interior of $\text{CH}_q(X_Y)$ is homeomorphic to the interior of Y , is depicted in Figure 3. This is equivalent to the condition that Y is q -compatible, which we now turn to define.

q -compatibility. Let $\hat{i}: Y \rightarrow X_Y$ be the lift of the inclusion map to the cover. We say that the subsurface Y of X is q -compatible if the interior of $\text{CH}_q(\Lambda_Y)$ is a disk. In this case, [20, Lemma 2.6] implies that $\hat{i}: Y \rightarrow X_Y$ is homotopic to a map $\hat{i}_q: Y \rightarrow \bar{X}_Y$ which restricts to a homeomorphism from $Y \setminus \partial_0 Y$ to

$$Y_q = \text{CH}_q(X_Y) \setminus \tilde{\mathcal{P}}_Y.$$

We recall that by [20, Lemma 5.1], if Y is q -compatible, then

- (1) the projection $\iota_q: Y \rightarrow \bar{X}$ of $\hat{i}_q$ to $\bar{X}$ is an embedding from $\text{int}(Y)$ into X which is homotopic to the inclusion,
- (2) $\hat{i}_q(\partial Y \setminus \partial_0 Y)$ does not pass through points of $\tilde{\mathcal{P}}_Y$.

See Figure 3. The embedded image $\iota_q(\text{int}(Y))$ is an open representative of Y in X and is denoted by $\text{int}_q(Y)$.




Figure 3. The image of a q -compatible subsurface Y in $\bar{X}_Y$ under $\hat{i}_q$. Open circles are points of $\tilde{\mathcal{P}}_Y$ (corresponding to the image of $\partial_0 Y$), and dots are singularities not contained in $\tilde{\mathcal{P}}_Y$. The ideal boundary of X_Y is in blue.

The following is our main tool for proving q -compatibility; it is [20, Proposition 5.2].

Proposition 2.5 (q -compatibility). *Let $Y \subset X$ be an essential subsurface.*

- (1) *If Y is nonannular and $d_Y(\lambda^-, \lambda^+) \geq 3$, then Y is q -compatible.*
- (2) *If Y is an annulus and $d_Y(\lambda^-, \lambda^+) \geq 4$, then Y is q -compatible. In this case, $\text{int}_q(Y)$ is a flat cylinder.*

We remark that the constants in [20, Proposition 5.2] are slightly different since d_Y was defined there to be the minimal distance between projections.

τ -compatibility. Next we focus on compatibility with respect to the veering triangulation. Call a q -compatible subsurface $Y \subset X$ *τ -compatible* if the map $\hat{\iota}_q: Y \rightarrow \bar{X}_Y$ is homotopic rel $\partial_0 Y$ to a map $\hat{\iota}_\tau: Y \rightarrow \bar{X}_Y$ which is an embedding on $Y' = Y \setminus \partial_0 Y$ such that

- (1) $\hat{\iota}_\tau$ takes each component of $\partial' Y = \partial Y \setminus \partial_0 Y$ to a simple curve in $\bar{X}_Y \setminus \tilde{\mathcal{P}}_Y$ composed of a union of τ -edges,
- (2) the map $\iota_\tau: Y \rightarrow \bar{X}$ obtained by composing $\hat{\iota}_\tau$ with $\bar{X}_Y \rightarrow \bar{X}$ restricts to an embedding from $\text{int}(Y)$ into X .

When the subsurface Y is τ -compatible, we set

$$\partial_\tau Y \equiv \iota_\tau(\partial' Y)$$

which is a collection of τ -edges with disjoint interiors. We call $\partial_\tau Y$ the *τ -boundary* of Y and consider it as a 1-complex of τ -edges in X . (At times we will also think of $\partial_\tau Y$ as a collection of disjoint τ -edges in the fully-punctured surface $\mathring{X}$.) Similar to the situation of a q -compatible subsurface, if Y is τ -compatible, then one component of $X \setminus \partial_\tau Y$ is an open subsurface isotopic to the interior of Y ; this is the image $\iota_\tau(\text{int}(Y))$ and is denoted by $\text{int}_\tau(Y)$. For future reference, we set $Y_\tau \subset X_Y$ to be the intersection of X_Y with the image of $\hat{\iota}_\tau$. By definition, the covering $X_Y \rightarrow X$ maps the interior of Y_τ homeomorphically onto $\text{int}_\tau(Y)$.

The following result is [20, Theorem 5.3].

Theorem 2.6 (τ -compatibility). *Let $Y \subset X$ be an essential subsurface.*

- (1) *If Y is nonannular and $d_Y(\lambda^-, \lambda^+) \geq 3$, then Y is τ -compatible.*
- (2) *If Y is an annulus and $d_Y(\lambda^-, \lambda^+) \geq 4$, then Y is τ -compatible.*

The comment following Proposition 2.5 also applies here.

3. Veering triangulations and subsurfaces

Here we study the connection between sections of the bundle $\mathcal{N} \rightarrow \mathring{X}$ and projections to subsurfaces of X . In brief, we tailor the theory of subsurface projections to the veering structure. This is accomplished in Lemma 3.6 and Proposition 3.7.

3.1. Sections of the veering triangulation

A *section* of the veering triangulation τ in $\mathcal{N}$ is a simplicial embedding $s: (\mathring{X}, \Delta) \rightarrow \mathcal{N}$ that is a section of the fibration $\Pi: \mathcal{N} \rightarrow \mathring{X}$. Here, Δ is an ideal triangulation of $\mathring{X}$, which by construction consists of τ -edges. We will also refer to the image of s in $\mathcal{N}$, which we often denote by T , as a section.

There is a bijective correspondence between sections of $\mathcal{N}$ and ideal triangulations of $\mathring{X}$ by τ -edges. More generally, we use the notation Π_* to denote the map that associates




Figure 4. Two τ -edges with $e > f$.

to any subcomplex of a section the corresponding union of τ -simplices of $\overset{\circ}{X}$, and we use Π^* to denote its inverse. In particular, if K is a union of disjoint τ -edges of $\overset{\circ}{X}$, $\Pi^*(K)$ is the subcomplex of $\mathcal{N}$ obtained by lifting its simplices to $\mathcal{N}$. Note that T and T' differ by a tetrahedron move in $\mathcal{N}$ if and only if the ideal triangulations $\Pi_*(T)$ and $\Pi_*(T')$ differ by a diagonal exchange. Here, an upward (downward) tetrahedron move on a section T replaces two adjacent faces at the bottom (top) of a tetrahedron with the two adjacent top (bottom) faces.

Since the fibers of $\Pi: \mathcal{N} \rightarrow \overset{\circ}{X}$ give $\mathcal{N}$ an oriented foliation by lines and each of these lines meets each section exactly once, we have the following observation: For each $x \in \mathcal{N}$ and each section T of $\mathcal{N}$, it makes sense to write $x \leq T$ or $x \geq T$ depending on whether x lies weakly below or above T along the orientation of the line through x . (Here, $x \leq T$ and $x \geq T$ imply that $x \in T$.) In fact, this ordering extends to each simplex of $\mathcal{N}$; we write $\sigma \leq T$ if $x \leq T$ for each $x \in \sigma$. Since we will use this fiberwise ordering for several (simplicial) constructions, it is important to note that it is consistent along simplices; that is, if $x \leq T$ and σ is the smallest simplex containing x , then $\sigma \leq T$. Finally, if K is a subcomplex of $\mathcal{N}$, then $K \leq T$ if for each simplex σ of K , $\sigma \leq T$.

In [20, Section 2.1], we define a strict partial order among τ -edges using their spanning rectangles: If e crosses f , we say that $e > f$ if e crosses the spanning rectangle of f from top to bottom, and f crosses the spanning rectangle of e from left to right (i.e., if the slope of e is greater than the slope of f). A priori, this partial order is defined in the universal cover $\tilde{X}$, but it projects consistently to $\overset{\circ}{X}$ and so defines a partial order of τ -edges there as well. See Figure 4.

This definition is consistent with the ordering of the simplices $\Pi^*(e)$, $\Pi^*(f)$ in $\mathcal{N}$, and, in particular, the following holds.

Lemma 3.1. *With notation as above, $K \leq T$ if and only if whenever an edge e of $\Pi_*(T)$ crosses an edge f of $\Pi_*(K)$, we have $f < e$.*

Proof. First, let f be any τ -edge in $\overset{\circ}{X}$, and let Δ be a triangulation of $\overset{\circ}{X}$ by τ -edges. By [20, Lemma 3.4], if e is an edge of Δ and $e > f$, then there is an edge of Δ crossing f which is downward flippable, meaning that there is a diagonal exchange of Δ replacing it with an edge of smaller slope. (An analogous statement holds if $f > e$.) Such a diagonal exchange results in a triangulation Δ' either containing f or still containing an edge e'

with $e' > f$. After finitely many downward diagonal exchanges, we arrive at a triangulation by τ -edges which contains the edge f . (See [20, Section 3] for details.) Translating this statement to $\mathcal{N}$, this means that starting with $\Pi^*(\Delta)$ there is a sequence of downward tetrahedron moves resulting in a section containing $\Pi^*(f)$. Hence, $\Pi^*(f) \leq \Pi^*(\Delta)$.

This, together with the corresponding result for when $K \geq T$, implies the lemma. ■

Given sections T_1 and T_2 , we use $U(T_1, T_2)$ to denote the subcomplex of $\mathcal{N}$ between them. Formally, $U(T_1, T_2)$ is the subcomplex of $\mathcal{N}$ which is the union of all simplices σ such that either $T_1 \leq \sigma \leq T_2$ or $T_2 \leq \sigma \leq T_1$.

It will be helpful to consider the lattice structure of sections. For sections T_1, T_2 , we denote their fiberwise *maximum* by $T_1 \vee T_2$. If we name the oriented fiber containing x by l_x , this is the subset $\{x \in \mathcal{N} : x = \max_l \{l_x \cap T_1, l_x \cap T_2\}\}$, where the max is taken with respect to the ordering on each l_x . See Figure 5.




Figure 5. The section T_1 is blue and the section T_2 is red. The section $T_1 \vee T_2$ is highlighted in yellow.

Lemma 3.2. *If T_1 and T_2 are sections, then $T_1 \vee T_2$ is a section.*

Proof. Since the restriction of Π to $T_1 \vee T_2$ is a homeomorphism to $\mathring{X}$, it suffices to show that $T_1 \vee T_2$ is a subcomplex of $\mathcal{N}$. Let $x \in T_1 \vee T_2$ and suppose that x is contained in T_1 . Then $x \geq T_2$ and so if σ is the minimal simplex of T_1 containing x , we see $\sigma \geq T_2$. Hence, $\sigma \subset T_1 \vee T_2$, and we conclude that $T_1 \vee T_2$ is indeed a section. ■

We can define the *minimum* $T_1 \wedge T_2$ of two sections similarly. With this terminology, it makes sense to say that $T_1 \vee T_2$ is the *top* of $U(T_1, T_2)$. More precisely, $T_1 \vee T_2 \subset U(T_1, T_2)$ and for every simplex $\sigma \subset U(T_1, T_2)$, $\sigma \leq T_1 \vee T_2$. Similarly, we say that $T_1 \wedge T_2$ is the *bottom* of $U(T_1, T_2)$. Further, using our definitions we see that

$$U(T_1, T_2) = U(T_1 \wedge T_2, T_1 \vee T_2).$$

Note that the part of T_1 that lies above T_2 is $T_1 \cap (T_1 \vee T_2)$.

Sections through a subcomplex. Let $E \subset \mathring{X}$ be a union of disjoint τ -edges, and set $K = \Pi^*(E)$ to be the corresponding subcomplex of $\mathcal{N}$. (Our primary example will be $E = \partial_\tau Y$ for a τ -compatible subsurface Y of X . In this situation, we think of $\partial_\tau Y$ as

a collection of τ -edges of $\overset{\circ}{X}$.) We define $T(E) = T(K)$ to be the set of sections of $\mathcal{N}$ which contain K as a subcomplex. Similarly, we define $\Delta(E) = \Delta(K)$ as the set of ideal triangulations of $\overset{\circ}{X}$ by τ -edges containing E . We recall the following two basic results from [20]. The first states simply that $T(E)$ is nonempty. It is [20, Lemma 3.2].

Lemma 3.3 (Extension lemma). *Suppose that E is a collection of τ -edges in $\overset{\circ}{X}$ with pairwise disjoint interiors. Then $T(E)$ is nonempty.*

The second [20, Proposition 3.3] states that $T(K)$ is always connected by tetrahedron moves. This includes, in particular, the case of $T(\emptyset)$, the set of all sections.

Lemma 3.4 (Connectivity). *If E is a collection of τ -edges in X with pairwise disjoint interiors, then $\Delta(E)$ is connected via diagonal exchanges. In terms of $\mathcal{N}$, for $K = \Pi^*(E)$, $T(K)$ is connected via tetrahedron moves.*

Moreover, if $T_1, T_2 \in T(K)$ with $T_1 \leq T_2$, then there is a sequence of upward tetrahedron moves from T_1 to T_2 through sections of $T(K)$.

As explained in [20, Corollary 3.6], whenever $E \neq \emptyset$, there is a well-defined *top* T^+ and *bottom* T^- of $T(E)$. That is, $T^+ \in T(E)$ and for any $T \in T(E)$, $T \leq T^+$.

φ -sections. Let q be a quadratic differential associated to a pseudo-Anosov homeomorphism φ . Recall that the deck transformation Φ of $\mathcal{N}$ is chosen to translate in the opposite direction of the flow.

Say that a section T of the veering triangulation τ is a φ -section if $\Phi(T) \leq T$. In other words, T is a φ -section if every τ -edge of $\Pi_*(\Phi(T)) = \varphi(\Pi_*(T))$ which crosses a τ -edge of $\Pi_*(T)$ does so with lesser slope. Note that if T is a φ -section, then $\Phi^j(T) \leq \Phi^i(T)$ for all $i \leq j$.

Agol's original construction produces a veering triangulation from a sequence of diagonal exchanges through φ -sections [1, Proposition 4.2]. In fact, he proves the following assertion.

Lemma 3.5 (Agol). *There is a sweep-out of τ through φ -sections. That is, there is a sequence $(T_i)_{i \in \mathbb{Z}}$ of φ -sections such that T_{i+1} is obtained from T_i by simultaneous upward tetrahedron moves.*

3.2. Projections to τ -compatible subsurfaces

In this section, we define a variant of the subsurface projections π_Y that is adapted to the simplicial structure of τ . In the hyperbolic setting, $\pi_Y(\alpha)$ can be defined using the geodesic representatives of the surface Y and the curve α . In our setting, we need to use the simplicial representative Y_τ of a τ -compatible surface Y and a collection of τ -edges representing α . The main result here will be Proposition 3.7, which shows, in a suitable setting, that the simplicial variant of the projection is uniformly close to the usual notion.

Recall first the notion of a proper graph G in a surface from Section 2.1.3 and its image $\mathcal{A}_S(G)$ in the arc graph (definition (2.3)).

If E is a collection of τ -edges of $\overset{\circ}{X}$ with disjoint interiors, then its closure in X , $\text{cl}_X(E)$, is a proper graph. This is the union of the corresponding saddle connections in X . In particular, if K is a subcomplex of a section, then $E = \Pi_*(K^{(1)})$ is such a collection of τ -edges and we make the notational abbreviation

$$\mathcal{A}_X(K) = \mathcal{A}_X(\text{cl}_X(\Pi_*(K^{(1)}))). \quad (3.1)$$

Note that for any section T , we have (Corollary 2.4) $\text{diam}(\mathcal{A}_X(T)) \leq D$, where $D = 15$.

Suppose that $Y \subset X$ is a τ -compatible nonannular subsurface and $G \subset X$ is a union of τ -edges with disjoint interiors (i.e., a *proper graph of τ -edges*). Then $\text{int}_\tau(Y) \cap G$ is a proper graph in $\text{int}_\tau(Y)$, and we set

$$\pi_Y^\tau(G) = \mathcal{A}_Y(\text{int}_\tau(Y) \cap G). \quad (3.2)$$

Note that this could in general be empty, if $\text{int}_\tau(Y) \cap G$ is not essential.

When Y is a τ -compatible annulus, $\text{int}_\tau(Y) \cap G$ is a collection of disjoint arcs each of which is contained in the interior of a τ -edge. Taking the preimages of these τ -edges in X_Y , we obtain the projection $\pi_Y^\tau(G)$ by associating to each such τ -edge a that joins *opposite* sides of $\partial Y_\tau \subset X_Y$, the collection of complete q -geodesics a^* in X_Y that contain it. Each of these geodesics gives a well-defined arc of $\mathcal{A}(Y) = \mathcal{A}(X_Y)$, and if there are no such τ -edges, then the projection is empty.

In general, this notion of subsurface projection is easily extended to a subcomplex K of a section of $\mathcal{N}$, in analogy with (3.1). We simply write

$$\pi_Y^\tau(K) = \pi_Y^\tau(\text{cl}_X(\Pi_*(K^{(1)}))).$$

Finally, we define the subsurface distance between subcomplexes K_1 and K_2 to be

$$d_Y(K_1, K_2) = d_Y(\pi_Y^\tau(K_1), \pi_Y^\tau(K_2)). \quad (3.3)$$

The following lemma establishes some important technical properties of τ -compatible subsurfaces. Key to the argument is the construction of $\text{int}_\tau(Y)$ from $\text{int}_q(Y)$ that appears in [20, Theorem 5.3] and is illustrated in Figure 6.

We remark that one difficulty in what follows is that the τ -representative $\text{int}_\tau(Y)$, unlike the q -representative $\text{int}_q(Y)$, is *not* convex with respect to the q metric.

Lemma 3.6. *Let Y and Z be τ -compatible subsurfaces of X , and let $G \subset X$ be a proper graph of τ -edges.*

- (1) *The diameter $\text{diam}_Y(\pi_Y^\tau(G))$ is bounded by $D = 15$. If Y is an annulus, then $\text{diam}_Y(\pi_Y^\tau(G)) \leq 3$.*
- (2) *If Y and Z are disjoint, then so are $\text{int}_\tau(Y)$ and $\text{int}_\tau(Z)$.*
- (3) *The subsurface $\text{int}_\tau(Y)$ is in minimal position with the foliation $\lambda^\pm$. In particular, the arcs of $\text{int}_\tau(Y) \cap \lambda^\pm$ agree with the arcs of $\pi_Y(\lambda^\pm)$.*




Figure 6. Left: $\text{int}_q(Y)$ is shaded, contains the blue singularities, and its boundary contains the black saddle connection. Middle: the “inner $\mathbf{t}$ -hull construction” isotopes the surface within itself; its new boundary consists of saddle connections (dark red) through the blue singularities. Right: $\text{int}_\tau(Y)$, whose boundary consists of τ -edges (blue), is then produced with the “outer $\mathbf{t}$ -hull construction”.

Proof. The graph $G \cap \text{int}_\tau(Y)$ has its vertices in $\text{int}_\tau(Y) \cap \text{sing}(q)$. By Gauss–Bonnet, $|\text{int}_\tau(Y) \cap \text{sing}(q)| \leq 2|\chi(Y)|$, and so (1) follows from Corollary 2.4 when Y is not an annulus.

If Y is an annulus, then recall that $\text{int}_q(Y)$ is a flat annulus and that $\text{int}_q(Y) \subset \text{int}_\tau(Y)$ since the process of going from the q -hull to the τ -hull for annuli only pushes outward (cf. [20, Remark 5.4]). Lifting to the annular cover X_Y , let a, b be τ -edges coming from G that join opposite sides of Y_τ . Then a and b are disjoint and any of their q -geodesic extensions a^*, b^* cross the maximal open flat annulus $\text{int}_q(Y)$ of X_Y in subsegments of a, b , respectively. Moreover, by a standard Gauss–Bonnet argument (e.g., [21, Lemma 3.8]), any two q -geodesic segments intersect at most once in any component of $X_Y \setminus \text{int}_q(Y)$. Hence, a^*, b^* intersect at most twice and so $\text{diam}_Y(\pi_Y^\tau(G)) \leq 3$ when Y is an annulus.

Items (2) and (3) follow exactly as in [20, Lemma 6.1]. As that lemma was proven only in the fully-punctured case, we note that in general for item (2) one must perform the inner $\mathbf{t}$ -hull construction (middle of Figure 6) as an intermediate step. However, since this pushes each surface within itself, it must preserve disjointness. For (3), the isotopy from $\text{int}_q(Y)$ to $\text{int}_\tau(Y)$ pushes along the leaves of λ^+ (or λ^-). Hence, the leaves of λ^+ (or λ^-) are in minimal position with $\text{int}_\tau(Y)$ since they are with $\text{int}_q(Y)$, by the local CAT(0) geometry. ■

We note that it follows from Lemma 3.6(3) (or directly from its proof) that if $\tilde{\lambda}^\pm$ are the lifts of $\lambda^\pm$ to $\tilde{X}$ and C_Y^τ is a component of the preimage of $\text{int}_\tau(Y)$, then the intersection of C_Y^τ with each leaf of $\tilde{\lambda}^\pm$ is connected.

The following proposition relates the projections defined here to the usual notion of subsurface projection.

Proposition 3.7. *Let Y, Z be τ -compatible subsurfaces of X , and assume $\pi_Z(\partial Y) \neq \emptyset$. Further assume that $d_Y(\lambda^-, \lambda^+) \geq 3$ (or if Y is an annulus that $d_Y(\lambda^-, \lambda^+) \geq 6$). Then*

$$\text{diam}_Z(\pi_Z^\tau(\partial_\tau Y) \cup \pi_Z(\partial Y)) \leq 7.$$

As usual, let X_Z denote the Z -cover of X . In this subsection, a *core* Z' of X_Z is a submanifold with boundary which is a complementary component of simple curves and

proper arcs such that $Z' \rightarrow X_Z$ is a homotopy equivalence. This definition includes the usual convex core in the hyperbolic metric as well as $Z_\tau \subset X_Z$ (Section 2.4). In general, $\partial Z' \subset X_Z$ is a collection of curves and arcs.

Let γ be any essential curve or proper arc in X_Z . Note that γ may not be in minimal position with $\partial Z'$, that is, there may be bigons between arcs of γ and $\partial Z'$. To handle this situation, we make the following definition: For some $k \geq 0$, we say that γ is in k -position with respect to $\partial Z'$ if k is the largest integer such that a collection of k nested subsegments of $\partial Z'$ cobound bigons with subsegments of γ whose interiors are contained in $X_Z \setminus Z'$. See Figure 7.




Figure 7. The curve/arc γ (blue) is in 4-position with respect to ∂Z , but not 3-position. Note that the arcs of ∂Z cobounding these bigons are nested.

The point of this condition is the following lemma.

Lemma 3.8. *Suppose that Z is a nonannular subsurface of X and that X_Z and Z' are as above. Let γ be an essential curve or proper arc in X_Z which is in k -position with respect to $\partial Z'$, and let y be an essential arc of $\gamma \cap Z'$. Then*

$$d_Z(y, \pi_Z(\gamma)) \leq 2 \log(2k) + 2.$$

Proof. To prove the lemma, we consider y as an arc of X_Z as follows: For each endpoint $p \in y \cap \partial Z'$ append to y a proper ray in X_Z starting at p which meets Z' only at p . Let y^* denote the resulting essential arc of X_Z . (Under the implicit/canonical identification $\mathcal{A}(\text{int}(Z')) \cong \mathcal{A}(X_Z)$, y and y^* are identified.) Now we claim that as isotopy classes of arcs in X_Z , y^* and γ have at most $2k$ essential intersections, thus proving the lemma.

For this, first push y^* slightly to one side of itself, so that y^* and γ are transverse. Then each point $p_i \in y^* \cap \partial Z'$ ($i = 1, 2$) is contained in no more than k nested subsegments of $\partial Z'$ which cobound bigons $\mathcal{B}_i$ with γ , as in the definition of k -position. See Figure 8.

Since each component of $X_Z \setminus \text{int}(Z')$ is a disk or an annulus, it follows that there is an isotopy supported in $X_Z \setminus \text{int}(Z' \cup \mathcal{B}_1 \cup \mathcal{B}_2)$ which removes all of the intersections between γ and y^* which are not contained in $\mathcal{B}_1 \cup \mathcal{B}_2$. Hence, γ and y^* have at most $2k$ essential intersections. This completes the proof. ■




Figure 8. The arc γ (green) is in 2-position with ∂Z . The arc γ^* (blue) intersects γ essentially no more than twice.

Because $\text{int}_\tau(Y)$ is an *open* subsurface representative of Y , it does not provide a good representative of ∂Y . So for the proof of Proposition 3.7, we do the following: Let Y_n ($n \geq 1$) be the exhaustion of $\text{int}_\tau(Y)$ by subsurfaces isotopic to Y obtained by removing the open ε_n -neighborhood of $\partial_\tau Y$. Here, $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$, and distance is taken with respect to the flat metric q . Note that by our definitions, each $\mathcal{A}(Y_n)$ is naturally identified with $\mathcal{A}(Y)$. When Y is not an annulus, we will use the property that, through this identification, for any curve or arc a in X , the collection of curves and arcs in $\mathcal{A}(Y)$ given by $a \cap Y_n$ eventually agrees with the collection given by $a \cap \text{int}_\tau(Y)$.

Proof of Proposition 3.7. First suppose that Z is not an annulus.

As above, let Y_n be the exhaustion of $\text{int}_\tau(Y)$ introduced above. To keep notation simple, we set $Y = Y_n$ for n sufficiently large (to be determined later). Note that by construction, ∂Y and $\partial_\tau Y$ are disjoint in X . Hence, if γ is any essential component of the preimage of ∂Y in X_Z , then γ is disjoint from the preimage of $\partial_\tau Y$. Letting y be any essential component of $\gamma \cap Z_\tau$, this shows that $d_Z(y, \pi_Z^\tau(\partial_\tau Y)) \leq 1$. Hence, it suffices to bound the distance between y and $\pi_Z(\gamma)$ in $\mathcal{A}(Z)$. This will follow from Lemma 3.8, once we show that γ is in 2-position with respect to ∂Z_τ in X_Z .

Suppose that this were not the case; that is, suppose that there is a point $p \in \partial Z_\tau$ which is contained in 3 nested subsegments of ∂Z_τ , each of which cobounds a bigon with a subarc of γ contained in $X_Z \setminus \text{int}(Z_\tau)$. We now lift this picture to $\tilde{X}$ to produce a component C_Z^τ of the preimage of Z_τ in $\tilde{X}$, a point $\tilde{p} \in \partial C_Z^\tau$, and arcs $\tilde{\gamma}_1, \tilde{\gamma}_2, \tilde{\gamma}_3$ in the preimage of γ that cobound bigons with nested subsegments of ∂C_Z^τ containing $\tilde{p}$. (The arcs are ordered so that $\tilde{\gamma}_1$ is innermost, i.e., closest to $\tilde{p}$, and $\tilde{\gamma}_3$ is outermost.) Let B be the bigon cobounded by a subarc of ∂C_Z^τ and $\tilde{\gamma}_3$.

Since these arcs are components of preimages of ∂Y , we can use them to produce a component C_Y of the preimage of Y in $\tilde{X}$ such that

- (1) some component of ∂C_Y contains $\tilde{\gamma}_1$ or $\tilde{\gamma}_2$,
- (2) $C_Y \cap B$ contains a component which is a disk whose sides alternate between subarcs of ∂C_Y and subarcs of ∂C_Z^τ .




Figure 9. The component C_Y (green) and its intersection with $l^\pm$ (blue). The intersection of the black zig-zag arc with $\tilde{X}$ give components of ∂C_Y^τ .

For simplicity, we assume that $\tilde{\gamma}_1$ is contained in the component of ∂C_Y in item (1) and denote the other component of $\partial C_Y \cap B$ that appears in item (2) by $\tilde{\eta}$. (It may be that $\tilde{\eta} = \tilde{\gamma}_2$.) See Figure 9.

Now since the leaves of $\lambda^\pm$ are in minimal position with respect to Z_τ (Lemma 3.6 and the comment that follows), the intersection of each leaf of $\tilde{\lambda}^\pm$ (the lifts of $\lambda^\pm$ to $\tilde{X}$) with C_Y^τ is connected. Hence, we may choose subrays $l^\pm$ in $\tilde{\lambda}^\pm$ which are based at $\tilde{p}$ and are disjoint from $\text{int}(C_Y^\tau)$. By condition (2) above, each of $l^\pm$ must pass first through $\tilde{\gamma}_1$ and then through $\tilde{\eta}$. Moreover, each of $l^\pm \cap C_Y$ is an arc joining distinct components of ∂C_Y , once n is sufficiently large, again by Lemma 3.6, this time applied to Y .

If Y is also nonannular, then we conclude that the two leaf segments $l^\pm \cap C_Y$ project to give homotopic components of $\lambda^+ \cap Y$ and $\lambda^- \cap Y$. Once n is sufficiently large, this gives homotopic components of $\lambda^+ \cap \text{int}_\tau(Y)$ and $\lambda^- \cap \text{int}_\tau(Y)$, and by Lemma 3.6, we see that $d_Y(\lambda^+, \lambda^-) \leq 2$. This contradiction shows that γ is in 2-position with ∂Z_τ and completes the proof when neither Y nor Z are annuli.

If Y is an annulus, then as remarked above, $\text{int}_q(Y)$ is a flat annulus and $\text{int}_q(Y) \subset \text{int}_\tau(Y)$. The argument above produces rays $l^\pm$ such that $l^\pm \cap C_Y$ are leaf segments of $\tilde{\lambda}^\pm$. We claim that these segments project to disjoint arcs of Y . Otherwise, there is a deck transformation g that stabilizes C_Y such that $(g \cdot l^- \cap l^+) \cap C_Y$ is nonempty. In this case, we have $g \cdot B \cap B \neq \emptyset$ (where B is the bigon from above), and so $g \cdot C_Y^\tau \cap C_Y^\tau \neq \emptyset$. Since Z is τ -compatible, this implies that g stabilizes C_Y^τ and contradicts our assumption that Y and Z overlap.

Hence, we conclude that the intersections $l^\pm \cap C_Y$ project to disjoint leaf segments in $Y \subset \text{int}_\tau(Y) \subset X_Y$. For n large as above, these leaves do not intersect again before exiting $\text{int}_\tau(Y)$ in X_Y . Since leaves of $\lambda^\pm$ intersect at most once outside the maximal open flat annulus $\text{int}(Y_q) \subset X_Y$, this produces representatives of $\pi_Y(\lambda^-)$ and $\pi_Y(\lambda^+)$ intersecting at most twice. Hence, $d_Y(\lambda^+, \lambda^-) \leq 5$ giving the same contradiction as before.

It remains to establish the proposition when Z is an annulus. Since $\pi_Z(\partial Y) \neq \emptyset$, there is a τ -edge in the lift of $\partial_\tau Y$ to X_Z that joins boundary components of Z_τ on opposite sides of the core curve of X_Z . Fix any such τ -edge a . Since $\text{int}_q(Z) \subset \text{int}_\tau(Z)$, a crosses $\text{int}_q(Z) \subset X_Z$ and has its endpoints in $X_Z \setminus \text{int}_q(Z)$. Recall the definition of $a^* \in \pi_Z^\tau(\partial_\tau Y)$ from the discussion following equation (3.2).

Let γ be any essential component of the lift of ∂Y to X_Z , and let γ_q be its geodesic representative in the q -metric. As before, the preimage of ∂Y in X_Z can be made disjoint from γ_q since Y is q -compatible. Since γ_q is a q -geodesic, its intersection with $\text{int}_q(Z) \subset X_Z$ is contained in a single saddle connection b . Moreover, because Y is τ -compatible, any saddle connection of $\partial_q Y \subset X$ intersects any τ -edge of $\partial_\tau Y \subset X$ at most once in its interior (see, e.g., [20, Theorem 5.3]), and so a and b intersect at most once. If a^* is any extension of a to a complete q -geodesic in X_Y , then we have by the same Gauss–Bonnet argument as for the proof of Lemma 3.6 that γ_q intersects a^* at most three times. We conclude that

$$\text{diam}_Z(\pi_Z^\tau(\partial_\tau Y) \cup \pi_Z(\partial Y)) \leq 2 + d_Z(a^*, \gamma_q) \leq 6,$$

and the proof of Proposition 3.7 is complete. $\blacksquare$

4. From immersions to covers

When defining the subsurface projection π_Y for a subsurface $Y \subset S$, we consider preimages of curves in the cover S_Y associated to Y . Of course, the same operation can be done for any cover of S corresponding to a finitely generated subgroup of $\pi_1(S)$. The main theorem of this section (Theorem 4.1, which is Theorem 1.3 in the introduction) gives a concrete explanation for why these more general projections do not capture additional information. This will be an essential ingredient for the proof of Theorem 7.1.

First, for a finitely generated subgroup $\Gamma < \pi_1(S)$, let $S_\Gamma \rightarrow S$ be the corresponding cover. If $Y \subset S_\Gamma$ is a compact core, the covering map restricted to Y is an immersion, and we say that $Y \rightarrow S$ *corresponds to* $\Gamma < \pi_1(S)$. Lifting to the cover induces a (partially defined) map of curve and arc graphs which we denote by $\pi_Y: \mathcal{A}(S) \rightarrow \mathcal{A}(Y)$, and we define $d_Y(\alpha, \beta)$ and diam_Y accordingly. Note that these constructions depend on Γ and not on the choice of Y , and that π_Y agrees with the usual subsurface projection when $Y \subset S$.

The goal of this section is the following theorem, which may be of independent interest.

Theorem 4.1 (Immersion to cover). *There is a uniform constant $M \leq 38$ satisfying the following: Let S be a surface and let $Y \rightarrow S$ be an immersion corresponding to a finitely-generated $\Gamma < \pi_1(S)$. Then either $\text{diam}_Y(\mathcal{A}(S)) \leq M$, or $Y \rightarrow S$ is homotopic to a finite cover $Y \rightarrow W$ for W a subsurface of S .*

Theorem 4.1 will follow as a corollary of the following statement.

Theorem 4.2. *Let $Y \rightarrow S$ be an immersion corresponding to $\Gamma < \pi_1(S)$, and let $\lambda^\pm$ be a transverse pair of foliations without saddle connections. If $d_Y(\lambda^-, \lambda^+) \geq 37$, then $Y \rightarrow S$ is homotopic to a finite cover $Y \rightarrow W$ for W a subsurface of S .*

Let q be a quadratic differential whose horizontal and vertical foliations are $\lambda^\pm$, and let X denote S endowed with q . For a curve or arc δ , denote its horizontal length with respect to q by $h_q(\delta)$ and its vertical length with respect to q by $v_q(\delta)$. For a homotopy class, we let h_q and v_q denote the minima over all representatives. Recall that a multicurve γ is *balanced* at q if $h_q(\gamma) = v_q(\gamma)$. For the quadratic differential q and a multicurve γ , there is always some time $t \in \mathbb{R}$ such that γ is balanced at q_t , where q_t is the image of q under the Teichmüller flow for time t .

We let $X_Y = X_\Gamma$ denote the associated cover, and recall from Section 2.4, the definition of the q -hull $Y_q \subset X_Y$. If $d_Y(\lambda^-, \lambda^+) \geq 5$, then Y_q is embedded in X_Y , by which we mean the map $\hat{\iota}_q$ is an embedding on $Y \setminus \partial_0 Y$ (Proposition 2.5). In the language of the previous section, Y_q is a core of X_Y , and $\partial Y_q \subset X_Y$ is a collection of locally geodesic curves and proper arcs (of finite q -length).

If α, β are curves or properly embedded arcs in X , then we let $j_Y(\alpha, \beta)$ denote the minimum, over all components a of $\pi_Y(\alpha)$ and b of $\pi_Y(\beta)$, of the number of intersection points of a and b . We may also use the same notation if α, β are laminations in S , or essential curves or properly embedded arcs in Y . The following inequality comes essentially from [21].

Lemma 4.3. *Suppose Y_q is embedded in the cover X_Y , and ∂Y is balanced with respect to q . Then for any essential curve or arc δ in Y_q , we have*

$$4 \cdot \ell_q(\delta) \geq j_Y(\lambda, \delta) \cdot \ell_q(\partial Y),$$

for λ equal to either λ^+ or λ^- .

Proof. We show the inequality for λ^+ . First recall that Y_q contains a union of maximal vertical strips $S_1, \dots, S_n$ with disjoint, singularity-free interiors having the property that $h_q(\partial Y) = 2 \sum_i h_q(S_i)$. Here, $h_q(S_i)$ denotes the width of the strip S_i . For details of this construction, see [21, Section 5]. If δ is an essential curve or arc of Y_q , then δ crosses each strip S_i at least $j_Y(\delta, \lambda^+)$ times since each strip is foliated by segments of λ^+ . Hence

$$h_q(\delta) \geq \sum_i h_q(S_i) \cdot j_Y(\delta, \lambda^+) = \frac{1}{2} h_q(\partial Y) \cdot j_Y(\delta, \lambda^+).$$

Since ∂Y is balanced, $h_q(\partial Y) = v_q(\partial Y)$ and so $\ell_q(\partial Y) \leq 2h_q(\partial Y)$. We conclude

$$\ell_q(\delta) \geq h_q(\delta) \geq \frac{1}{4} \ell_q(\partial Y) \cdot j_Y(\delta, \lambda^+)$$

as required. ■

We now proceed with the proof of Theorem 4.2.

Proof of Theorem 4.2. We may suppose that $d_Y(\lambda^-, \lambda^+) \geq 5$ so that Y_q is embedded in X_Y . If Y is an annulus, then $\text{int}(Y_q) \subset X_Y$ is a flat annulus which must cover a flat annulus in X . So we now suppose that Y is not an annulus. We may further assume, applying the Teichmüller flow to q if necessary, that ∂Y_q is balanced.

Let P be either S^1 or $\mathbb{R}$. If $\gamma: P \rightarrow X_Y$ is a parameterization of a boundary component of Y_q (see Figure 3), then we say that a *re-elevation* of γ is any lift to X_Y of $\mathbb{R} \rightarrow P \xrightarrow{\gamma} X_Y \xrightarrow{p} X$, where $\mathbb{R} \rightarrow P$ is the universal cover, and we say that the re-elevation is *essential* if it meets Y_q essentially. From nonpositive curvature of the metric, a re-elevation is essential if and only if it meets $\text{int}(Y_q)$.

We note that with this terminology, $Y \rightarrow X$ covers a subsurface of X if and only if there are no essential re-elevations of components of ∂Y_q , since in this case

$$p^{-1}(p(\partial Y_q)) \cap Y_q = \partial Y_q$$

(see [20, Lemma 6.6]). Thus our goal now is to prove that if $d_Y(\lambda^-, \lambda^+)$ is sufficiently large, independent of Y and X , then there are no essential re-elevations of ∂Y_q .

Let $g: \mathbb{R} \rightarrow X_Y$ be a re-elevation of γ , let $(a, b) \subset \mathbb{R}$ be a component of $g^{-1}(\text{int}_q Y)$, and let $\tilde{\gamma}$ denote the restriction of g to $[a, b]$. This is a geodesic path (possibly with self-intersection) with endpoints in ∂Y_q , and if g is essential, then $\tilde{\gamma}$ may be chosen to be essential.

We now look for a restriction of $\tilde{\gamma}$ to an essential simple arc or curve. Let $d \in [a, b]$ be the supremum over $t \in [a, b]$ for which $\tilde{\gamma}|_{[a, t]}$ is an embedding.

Case 1: $d < b$. Then there exists $c \in [a, d)$ such that $\tilde{\gamma}(d) = \tilde{\gamma}(c)$. Thus $x = \tilde{\gamma}|_{[c, d]}$ is an embedded loop.

Case 1a: x is an essential loop, which we name σ . In the case where $P = S^1$, divide $\mathbb{R}$ into fundamental domains for the covering map $\mathbb{R} \rightarrow S^1$, so that c is a boundary point of a fundamental domain, and let $l \geq 0$ be the number of full fundamental domains contained in $[c, d]$. When $P = \mathbb{R}$, set $l = 0$. Hence, as γ is a component of $\partial_q Y$,

$$\ell_q(\sigma) \leq \ell_q(\tilde{\gamma}|_{[c, d]}) \leq (l + 1) \cdot \ell_q(\partial_q Y).$$

By Lemma 4.3,

$$j_Y(\lambda^\pm, \sigma) \cdot \ell_q(\partial Y) \leq 4 \cdot \ell_q(\sigma) \leq 4(l + 1) \cdot \ell_q(\partial Y),$$

and so $j_Y(\lambda^\pm, \sigma) \leq 4(l + 1)$.

The Gauss–Bonnet theorem for the Euclidean cone metric on Y_q implies that the number of singularities in the interior of Y_q is no more than $2|\chi(Y)|$, and since $\tilde{\gamma}|_{[c, d]}$ is an embedded loop, it visits each of these singularities of Y_q at most once. As $\tilde{\gamma}|_{[c, d]}$ contains at least l singularities interior to Y_q , we obtain $l \leq 2|\chi(Y)|$. Combining this with the inequality arrived at above, we conclude that

$$j_Y(\lambda^\pm, \sigma) \leq 4(2 \cdot |\chi(Y)| + 1).$$

We now invoke Lemma 2.3 (and in particular equation (2.1)) to conclude that

$$d_Y(\lambda^\pm, \sigma) \leq 15.$$

Hence, $d_Y(\lambda^-, \lambda^+) \leq 30$ as required.

Case 1b: If $x = \tilde{\gamma}|_{[c,d]}$ is inessential, it still cannot be null-homotopic since g is a geodesic path, so it must be peripheral. That is, either x bounds a punctured disk in Y_q or x cobounds an annulus A with a boundary component u of Y_q . We claim that in the latter case the endpoint $\tilde{\gamma}(a)$ does not lie on u .

Suppose otherwise. Then there are two possibilities. If $\tilde{\gamma}|_{[d,d+\varepsilon)}$ does not enter A , then x is a q -geodesic loop: at $\tilde{\gamma}(c)$ it subtends an angle of at least π inside A , and at $\tilde{\gamma}(d)$ it subtends an angle of at least π outside A . (See Figure 10.) But this is a contradiction, x cannot be a geodesic representative of u and is not equal to u .




Figure 10. When $\tilde{\gamma}$ meets itself but glances off the annulus A , we obtain a geodesic loop homotopic to the boundary component u .

If $\tilde{\gamma}|_{[d,d+\varepsilon)}$ enters A for small ε , then we claim there is an immersed q -geodesic bigon cobounded by arcs of $\tilde{\gamma}$, which contradicts nonpositive curvature. Indeed, thicken A slightly to an annulus A' , and let $t > d$ be the smallest value for which $\tilde{\gamma}(t)$ meets $\partial A' \setminus u$. Consider the lifts of $\tilde{\gamma}|_{[a,t]}$ to the universal cover $\tilde{A}'$ of A' . This cover is an infinite strip, and the lifts form a $\mathbb{Z}$ -periodic sequence of arcs connecting the two boundary components. One may order the lifts by their endpoints on (either) boundary component, and then we see that, since each lift crosses the lift above it at a preimage of $\tilde{\gamma}(d)$, two consecutive lifts must intersect at least twice. This produces an immersed bigon downstairs, and our contradiction (see Figure 11).

We conclude that the concatenation σ of $\tilde{\gamma}|_{[a,d]}$ followed by $\tilde{\gamma}|_{[a,c]}$ traversed in the opposite direction is an essential arc of Y_q which is homotopic to an embedded arc (Figure 12). Moreover, since x is embedded, σ meets no singularity interior to Y_q more than twice. Hence, if we let $l \geq 0$ be the number of fundamental domains in $[a, d]$, defined




Figure 11. On the right is the annulus A formed when γ crosses itself at its first intersection point. Its thickening A' is indicated by the dotted circle. At left is the universal cover of A' where one can see the bigon between translates of the lift of $\tilde{\gamma}$.




Figure 12. The concatenation σ is formed from the self-intersection of $\tilde{\gamma}$. The inner circle represents u or the puncture.

exactly as above, then $\ell_q(\sigma) \leq 2\ell_q(\tilde{\gamma}|_{[a,d]}) \leq 2(l+1) \cdot \ell_q(\partial Y)$. Using the same reasoning as in the previous case, we conclude that $j_Y(\lambda^\pm, \sigma) \leq 8(l+1)$. Using the fact that σ meets no singularity of Y_q more than twice, we have $l \leq 4|\chi(Y)|$ and hence

$$j_Y(\lambda^\pm, \sigma) \leq 8(4 \cdot |\chi(Y)| + 1).$$

This time we apply equation (2.2) to conclude that $d_Y(\lambda^\pm \sigma) \leq 18$ and so

$$d_Y(\lambda^-, \lambda^+) \leq 36.$$

Case 2: $d = b$. This final case is handled just like case 1a, except that $\sigma = \tilde{\gamma}$ is an essential arc with embedded interior, rather than an essential loop.

This completes the proof of Theorem 4.2. ■

Theorem 4.1 now follows as a corollary.

Proof of Theorem 4.1. Suppose that Y does not cover a subsurface of S . Let α, β be two curves with nontrivial projection to Y , and let $(\lambda_n^\alpha)_{n \geq 0}$ and $(\lambda_n^\beta)_{n \geq 0}$ be any sequences of filling laminations such that α is contained in the Hausdorff limit of λ_n^α and β is contained in the Hausdorff limit of λ_n^β . For example, if $f: S \rightarrow S$ is any pseudo-Anosov homeomorphism, then we can take the stable and unstable laminations of $T_\alpha^n \circ f \circ T_\beta^n$ which are pseudo-Anosov for all but finitely many $n \in \mathbb{Z}$. Here T_γ denotes the Dehn twist about the curve γ .

Now let $q_n = q(\lambda_n^\alpha, \lambda_n^\beta)$ be the holomorphic quadratic differential whose vertical and horizontal foliations are determined by λ_n^α and λ_n^β , respectively. By Theorem 4.2, then $d_Y(\lambda_n^\alpha, \lambda_n^\beta) \leq 36$ for each $n \in \mathbb{Z}$. However, for large enough $n \geq 0$, $\pi_Y(\alpha) \subset \pi_Y(\lambda_n^\alpha)$ and $\pi_Y(\beta) \subset \pi_Y(\lambda_n^\beta)$, and we conclude that $d_Y(\alpha, \beta) \leq 36$. Since α and β were arbitrary curves with nontrivial projection to Y , we conclude that $\text{diam}_Y(\mathcal{A}(S)) \leq 38$. ■

5. Uniform bounds in the veering triangulation

In this section, we produce two estimates relating sections of the veering triangulation to subsurface projections in the *original* surface X . In Proposition 5.1, we show that for a τ -compatible subsurface Y of X , the π_Y^τ -projections of the top and bottom sections of $T(\partial_\tau Y)$ are near the projections of λ^+ and λ^- , respectively, to $\mathcal{A}(Y)$. In Lemma 5.2, we show that if sections T_1 and T_2 have sufficiently far apart π_Y^τ -projections to $\mathcal{A}(Y)$, then they must differ by at least $|\chi(Y)|$ tetrahedron moves in the veering triangulation. Both estimates will be important ingredients in the proof of Theorem 6.1.

5.1. Top/bottom of the pocket and distance to $\lambda^\pm$

Let $Y \subset X$ be a τ -compatible subsurface as in Section 2.4. Since $\partial_\tau Y$ is a collection of edges of τ with disjoint interiors in X , Lemma 3.3 gives that $T(\partial_\tau Y) \neq \emptyset$. Let $T^\pm \in T(\partial_\tau Y)$ denote the top and bottom sections containing $\partial_\tau Y$.

The following proposition is analogous to [20, Proposition 6.2] in the fully-punctured setting. However, more work is needed here to relate the projection of T^+ to the projection of λ^+ in the curve and arc graph of Y .

Proposition 5.1 (Compatibility with τ and distance to $\lambda^\pm$). *Let $T^\pm$ be the top and bottom sections in $T(\partial_\tau Y)$. For any section $Q \geq T^+$, $d_Y(Q, \lambda^+) \leq D + 1$ and for any $Q \leq T^-$, $d_Y(Q, \lambda^-) \leq D + 1$.*

Recall that $d_Y(Q, \lambda^-)$ stands for $d_Y(\pi_Y^\tau(Q), \pi_Y(\lambda^-))$.

Proof. We begin by remarking that if $\text{int}_\tau(Y)$ contains no singularities of q other than punctures (i.e., if $\mathcal{P}_Y = \text{sing}(q) \cap \text{int}_\tau(Y)$ in $\bar{X}$), then the argument from [20, Proposition 6.2] carries through and gives a better constant. This includes the situation where Y is an annulus.

In the general (nonannular) case, we show that there is an embedded edge path p in $\Pi_*(Q) = \Delta_Q$ which projects to an essential arc of $\text{int}_\tau(Y)$ and is isotopic to a properly embedded arc of $\text{int}_\tau(Y) \cap \lambda^+$. Together with Lemma 3.6, this shows that

$$d_Y(Q, \lambda^+) = d_Y(\pi_Y^\tau(Q), \lambda^+) = \text{diam}_Y(\pi_Y^\tau(Q)) + \text{diam}_Y(\lambda^+) \leq D + 1$$

as required (the proof for λ^- is identical).

Let $T_0, T_1, T_2, \dots$ be any sequence of sections through upward tetrahedron moves starting with $T_0 = T^+$ such that $Q = T_j$ for some $j \geq 0$ (Lemma 3.4), and let $\Delta_i = \Pi_*(T_i)$ be the corresponding τ -triangulations of X . Note that Δ_{i+1} is obtained from Δ_i by a single diagonal exchange.

Since $T_0 \in T(\partial_\tau Y)$, Δ_0 contains a subcomplex Y_0 which triangulates the image of Y_τ under the covering $X_Y \rightarrow X$. We inductively construct a sequence of subcomplexes $Y_0 \supset Y_1 \supset \dots \supset Y_{n+1}$ of $\Delta_0 \subset X$ satisfying certain conditions. For this, we say that each Y_i comes with a *boundary* $\partial_\tau Y_i$, which is defined inductively below, and we set $\text{int}_\tau(Y_i) = Y_i \setminus \partial_\tau Y_i$. For Y_0 , we have $\partial_\tau Y_0 = \partial_\tau Y$ so that $\text{int}_\tau(Y) = Y_0 \setminus \partial_\tau Y_0$ is isotopic to Y .

To construct Y_{i+1} from Y_i observe that the upward exchange from Δ_i to Δ_{i+1} either occurs along an edge not meeting Y_i , in which case we set $Y_{i+1} = Y_i$, or it must occur along an edge e_i of $\partial_\tau Y_i$. Otherwise, the edge e_i lies in the interior of $Y_i \subset Y_0$ where it is wider (with respect to q) than the other edges in its two adjacent triangles. Hence, the same must be true in the triangulation Δ_0 , and we see that $\Pi^*(e_i)$ is upward flippable in T^+ . This gives a section $T' \in T(\partial_\tau Y)$ with $T^+ < T'$, contradicting the definition of T^+ .

Let Q_i be the quadrilateral in Δ_i whose diagonal is e_i . Writing Q_i as two triangles adjacent along e_i , at least one of them, call it D_i , is contained in Y_i , as in Figure 13. If D_i is the only triangle in Y_i , set $Y_{i+1} = Y_i \setminus (\text{int}(D_i) \cup e_i)$ and $\partial_\tau Y_{i+1} = (\partial_\tau Y_i \cup \partial D_i) \setminus e_i$.




Figure 13. The first step: The upward diagonal exchange along the boundary and removing the triangle D_i . Vertices lie in $\text{sing}(q) \cup \mathcal{P}$.

If the other triangle D'_i of Q_i is also in Y_i , set $Y_{i+1} = Y_i \setminus \text{int}(Q_i)$ and $\partial_\tau Y_{i+1} = \partial_\tau Y_i \cup \partial Q_i \setminus e_i$.

Let v_i be the vertex of D_i opposite e_i (and v'_i the vertex of D'_i opposite e_i if we are in the second case). If v_i (or v'_i) is contained in $\partial_\tau Y_i \cup \mathcal{P}$, or $v_i = v'_i$, then we set $n = i$, so that Y_{n+1} is the last step of the construction.




Figure 14. An example sequence of Y_i (indicated in blue). A component K of $\text{int}_\tau(Y_0) \setminus Y_i$ is indicated with its λ^+ foliation in grey. Here σ_K , dotted, is the edge e_0 .

The construction has the following properties for $i \leq n + 1$ (see Figure 14):

- (1) Y_i is a subcomplex of each of $\Delta_0, \Delta_1, \dots, \Delta_i$. Moreover, the diagonal exchange from Δ_{i-1} to Δ_i replaces an edge e_{i-1} in either $\partial_\tau Y_{i-1}$ or in the complement of Y_{i-1} .
- (2) Each component K of $\text{int}_\tau(Y_0) \setminus Y_i$ is a disk foliated by leaves of λ^+ . Moreover, ∂K is composed of two arcs p_K and σ_K , where σ_K is a τ -edge in $\partial_\tau Y$ and p_K is a path in $\partial_\tau Y_i$, and each leaf of the foliation of K meets both p_K and σ_K at its endpoints.
- (3) For each component K of $\text{int}_\tau(Y_0) \setminus Y_i$ and each interior vertex v of p_K , there is at least one edge e of Δ_i entering K from v , and every such edge crosses σ_K from top to bottom, so that $e > \sigma_K$.

First note that Y_0 satisfies the properties for $i = 0$ and that property (1) holds for all $i \geq 0$ by construction.

Assume property (2) holds for $i \leq n$, and let us prove it for $i + 1$. Let D_i be the disk removed to obtain Y_{i+1} , and let p_{D_i} be the other two sides of ∂D_i .

First suppose e_i is in $\partial_\tau Y_0$ (which in particular holds for $i = 0$). The two edges in p_{D_i} cannot be in the boundary of any other component K' of $\text{int}_\tau(Y_0) \setminus Y_i$, because then for some $j < i$, there would have been a D_j whose third vertex v_j was either a puncture or on $\partial_\tau Y_0$, which implies $j > n$. Thus D_i itself is a component of $\text{int}_\tau(Y_0) \setminus Y_{i+1}$, and (2) evidently holds, with $\sigma_{D_i} = e_i$.

Now if e_i is not in $\partial_\tau Y_0$, it must lie in p_K for some component K of $\text{int}_\tau(Y_0) \setminus Y_i$. We have again that p_{D_i} cannot share edges with any other component of $\text{int}_\tau(Y_0) \setminus Y_i$, and so we obtain a component K' of $\text{int}_\tau(Y_0) \setminus Y_{i+1}$ by adjoining $D_i \setminus p_{D_i}$ to K . The boundary path $p_{K'}$ is obtained from p_K by replacing e_i by p_{D_i} , and the λ^+ foliation of K extends across e_i to K' . The edge $\sigma_{K'}$ is just σ_K .




Figure 15. The proof of property (3).

Now we consider property (3). Breaking up into cases as in the proof of (2), consider first the case that e_i is in $\partial_\tau Y_0$ (Figure 15 (a)). Let Q_i be the quadrilateral defining the move $\Delta_i \rightarrow \Delta_{i+1}$. Then the new edge e'_i is the other diagonal of Q_i and must cross e_i . This shows that (3) holds for the new component $K = \text{int}(D_i)$ of $\text{int}_\tau(Y_0) \setminus Y_{i+1}$.

If e_i is in p_K for some component K of $\text{int}_\tau(Y_0) \setminus Y_i$, consider Q_i again (Figure 15 (b)), and let D'_i be the complementary triangle to D_i in Q_i . The two complementary edges to e_i in $\partial D'_i$ must, by inductively applying (3), cross σ_K . It follows that the new edge e'_i also crosses σ_K , thus proving (3) for $i + 1$. Here, we are using the general fact that if an edge of a τ -triangle crosses σ from top to bottom, then so does the tallest edge of that triangle.

Finally, suppose e_i is outside of Y_i . If Q_i is not adjacent to Y_i at all, then nothing changes and (3) continues to hold. If Q_i shares a vertex w with p_K for some component K of $\text{int}_\tau(Y_0) \setminus Y_i$, we have a configuration like Figure 15 (c) or (d). In (c), w is on e_i . The new edge e'_i , which connects to another vertex on p_K , must therefore also cross σ_K (by transitivity of $<$). Moreover, w is still adjacent to one of the boundary edges of Q_i which also passes through σ_K . Thus (3) is preserved at all vertices of p_K . In (d), w is the vertex meeting e'_i and again we must have e'_i crossing σ_K . This concludes the proof of properties (1)–(3).

The sequence terminates with the diagonal exchange from Δ_n to Δ_{n+1} along an edge e_n whose associated triangle D_n has its opposite vertex v_n in either $\mathcal{P}$ or the boundary of Y_n . There are either one or two components K of $\text{int}_\tau(Y_0) \setminus Y_{n+1}$ whose boundary path p_K contains v_n . Examples of these two cases are shown in Figure 16.

For each such component K , let l_K be the leaf of λ^+ adjacent to v_n and exiting K through $\sigma_K \subset \partial_\tau Y_0$, as in property (2). Let q_K be a path along p_K from v_n to one of the endpoints of p_K . When there is only one component K , let $l = l_K$ and $q = q_K$. If there are two components K, K' , let $l = l_K \cup l_{K'}$ and $q = q_K \cup q_{K'}$.

By Lemma 3.6, l is an embedded essential arc of $\text{int}_\tau Y$, and q and l define the same element of $\mathcal{A}(Y)$.




Figure 16. Two examples of Y_{n+1} and the corresponding λ^+ arcs and Δ_i arcs. In the first case, the point v_n is a puncture so only one leaf l_K and path q_K is needed. In the second case, v_n lies on the boundary of two components of $\text{int}_\tau(Y_0) \setminus Y_{n+1}$.

Recall that $\Pi_*(Q) = \Delta_j$. If $j \leq n + 1$, we know that Y_{n+1} is a subcomplex of Δ_j by (1), and hence we obtain a bound

$$d_Y(\Delta_j, \lambda^+) \leq D + 1.$$

If $j > n + 1$, we must consider what happens after further transitions. By (3), for each component K of $\text{int}_\tau(Y_0) \setminus Y_{n+1}$ adjacent to v_n , there is an edge f_{n+1} of Δ_{n+1} adjacent to v_n , and passing through K and exiting through σ_K . In particular, the defining rectangle R_{n+1} of f_{n+1} passes through the rectangle of σ_K from top to bottom. We claim that there is such an f_i and R_i for each $i \geq n + 1$. If this holds for i , and the move $\Delta_i \rightarrow \Delta_{i+1}$ does not replace f_i , then we can let $R_{i+1} = R_i$ and $f_{i+1} = f_i$. If f_i is replaced, there is a quadrilateral Q_i in Δ_i of which f_i is a diagonal, and one of the two edges of ∂Q_i adjacent to v_n must also cross the σ_K rectangle from top to bottom. (See Figure 17.)

Therefore, we can use these edges of Δ_j (either one or two depending on the number of components K adjacent to v_n) to give an essential path in Δ_j which gives the same element of $\mathcal{A}(Y)$ as the leaf path l . We conclude again that

$$d_Y(\Delta_j, \lambda^+) \leq D + 1,$$

which completes the proof of Proposition 5.1. ■

5.2. Sweeping (slowly) through pockets

The following lemma states that in order to move a definite distance in the curve graph of $Y \subset X$ a certain number of edges, linear in the complexity of Y , need to be flipped. It will be needed for the proof of Theorem 6.1.




Figure 17. For $i > n$, we always find an edge in Δ_i starting at v_n and exiting K through σ_K .

Lemma 5.2 (Complexity slows progress). *Suppose that $T_1, T_2 \in T(\partial_\tau Y)$ are connected by no more than $|\chi(Y)|$ diagonal exchanges through sections of $T(\partial_\tau Y)$. Then*

$$d_Y(T_1, T_2) \leq 2D.$$

Proof. We begin with the following claim.

Claim 1. *Let $\bar{Y}$ be a compact surface containing a finite (possibly empty) set $\mathcal{P}_Y \subset \text{int}(\bar{Y})$, let $Y = \bar{Y} \setminus \mathcal{P}_Y$, and assume that $\chi(Y) \leq -1$. Let $\bar{\Delta}$ be a triangulation of $\bar{Y}$ whose vertex set contains $\mathcal{P}_Y$, and consider the properly embedded graph $\Delta = \bar{\Delta} \setminus \mathcal{P}_Y \subset Y$. If E is a collection of nonboundary edges of Δ with $|E| \leq |\chi(Y)|$, there is $c \in \mathcal{A}(Y)$ that is nearly simple in Δ and traverses no edge of E .*

Recall from Section 2.1.3 that c is nearly simple in Δ if it is properly homotopic to path or curve in Δ that visits no vertex more than twice. Note that the conclusion of the claim is equivalent to the statement that the proper graph $\Delta^{(1)} \setminus \text{int}(E)$ is essential in Y .

Applying the claim to $\text{int}_\tau(Y)$ gives $c \in \mathcal{A}(Y)$ which is nearly simple in the graph $\text{cl}_X(\Pi_*((T_1 \cap T_2)^{(1)}))$. Hence, $d_Y(T_1, T_2) \leq 2D$.

Now it suffices to prove the claim.

Proof of Claim 1. First, we blow up the punctures $\mathcal{P}_Y$ to boundary components. That is, each $v \in \mathcal{P}_Y$ is replaced by a subdivided circle S_v containing a vertex for each edge adjacent to v . Continue to call the resulting surface Y and note that Δ induces a natural cell structure on Y , which we continue to denote by Δ , made up of n -gons with $n \leq 6$. Obviously, $\chi(Y)$ is unchanged.

Considering E as a possibly disconnected subgraph of Δ , let E' be the edges of E remaining after removing the components of E which are contractible in Y and do not

meet ∂Y . Hence, every component of $Y \setminus E'$ is π_1 -injective. We first claim that some component of $Y \setminus E'$ is neither a disk nor a boundary parallel annulus. Letting $G = E' \cup \partial Y$, we have $\chi(Y) = \chi(Y \setminus G) + \chi(G)$ since $Y \setminus G$ is adjoined to G along circles. Now $\chi(G) = \chi(\partial Y) + v - e = v - e$, where v is the number of vertices in $E' \setminus \partial Y$ and $e = |E'|$. If $Y \setminus E'$ consists of only $d \geq 0$ disks and $a \geq 0$ annuli, then we have $\chi(Y \setminus G) = \chi(Y \setminus E') = d$, and so $\chi(Y) = d + v - e$. Since an annulus must have a vertex of $E' \setminus \partial Y$ in its boundary, we note that if $d = 0$, then $v > 0$, and hence $d + v \geq 1$. We thus have $\chi(Y) \geq 1 - |E'|$, and so $|\chi(Y)| \leq |E'| - 1 \leq |E| - 1$, a contradiction.

Now let U' be some component of $Y \setminus E'$ which is neither a disk nor a boundary parallel annulus. Hence, there is an essential simple closed curve γ' of Y contained in U' . As $U = U' \setminus E$ has a component corresponding to U' minus a collection of disks, γ' is homotopic to a simple curve γ in U . Let c be a cell of Δ which γ crosses. Since c is a blown-up triangle, the edges that γ crosses, which are in the complement of E , are connected along either vertices of c or arcs of ∂Y (which are also not in E). Thus γ can be deformed to a curve in Δ that does not traverse the edges of E . This shows that the proper graph $\Delta^{(1)} \setminus \text{int}(E)$ is essential in Y and the claim follows. ■

This completes the proof of Lemma 5.2. ■

6. Uniform bounds in a fibered face

In this section, we prove the first main theorem of the paper. For a surface Y , recall that $|\chi'(Y)| = \max\{|\chi(Y)|, 1\}$.

Theorem 6.1 (Bounding projections for M). *Let M be a hyperbolic fibered 3-manifold with fibered face $\mathbf{F}$. Then for any fiber S contained in $\mathbb{R}_+\mathbf{F}$ and any subsurface Y of S ,*

$$|\chi'(Y)| \cdot (d_Y(\lambda^-, \lambda^+) - 16D) \leq 2D|\mathbf{F}|,$$

where $|\mathbf{F}|$ is the number of tetrahedra of the veering triangulation associated to $\mathbf{F}$.

The proof of Theorem 6.1 requires the construction of an embedded subcomplex of $(\overset{\circ}{M}, \tau)$ corresponding to Y whose size is roughly $|\chi'(Y)| \cdot d_Y(\lambda^-, \lambda^+)$.

6.1. Pockets and the approach from the fully-punctured case

Let us assume from now on that $d_Y(\lambda^-, \lambda^+)$ is sufficiently large for Y to be τ -compatible by Theorem 2.6. We first describe an approach directly extending the argument from [20], and explain where it runs into trouble.

Recall from Section 3.1 the definition of $T(\partial_\tau Y)$, and its top and bottom sections T^+ and T^- . For any two sections $T_1, T_2 \in T(\partial_\tau Y)$, we have the region $U(T_1, T_2)$ in $\mathcal{N}$ between them. The closure of the open subsurface $\text{int}_\tau(Y)$ in X is the subspace $\mathcal{R}_Y$ which is the image of $Y_\tau \subset X_Y$ under the covering map $X_Y \rightarrow X$. Using this, we further define

$$U_Y(T_1, T_2) = U(T_1, T_2) \cap \Pi^{-1}(\mathcal{R}_Y) \quad (6.1)$$

to be the region between T_1 and T_2 that lies above Y . Note that $U_Y(T_1, T_2)$ is a subcomplex of $\mathcal{N}$. We sometimes call this the *pinched pocket for Y* (between T_1 and T_2). We let

$$U_Y = U_Y(T^-, T^+)$$

denote the *maximal pocket for Y* .

By Proposition 5.1, $d_Y(T^-, T^+)$ is close to $d_Y(\lambda^-, \lambda^+)$. Thus when these are sufficiently large, by Lemma 5.2 we obtain a lower bound on the number of transitions between T^- and T^+ in $\mathcal{N}$ that project to $\mathcal{R}_Y$, and in particular a lower bound on the number of tetrahedra in U_Y . If U_Y were to embed in $\overset{\circ}{M}$ (equivalently, if U_Y were disjoint from all its translates by Φ), this would give us what we need. It is not in general true, so instead we must restrict to a suitable sub-region of U_Y .

In the fully-punctured case, we can use the fact that, since every edge of τ represents an element of $\mathcal{A}(S)$, any intersection between U_Y and $\Phi^k(U_Y)$ would project to a nonempty $\pi_Y(\varphi^k(\partial Y))$. Lemma 2.2 implies that, for $k > 0$, whenever $\pi_Y(\varphi^k(\partial Y))$ is nonempty, it is close to $\pi_Y(\lambda^-)$. Now by restricting U_Y to a subcomplex V_Y whose top and bottom surfaces are sufficiently far in $\mathcal{A}(Y)$ from λ^+ and λ^- , and applying this argument to V_Y , we see that V_Y cannot meet $\Phi^k(V_Y)$ at all.

In the general situation, since some singularities of q are not punctures, not every collection of τ -edges is essential and we are faced with the possibility that U_Y and $\Phi^k(U_Y)$ can intersect in large but homotopically inessential subcomplexes whose location is hard to control. This is the main difficulty.

6.2. Isolation via φ -sections

We begin, therefore, with the following construction. Let T_0 denote a φ -section (Section 3.1), so that $\varphi^k(T_0) \leq T_0$ for $k > 0$. Define

$$N = N_Y = \min\{i > 0: \varphi^i(Y) \text{ overlaps } Y\}. \quad (6.2)$$

Now consider the (possibly empty) region

$$R(T_0) = \text{int}(U(\Phi^N(T_0), T_0)) \cap \text{int}(U_Y). \quad (6.3)$$

It easily satisfies the following embedding property.

Proposition 6.2. *The restriction of the covering map $\mathcal{N} \rightarrow \overset{\circ}{M}$ to $R(T_0)$ is an embedding.*

Proof. We must show that $\Phi^i(R(T_0))$ is disjoint from $R(T_0)$ for all $i > 0$ (the case for $i < 0$ immediately follows). For $i \geq N$, this is taken care of by the first term of the intersection since T_0 was chosen to be a φ -section.

For $0 < i < N$, we have that $\text{int}_\tau(Y)$ and $\text{int}_\tau(\varphi^i(Y))$ are disjoint since the surfaces have no essential intersection, using Lemma 3.6. Therefore, the interiors of U_Y and $\Phi^i(U_Y)$ are disjoint as well. ■

What remains now is to choose T_0 so that a lower bound on $d_Y(\lambda^-, \lambda^+)$ implies a lower bound on the number of tetrahedra in $R(T_0)$. That is, we want to view $R(T_0)$ as the interior of a “pocket” between two sections, whose projections to $\mathcal{A}(Y)$ are close to $\lambda^\pm$. For this we will need to describe $R(T_0)$ from a different point of view.

6.3. Y -projections of sections

For any section T of $\mathcal{N}$, there is a corresponding section $T^Y \in T(\partial_\tau Y)$ obtained by pushing T below T^+ and above T^- . More formally,

$$T^Y = T^+ \wedge (T^- \vee T) = T^- \vee (T^+ \wedge T).$$

To see what this does it is helpful to consider it along the fibers of Π which we recall are oriented lines. For any $[a, b] \subset \mathbb{R}$, the map $x \mapsto b \wedge (a \vee x)$ is simply retraction of $\mathbb{R}$ to $[a, b]$, and is equal to $x \mapsto a \vee (b \wedge x)$.

Now we can use this projection to extend the notation $U_Y(T_1, T_2)$ (defined in (6.1)) to sections which are not necessarily in $T(\partial_\tau Y)$ by setting

$$\hat{U}_Y(T_1, T_2) = U_Y(T_1^Y, T_2^Y) \subset U_Y.$$

This construction is related to the region $R(T_0)$ defined in (6.3) by the following lemma.

Lemma 6.3. *We have*

$$\text{int}(\hat{U}_Y(T_1, T_2)) = \text{int}(U(T_1, T_2)) \cap \text{int}(U_Y). \quad (6.4)$$

Proof. First consider what happens in each Π -fiber, which is just a statement about projections in $\mathbb{R}$: If $J = [s, t]$ is an interval in $\mathbb{R}$, we have, as above, the retraction $\pi_J(u) = s \vee (t \wedge u)$, and for any other interval I , we immediately find

$$\text{int}(\pi_J(I)) = \text{int}(I) \cap \text{int}(J). \quad (6.5)$$

To apply this to our situation, note first that both the left- and right-hand sides of equality (6.4) are contained in $\Pi^{-1}(\text{int}_\tau(Y))$. This is because U_Y is in $\Pi^{-1}(\mathcal{R}_Y)$ and in $U(T^-, T^+)$, which means that for each $x \in \mathcal{R}_Y \setminus \text{int}_\tau(Y) = \partial_\tau Y$, $U_Y \cap \Pi^{-1}(x)$ is a single point, and hence not in the interior.

Now since Π is a fibration, for any two sections T, T' and region Z of the form $U(T, T') \cap \Pi^{-1}(\Omega)$ where Ω is open, we have

$$\text{int}(Z) \cap \Pi^{-1}(x) = \text{int}(Z \cap \Pi^{-1}(x)).$$

For $x \in \text{int}_\tau(Y)$, applying this to the left-hand side of (6.4) we see that

$$\text{int}(\hat{U}_Y(T_1, T_2)) \cap \Pi^{-1}(x) = \text{int}(\hat{U}_Y(T_1, T_2) \cap \Pi^{-1}(x)).$$




Figure 18. The section T , its intersection with $U(T^-, T^+)$, and the decomposition $T = K_+ \cup K_0 \cup K_-$.

On the right-hand side, we obtain

$$\text{int}(U(T_1, T_2)) \cap \Pi^{-1}(x) = \text{int}(U(T_1, T_2) \cap \Pi^{-1}(x))$$

and

$$\text{int}(U_Y) \cap \Pi^{-1}(x) = \text{int}(U_Y \cap \Pi^{-1}(x)).$$

This reduces the equality to a fiberwise equality, where it follows from (6.5). $\blacksquare$

6.4. Relation between T^Y and π_Y

A crucial point now is to show that the operation $T \rightarrow T^Y$ does not alter the projection to $\mathcal{A}(Y)$ by too much.

Proposition 6.4. *For all sections T of $\mathcal{N}$, $d_Y(T, T^Y) \leq 4D$.*

Here $d_Y(T, T^Y)$ is meant in the sense of (3.3).

Proof. Write T as a union of three subcomplexes, $T = K_+ \cup K_0 \cup K_-$, where

$$K_0 = T \cap U(T^-, T^+), \quad K_+ = T \cap (T \vee T^+), \quad K_- = T \cap (T \wedge T^-).$$

See Figure 18. Since $T^+, T^- \in T(\partial_\tau Y)$, $\Pi_*(K_0)$ is a subcomplex of τ -edges in X that do not cross $\partial_\tau Y$.

Lemma 6.5. *If $\pi_Y^\tau(K_+)$ and $\pi_Y^\tau(K_-)$ are empty, then $\Pi_*(K_0)$ contains a punctured spine for Y .*

A punctured spine for Y is a subspace which is a retract of Y minus a union of disjoint disks. In other words, the conclusion of Lemma 6.5 implies that every essential curve in Y is homotopic into $\Pi_*(K_0)$. In particular, $\pi_Y^\tau(K_0)$ is nonempty and so is every $\pi_Z^\tau(K_0)$ for Z a τ -compatible subsurface that overlaps with Y .

Proof of Lemma 6.5. The statement that $\pi_Y^\tau(K_+)$ is empty means that, after projecting K_+ by Π into X and intersecting with $\text{int}_\tau(Y)$, we obtain components which are inessential subcomplexes, meaning they do not contain any essential curves or proper

arcs. (When Y is an annulus, this means that no τ -edge from K_+ joins opposite sides of the open annulus $\text{int}_\tau(Y)$.)

Now $K_+ \setminus K_0$ and $K_- \setminus K_0$ are open in T and disjoint, so their projections to X intersect $\text{int}_\tau(Y)$ in a collection of disjoint open sets each of which is inessential in the above sense. It follows that the complement of these open sets, which is $\Pi(K_0) \cap \text{int}_\tau(Y)$, intersects every essential curve and proper arc in $\text{int}_\tau(Y)$. Hence it contains a punctured spine. ■

If $T \cap U_Y$ projects to an essential subcomplex of Y_τ , then the proposition follows since T and T^Y both contain $T \cap U_Y$.

If $T \cap U_Y$ is inessential, then, by Lemma 6.5, at least one of $\pi_Y^\tau(K_\pm)$ is nonempty. Suppose $\pi_Y^\tau(K_+)$ is nonempty.

Since $K_+ = T \cap (T \vee T^+)$, we immediately have

$$d_Y(T, T \vee T^+) \leq D.$$

Since $T^+ \leq T \vee T^+$, Proposition 5.1 tells us that

$$d_Y(T \vee T^+, \lambda^+) \leq D.$$

Thus

$$d_Y(T, \lambda^+) \leq 2D.$$

Now note that $T^Y \cap T^+$ is equal to the part of T^+ lying below T , hence its Π -projection to X has the same image as K_+ . It follows that $\pi_Y^\tau(T^Y \cap T^+)$ is nonempty, therefore

$$d_Y(T^Y, T^+) \leq D.$$

Since Proposition 5.1 again gives us

$$d_Y(T^+, \lambda^+) \leq D,$$

we combine all of these to conclude

$$d_Y(T, T^Y) \leq 4D.$$

This completes the proof of Proposition 6.4. ■

Remark 6.6. Note that the proof of Proposition 6.4 shows that if $\pi_Y^\tau(K_+)$ is nonempty, then $d_Y(T, \lambda^+) \leq 2D$. The corresponding statement also holds if $\pi_Y^\tau(K_-)$ is nonempty.

6.5. Finishing the proof of Theorem 6.1

We assume that $d_Y(\lambda^-, \lambda^+) \geq 10D$. To complete the argument, we choose a φ -section T_0 with the property that $3D \leq d_Y(T_0, \lambda^+) \leq 5D$. Such a section exists by Lemma 3.5 and Proposition 5.1. Let $N = N_Y$ be as in (6.2).

Lemma 6.7. *With notation as above,*

$$d_Y(\Phi^N(T_0), \lambda^-) \leq 2D + 11.$$

Proof. From Lemma 2.2, we have

$$d_Y(\varphi^N(\partial Y), \lambda^-) \leq 4.$$

From Proposition 3.7, we have that since $\varphi^N(Y)$ overlaps with Y ,

$$d_Y(\partial\varphi^N(Y), \pi_Y^\tau(\partial_\tau\varphi^N(Y))) \leq 7.$$

Since $\varphi^N(\partial_\tau Y) = \partial_\tau(\varphi^N(Y))$, we have

$$d_Y(\varphi^N(\partial Y), \pi_Y^\tau(\varphi^N(\partial_\tau Y))) \leq 7.$$

Let K_0 , K_+ and K_- be subcomplexes of T_0 as in the proof of Proposition 6.4. Since $d_Y(T_0, \lambda^+) \geq 3D$ and $d_Y(T_0, \lambda^-) \geq 10D - 5D$ by choice of T_0 and the assumption on $d_Y(\lambda^-, \lambda^+)$, the proof of Proposition 6.4 (see Remark 6.6) tells us that both $\pi_Y^\tau(K^+)$ and $\pi_Y^\tau(K^-)$ are empty. So by Lemma 6.5, $T_0 \cap U_Y$ must contain a punctured spine for Y .

Now since $\varphi^N(Y)$ intersects Y essentially, it must be that $\pi_Y^\tau(\Phi^N(T_0 \cap U_Y))$ is nonempty. Since the triangulation of T_0^Y contains both $T_0 \cap U_Y$ and $\partial_\tau Y$, we have

$$d_Y(\pi_Y^\tau(\Phi^N(T_0 \cap U_Y)), \pi_Y^\tau(\Phi^N(\partial_\tau Y))) \leq D.$$

Combining these inequalities and observing that $T_0 \cap U_Y \subset T_0$, we obtain the desired inequality. $\blacksquare$

We now define the *isolated pocket* for Y to be

$$V = V_Y = \hat{U}_Y(\Phi^N(T_0), T_0). \quad (6.6)$$

Lemma 6.3 implies that the interior of V is equal to $R(T_0)$, and Proposition 6.2 therefore implies the following corollary.

Corollary 6.8. *The covering map $\mathcal{N} \rightarrow \mathring{M}$ embeds $\text{int}(V)$ in $\mathring{M}$.*

Thus we can complete the proof of Theorem 6.1 with the following proposition.

Proposition 6.9. *The isolated pocket for Y satisfies*

$$|V| \geq \frac{1}{2D} |\chi'(Y)| \cdot (d_Y(\lambda^-, \lambda^+) - 16D).$$

Proof. By definition,

$$V = U_Y(\Phi^N(T_0)^Y, T_0^Y) \subset U_Y.$$

Moreover, we claim that the sections defining this region satisfy the following:

$$d_Y(T_0^Y, \lambda^+) \leq 9D \quad \text{and} \quad d_Y(\Phi^N(T_0)^Y, \lambda^-) \leq 7D. \quad (6.7)$$

The first inequality follows directly from the assumption that $d_Y(T_0, \lambda^+) \leq 5D$ and Proposition 6.4. The second inequality follows from Lemma 6.7 and another application of Proposition 6.4. (Here we have used that $11 \leq D$.)

So we see that $d_Y(\Phi^N(T_0)^Y, T_0^Y) \geq d_Y(\lambda^-, \lambda^+) - 16D$. Hence, to prove the proposition, it suffices to show that

$$\frac{|V|}{|\chi'(Y)|} \geq \frac{d_Y(\Phi^N(T_0)^Y, T_0^Y)}{2D}. \quad (6.8)$$

Set $d = d_Y(\Phi^N(T_0)^Y, T_0^Y)$. If Y is not an annulus, then Lemma 5.2 implies that at least $|\chi'(Y)| \cdot d/2D$ upward tetrahedron moves through tetrahedra of U_Y are needed to connect $\Phi^N(T_0)^Y$ to T_0^Y . If Y is an annulus, the same is true since $|\chi'(Y)| = 1$ and any triangulation of Y has at least two edges joining opposite boundary components. As each of these tetrahedra lie in V by definition, this establishes equation (6.8) and completes the proof. ■

7. The subsurface dichotomy

In this section, we prove the second of our main theorems.

Theorem 7.1 (Subsurface dichotomy). *Let M be a hyperbolic fibered 3-manifold, and let S and F be fibers of M which are contained in the same fibered face. If $W \subset F$ is a subsurface of F , then either W is homotopic through surfaces transverse to the flow to an embedded subsurface $W' \subset S$ of S with*

$$d_{W'}(\lambda^-, \lambda^+) = d_W(\lambda^-, \lambda^+)$$

or the fiber S satisfies

$$9D \cdot |\chi(S)| \geq d_W(\lambda^-, \lambda^+) - 16D.$$

Punctures and blowups

Fix a fibered face $\mathbf{F}$ of M and denote the corresponding veering triangulation of $\mathring{M}$ by τ . Starting with a fiber F of M in the face $\mathbf{F}$, let $\mathring{F}$ be the fully-punctured fiber, that is, $\mathring{F} = F \setminus \text{sing}(q)$. Also let $\mathcal{N}_F$ be the infinite cyclic cover of $\mathring{M}$ corresponding to the fiber $\mathring{F}$ together with its veering triangulation (the preimage of τ), as in Section 3.1.

For any section T of $\mathcal{N}_F$, let $h_{F,T}^\circ: \mathring{F} \rightarrow \mathring{M}$ be the simplicial map obtained by composing the section with the covering map $\mathcal{N}_F \rightarrow \mathring{M}$. We want to describe a natural way to obtain a map $h_{F,T}: F \rightarrow M$ by filling in punctures.

Let $\check{F}$ be the partial compactification of $\mathring{F}$ to a surface with boundary obtained by adjoining the links of ideal vertices in $\text{sing}(q) \setminus \mathcal{P}$, in the simplicial structure on $\mathring{F}$ induced from T . In other words, we add a circle to each puncture. Similarly, let $\check{M}$ be

the manifold with torus boundaries obtained by adding links for the ideal vertices of τ associated to singular orbits of $\text{sing}(q) \setminus \mathcal{P}$. Then $h_{\check{F},T}^\circ$ extends continuously to a proper map $h_{\check{F},T}: \check{F} \rightarrow \check{M}$.

We obtain (a copy of) M from $\check{M}$ by adjoining solid tori to the boundary components, and a copy of F from $\check{F}$ by adjoining disks. Now by construction (since $\check{F}$ comes from puncturing the fiber F of M), the boundary components of $\check{F}$ map to meridians of the tori. It follows that $h_{\check{F},T}^\circ$ can be extended to a map $h_{F,T}: F \rightarrow M$ which maps the disks into the solid tori.

Now assume that W is a τ -compatible subsurface of F . When $T \in T(\partial_\tau W)$, we can restrict this construction to W as follows: First, let $\check{W}_\tau$ be subsurface obtained from $W_\tau \subset F_W$ by puncturing along all singularities. Then the covering $F_W \rightarrow F$ restricts to a map $\check{W}_\tau \rightarrow \check{F}$ which sends the interior of $\check{W}_\tau$ homeomorphically onto $\text{int}_\tau(W) \setminus \text{sing}(q)$. Composing this map with $h_{\check{F},T}^\circ$, we obtain $h_{\check{W},T}^\circ$. (Note that since $T \in T(\partial_\tau W)$, T naturally induces an ideal triangulation of $\check{W}_\tau$.) Restricting the above construction to $\check{W}_\tau$, we obtain $\check{W}$, $h_{\check{W},T}$ and $h_{W,T}$. This is done in such a way that each ideal point of ∂W_τ is replaced with an arc when forming $\check{W}$ and a “half-disk” when forming W .

If S is another fiber in the same face, together with a given section of $\mathcal{N}_S$, we define h_S° , $h_{\check{S}}$ and h_S in the same way. (Since we will not vary the section of $\mathcal{N}_S$, we do not include it in the notation.)

Intersection locus

Now consider the locus $h_{W,T}^{-1}(h_S^\circ(\check{S}))$, which is a τ -simplicial subcomplex of $\check{W}_\tau$. Its completion in $W_\tau \subset F_W$ (with respect to the underlying q -metric on F) is obtained by adjoining points of $\text{sing}(q) \setminus \mathcal{P}$, and we say that this completion is *inessential* if each of its components can be deformed to a point or to a boundary (or puncture) of W . We say it is *essential* if it is not inessential. In other words, the completion is essential if its 1-skeleton is an essential proper graph of W_τ .

Lemma 7.2 (Essential intersection). *Let F and S be fibers in the same face $\mathbf{F}$ of M . Suppose $W \subset F$ is τ -compatible subsurface of F such that $\pi_1(W)$ is not contained in $\pi_1(S)$. Then, for any section T of $\mathcal{N}_F$ in $T(\partial_\tau W)$, the completion of $h_{W,T}^{-1}(h_S^\circ(\check{S}))$ in W_τ is essential.*

Proof. After obtaining the blowups and maps, as above, we first claim that

$$h_{W,T}^{-1}(h_S(S))$$

is essential. The argument for this is similar to [20, Lemma 2.9], although there it was assumed that $h_W = h_{W,T}$ embeds W into M . In details, if the preimage was not essential, then each component is homotopic into a disk, or homotopic into the ends of W . It follows that h_W is homotopic to a map h'_W whose image misses $h_S(S)$ entirely—just precompose with an isotopy of W into itself which lands in the complement of $h_W^{-1}(h_S(S))$.

But if h'_W misses $h_S(S)$, we can conclude that $\pi_1(W)$ is in $\pi_1(S)$: Letting η denote the cohomology class dual to S in M , the fact that h'_W misses $h_S(S)$ implies that η vanishes on $\pi_1(W)$. Hence, if $h_W^{-1}(h_S(S))$ is inessential, then $\pi_1(W) \leq \pi_1(S)$.

Now note that $h_W^{-1}(h_S(S))$ is contained a small neighborhood of the completion of $h_W^{-1}(h_S(S))$. Indeed, each component of the completion of $h_W^{-1}(h_S(S))$ can be obtained from a component of $h_W^{-1}(h_S(S))$ by collapsing the adjoined disks back to singularities. We conclude that the completion of $h_W^{-1}(h_S(S))$ contains an essential component as well. This is what we wanted to prove. ■

Sections of V_W

Assume that $d_W(\lambda^-, \lambda^+) \geq 10D$. Then W has an isolated pocket

$$V_W = U_W(\varphi^N(T)^W, T^W) \subset \mathcal{N}_F,$$

as in equation (6.6). By Corollary 6.8, the restriction $\text{int}(V_W) \rightarrow \mathring{M}$ of the covering $\mathcal{N}_F \rightarrow \mathring{M}$ is an embedding. Now fix a sequence

$$\varphi^N(T)^W = T_0, T_1, \dots, T_n, T_{n+1} = T^W$$

of sections of $T(\partial_\tau W)$ such that $T_i \rightarrow T_{i+1}$ is a tetrahedron move in U_W for $i < n$, and $\Pi_*(T_n)$ and $\Pi_*(T_{n+1})$ restrict to the same triangulation of $\mathring{W}_\tau$ (note that $T_n \rightarrow T_{n+1}$ is not a tetrahedron move, because the triangulations can be different outside W_τ). That such a sequence exists follows from Lemma 3.4.

Using these sections, we define the subcomplex $W_i = T_i \cap U_W$ of T_i . By construction, the tetrahedron between T_i and T_{i+1} lies between W_i and W_{i+1} and so is contained in $\text{int}(V_W)$.

Denote the map h_{W_i, T_i} associated to the section T_i by $h_i: \mathring{W}_\tau \rightarrow M$.

With this setup, we can complete the proof of Theorem 7.1.

Proof of Theorem 7.1. We may assume that $d_W(\lambda^+, \lambda^-) \geq 10D$.

First, suppose that $\pi_1(W)$ is not contained in $\pi_1(S)$. Then by Lemma 7.2, the subcomplex $h_i^{-1}(h_S(S))$ of W_i has essential completion in W_τ for each $0 \leq i \leq n$. Denote the 1-skeleton of the completion of $h_i^{-1}(h_S(S))$ by p_i and note that its π_W^τ -projection is a nontrivial subset of $\mathcal{A}(W)$ whose diameter is bounded by D . Each transition from p_i to p_{i+1} corresponds to a tetrahedron move, where p_i contains the bottom edge and p_{i+1} the top edge. Because the tetrahedra are in V_W , which embeds in $\mathring{M}$, their top edges in $\mathring{M}$ are all distinct (any edge is the top edge of a unique tetrahedron). Since all these edges are in the image of $h_S(S)$, we find that their number is bounded above by $9|\chi(S)|$. Hence, $n \leq 9|\chi(S)|$.

Since p_0 and p_n are in the triangulations associated to the bottom and top of V_W , respectively, we have (equation (6.7))

$$d_W(p_0, p_n) \geq d_W(\lambda^-, \lambda^+) - 16D.$$

For each transition $p_i \rightarrow p_{i+1}$, observe that $\text{diam}_W(p_i \cup p_{i+1}) \leq D$ since the proper graph $p_i \cup p_{i+1}$ of $\text{int}_\tau(W)$ has at most $2|\chi(W)| + 1$ vertices (Corollary 2.4). Combining these facts, we obtain

$$9|\chi(S)| \geq \frac{d_W(\lambda^-, \lambda^+) - 16D}{D},$$

which completes the proof in this case.

Otherwise, $\pi_1(W) \leq \pi_1(S)$, and we finish the proof as in [20, Theorem 1.2] using Theorem 4.2 in place of the special case obtained there. First, recall that by [20, Lemma 2.8], the quantity $d_W(\lambda^-, \lambda^+)$ depends only on the conjugacy class $\pi_1(W) \leq \pi_1(M)$ and the fibered face $\mathbf{F}$. If we lift W to the S -cover of M , we see that projecting along flow lines gives an immersion $W \rightarrow S$ that induces the inclusion $\pi_1(W) \leq \pi_1(S)$ up to conjugation. Since

$$d_W(\lambda^-, \lambda^+) \geq 10D \geq 37,$$

Theorem 4.2 implies that the map $W \rightarrow S$ factors up to homotopy through a finite cover $W \rightarrow W'$ for a subsurface W' of S . Since $\pi_1(W) < \pi_1(F)$, the map

$$\pi_1(W') \rightarrow \pi_1(M)/\pi_1(F)$$

factors through $\pi_1(W')/\pi_1(W)$ which is finite. Since $\pi_1(M)/\pi_1(F) = \mathbb{Z}$, this implies $\pi_1(W')$ maps to the identity in $\pi_1(M)/\pi_1(F)$ so $\pi_1(W') < \pi_1(F)$. But this implies that $\pi_1(W) = \pi_1(W')$. Hence, the cover $W \rightarrow W'$ has degree 1 and so $d_W(\lambda^-, \lambda^+) = d_{W'}(\lambda^-, \lambda^+)$. This completes the proof of Theorem 7.1. ■

Acknowledgments. The authors thank the referee for comments that improved our exposition.

Funding. Y. N. Minsky was partially supported by NSF grants DMS-1610827 and DMS-2005328, and S. J. Taylor was partially supported by NSF grants DMS-1744551 and DMS-2102018 and the Sloan Foundation.

References

- [1] I. Agol, [Ideal triangulations of pseudo-Anosov mapping tori](#). In *Topology and geometry in dimension three*, pp. 1–17, Contemp. Math. 560, American Mathematical Society, Providence, RI, 2011 Zbl [1335.57026](#) MR [2866919](#)
- [2] I. Agol, Comparing layered triangulations of 3-manifolds which fiber over the circle. 2012, MathOverflow discussion, <http://mathoverflow.net/questions/106426>, visited on 21 December 2023
- [3] I. Agol, [The virtual Haken conjecture \(with an appendix by Ian Agol, Daniel Groves and Jason Manning\)](#). *Doc. Math.* **18** (2013), 1045–1087 Zbl [1286.57019](#) MR [3104553](#)
- [4] T. Aougab, [Uniform hyperbolicity of the graphs of curves](#). *Geom. Topol.* **17** (2013), no. 5, 2855–2875 Zbl [1273.05050](#) MR [3190300](#)

- [5] J. Behrstock, M. F. Hagen, and A. Sisto, [Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups](#). *Geom. Topol.* **21** (2017), no. 3, 1731–1804 Zbl 1439.20043 MR 3650081
- [6] J. Behrstock, B. Kleiner, Y. Minsky, and L. Mosher, [Geometry and rigidity of mapping class groups](#). *Geom. Topol.* **16** (2012), no. 2, 781–888 Zbl 1281.20045 MR 2928983
- [7] M. Bestvina, K. Bromberg, and K. Fujiwara, [Constructing group actions on quasi-trees and applications to mapping class groups](#). *Publ. Math. Inst. Hautes Études Sci.* **122** (2015), 1–64 Zbl 1372.20029 MR 3415065
- [8] M. Bestvina and M. Feighn, [Subfactor projections](#). *J. Topol.* **7** (2014), no. 3, 771–804 Zbl 1346.20054 MR 3252963
- [9] B. H. Bowditch, [Uniform hyperbolicity of the curve graphs](#). *Pacific J. Math.* **269** (2014), no. 2, 269–280 Zbl 1330.57023 MR 3238474
- [10] J. F. Brock, R. D. Canary, and Y. N. Minsky, [The classification of Kleinian surface groups, II: The ending lamination conjecture](#). *Ann. of Math. (2)* **176** (2012), no. 1, 1–149 Zbl 1253.57009 MR 2925381
- [11] M. T. Clay, C. J. Leininger, and J. Mangahas, [The geometry of right-angled Artin subgroups of mapping class groups](#). *Groups Geom. Dyn.* **6** (2012), no. 2, 249–278 Zbl 1245.57004 MR 2914860
- [12] D. Fried, [Fibrations over \$S^1\$ with pseudo-Anosov monodromy](#). Astérisque 66–67, Société Mathématique de France, Paris, 1979 Zbl 0446.57023 MR 568308
- [13] D. Fried, [The geometry of cross sections to flows](#). *Topology* **21** (1982), no. 4, 353–371 Zbl 0594.58041 MR 670741
- [14] F. Guéritaud, [Veering triangulations and Cannon–Thurston maps](#). *J. Topol.* **9** (2016), no. 3, 957–983 Zbl 1354.57025 MR 3551845
- [15] J. Hempel, [3-manifolds as viewed from the curve complex](#). *Topology* **40** (2001), no. 3, 631–657 Zbl 0985.57014 MR 1838999
- [16] H. Masur and S. Schleimer, [The geometry of the disk complex](#). *J. Amer. Math. Soc.* **26** (2013), no. 1, 1–62 Zbl 1272.57015 MR 2983005
- [17] H. A. Masur and Y. N. Minsky, [Geometry of the complex of curves. II. Hierarchical structure](#). *Geom. Funct. Anal.* **10** (2000), no. 4, 902–974 Zbl 0972.32011 MR 1791145
- [18] C. T. McMullen, [Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations](#). *Ann. Sci. Éc. Norm. Supér (4)* **33** (2000), no. 4, 519–560 Zbl 1013.57010 MR 1832823
- [19] Y. Minsky, [The classification of Kleinian surface groups, I: Models and bounds](#). *Ann. of Math. (2)* **171** (2010), no. 1, 1–107 Zbl 1193.30063 MR 2630036
- [20] Y. N. Minsky and S. J. Taylor, [Fibered faces, veering triangulations, and the arc complex](#). *Geom. Funct. Anal.* **27** (2017), no. 6, 1450–1496 Zbl 1385.57025 MR 3737367
- [21] K. Rafi, [A characterization of short curves of a Teichmüller geodesic](#). *Geom. Topol.* **9** (2005), 179–202 Zbl 1082.30037 MR 2115672
- [22] K. Rafi, [A combinatorial model for the Teichmüller metric](#). *Geom. Funct. Anal.* **17** (2007), no. 3, 936–959 Zbl 1129.30031 MR 2346280
- [23] K. Rafi and S. Schleimer, [Covers and the curve complex](#). *Geom. Topol.* **13** (2009), no. 4, 2141–2162 Zbl 1166.57013 MR 2507116
- [24] S. Schleimer, Notes on the complex of curves. 2006, <https://homepages.warwick.ac.uk/~masgar/Maths/notes.pdf>, visited on 21 December 2023
- [25] A. Sisto and S. J. Taylor, [Largest projections for random walks and shortest curves in random mapping tori](#). *Math. Res. Lett.* **26** (2019), no. 1, 293–321 Zbl 1453.20062 MR 3963985

- [26] W. P. Thurston, A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59** (1986), no. 339, 99–130 Zbl [0585.57006](#) MR [823443](#)

Received 11 March 2022.

Yair N. Minsky

Department of Mathematics, Yale University, PO Box 208283, New Haven, CT 06520, USA;
yair.minsky@yale.edu

Samuel J. Taylor

Department of Mathematics, Temple University, 1805 North Broad Street, Philadelphia, PA 19122, USA; samuel.taylor@temple.edu