

FIBERED FACES, VEERING TRIANGULATIONS, AND THE ARC COMPLEX

YAIR N. MINSKY AND SAMUEL J. TAYLOR

ABSTRACT. We study the connections between subsurface projections in curve and arc complexes in fibered 3-manifolds and Agol’s veering triangulation. The main theme is that large-distance subsurfaces in fibers are associated to large simplicial regions in the veering triangulation, and this correspondence holds uniformly for all fibers in a given fibered face of the Thurston norm.

1. INTRODUCTION

Let M be a 3-manifold fibering over the circle with fiber S and pseudo-Anosov monodromy f . The stable/unstable laminations λ^+, λ^- of f give rise to a function on the essential subsurfaces of S ,

$$Y \mapsto d_Y(\lambda^+, \lambda^-),$$

where d_Y denotes distance in the curve and arc complex of Y between the lifts of λ^\pm to the cover of S homeomorphic to Y . This distance function plays an important role in the geometry of the mapping class group of S [MM00, BKMM12, MS13], and in the hyperbolic geometry of the manifold M [Min10, BCM12].

In this paper we study the function d_Y when M is fixed and the fibration is varied. The fibrations of a given manifold are organized by the faces of the unit ball of Thurston’s norm on $H_2(M, \partial M)$, where each *fibered face* \mathcal{F} has the property that every irreducible integral class in the open cone $\mathbb{R}_+ \mathcal{F}$ represents a fiber. There is a pseudo-Anosov flow which is transverse to every fiber represented by \mathcal{F} , and whose stable/unstable laminations $\Lambda^\pm \subset M$ intersect each fiber to give the laminations associated to its monodromy. With this we note that the projection distance $d_Y(\lambda^+, \lambda^-)$ can be defined for any subsurface Y of any fiber in \mathcal{F} . We use $d_Y(\Lambda^+, \Lambda^-)$ to denote this quantity.

Our main results give explicit connections between d_Y and the *veering triangulation* of M , introduced by Agol [Ago11] and refined by Guéritaud [Gué15], with the main feature being that when d_Y satisfies explicit lower bounds, a thickening of Y is realized as an embedded subcomplex of the veering triangulation. In this way, the “complexity” of the monodromy f is visible directly in the triangulation in a way that is independent of the choice of fiber in the face \mathcal{F} . This is in contrast with the results of [BCM12] in which the estimates relating d_Y to the hyperbolic geometry of M are heavily dependent on the genus of the fiber.

Date: November 8, 2017.

This work was partially supported by NSF grants DMS-1311844 and DMS-1400498.

The results are cleanest in the setting of a *fully-punctured* fiber, that is when the singularities of the monodromy f are assumed to be punctures of the surface S (one can obtain such examples by starting with any M and puncturing the singularities and their flow orbits). All fibers in a face \mathcal{F} are fully-punctured when any one is, and in this case we say that \mathcal{F} is a *fully-punctured face*. In this setting M is a cusped manifold and the veering triangulation τ is an ideal triangulation of M .

We obtain bounds on $d_W(\Lambda^+, \Lambda^-)$ that hold for W in any fiber of a given fibered face:

Theorem 1.1 (Bounding projections over a fibered face). *Let M be a hyperbolic 3-manifold with fully-punctured fibered face \mathcal{F} and veering triangulation τ . For any essential subsurface W of any fiber of \mathcal{F} ,*

$$\alpha \cdot (d_W(\Lambda^-, \Lambda^+) - \beta) < |\tau|,$$

where $|\tau|$ is the number of tetrahedra in τ , $\alpha = 1$ and $\beta = 10$ when W is an annulus and $\alpha = 3|\chi(W)|$ and $\beta = 8$ when W is not an annulus.

Note that this means that $d_W \leq |\tau| + 10$ for each subsurface W , no matter which fiber W lies in. Further, the complexity $|\chi(W)|$ of any subsurface W of any fiber of \mathcal{F} with $d_W(\Lambda^+, \Lambda^-) \geq 9$ is also bounded in terms of M alone.

In addition, given one fiber with a collection of subsurfaces of large d_W , we obtain control over the appearance of high-distance subsurfaces in all other fibers:

Theorem 1.2 (Subsurface dichotomy). *Let M be a hyperbolic 3-manifold with fully-punctured fibered face \mathcal{F} and suppose that S and F are each fibers in $\mathbb{R}_+\mathcal{F}$. If W is a subsurface of F , then either W is isotopic along the flow to a subsurface of S , or*

$$3|\chi(S)| \geq d_W(\Lambda^-, \Lambda^+) - \beta,$$

where $\beta = 10$ if W is an annulus and $\beta = 8$ otherwise.

One can apply this theorem with S taken to be the smallest-complexity fiber in \mathcal{F} . In this case there is some finite list of “large” subsurfaces of S , and for all other fibers and all subsurfaces W with d_W sufficiently large, W is already accounted for on this finite list.

For a sample application of [Theorem 1.2](#), let W be an essential annulus with core curve w in a fiber F of M and suppose that $d_W(\Lambda^-, \Lambda^+) \geq K$ for some $K > 10$. (We note that it is easy to construct explicit examples of M with $d_W(\Lambda^-, \Lambda^+)$ as large as one wishes by starting with a pseudo-Anosov homeomorphism of F with large twisting about the curve w .) If w is trivial in $H_1(M)$, then [Theorem 1.2](#) (or more precisely [Corollary 6.7](#)) implies that w is actually isotopic to a simple closed curve in *every* fiber in the open cone $\mathbb{R}_+\mathcal{F}$ containing F . When w is nontrivial in $H_1(M)$ it determines a codimension-1 hyperplane P_w in $H^1(M) = H_2(M, \partial M)$ consisting of cohomology classes which vanish on w . For each fiber S of $\mathbb{R}_+\mathcal{F}$ either S is contained in P_w , in which case w is isotopic to a simple closed curve in S as before, or S lies outside of P_w and $|\chi(S)| \geq \frac{K-10}{3}$. We remark that the second alternative is non-vacuous so long as $H^1(M)$ has rank at least 2.

The general (non-fully-punctured) setting is also approachable with our techniques, but a number of complications arise and the connection to the veering triangulation of the fully-punctured manifold is much less explicit. An extension of the results in this paper to the general setting will be the subject of a subsequent paper.

Pockets in the veering triangulation. When Y is a subsurface of a fiber X in \mathcal{F} and $d_Y(\Lambda^+, \Lambda^-) > 1$, we show ([Theorem 5.3](#)) that Y is realized simplicially in the veering triangulation lifted to the cover $X \times \mathbb{R}$. If $d_Y(\Lambda^+, \Lambda^-)$ is even larger then this realization can be thickened to a “pocket”, which is a simplicial region bounded by two isotopic copies of Y . With sufficiently large assumptions this pocket can be made to embed in M as well, and this is our main tool for connecting arc complexes to the veering triangulation and establishing [Theorems 1.1](#) and [1.2](#):

Theorem 1.3. *Suppose Y is a subsurface of a fiber X with $d_Y(\lambda^-, \lambda^+) > \beta$, where $\beta = 8$ if Y is nonannular and $\beta = 10$ if Y is an annulus. Then there is an embedded simplicial pocket V in M isotopic to a thickening of Y , and with $d_Y(V^-, V^+) \geq d_Y(\lambda^-, \lambda^+) - \beta$.*

In this statement, V^+ and V^- refer to the triangulations of the top and bottom surfaces of the pockets, regarded as simplices in the curve and arc complex $\mathcal{A}(Y)$. Also, $d_Y(V^-, V^+)$ denotes the smallest d_Y -distance between an arc of V^- and an arc of V^+ .

The veering triangulation in fact recovers a number of aspects of the geometry of curve and arc complexes in a fairly concrete way. As an illustration we prove

Theorem 1.4. *In the fully punctured setting, the arcs of the veering triangulation form a geodesically connected subset $\mathcal{A}(\tau)$ of the curve and arc graph, in the sense that any two points in $\mathcal{A}(\tau)$ are connected by a geodesic that lies in $\mathcal{A}(\tau)$.*

Hierarchies of pockets. One is naturally led to generalize [Theorem 1.3](#) from a result embedding one pocket at a time to a description of all pockets at once. Indeed [Proposition 6.5](#) tells us that whenever subsurfaces Y and Z of X have large enough projection distances and are not nested, they have associated pockets V_Y and V_Z which are disjoint in $X \times \mathbb{R}$. These facts, taken together with [Theorem 1.4](#), strongly suggest that the veering triangulation τ encodes the hierarchy of curve complex geodesics between λ^\pm as introduced by Masur-Minsky in [[MM00](#)]. We expect that, using a version of [Theorem 1.4](#) that applies to subsurfaces and adapting the notion of “tight geodesic” from [[MM00](#)], one can carry out a hierarchy-like construction within the veering triangulation and recover much of the structure found in [[MM00](#)], with more concrete control, at least in the fully-punctured setting. We plan to explore this approach in future work.

Related and motivating work. The theme of using fibered 3-manifolds to study infinite families of monodromy maps is deeply explored in McMullen [[McM00](#)] and Farb-Leininger-Margalit [[FLM11](#)], where the focus is on Teichmüller translation distance.

Distance inequalities analogous to [Theorem 1.2](#), in the setting of Heegaard splittings rather than surface bundles, appear in Hartshorn [[Har02](#)], and then more fully in Scharlemann-Tomova [[ST06](#)]. Bachman-Schleimer [[BS05](#)] use Heegaard surfaces to give bounds on the curve-complex translation distance of the monodromy of a fibering. All of these bounds apply to entire surfaces and not to subsurface projections. In Johnson-Minsky-Moriah [[JMM10](#)], subsurface projections are considered in the setting of Heegaard splittings. A basic difficulty in these papers which we do not encounter here is the compressibility of the Heegaard surfaces, which makes it tricky to control essential intersections. On the other hand, unlike the surfaces and handlebodies that are used to obtain control in the Heegaard setting, the foliations we consider here are infinite objects, and the connection between them and finite arc systems in the surface is a priori dependent on the fiber complexity. The veering triangulation provides a finite object that captures this connection in a more uniform way.

The totally-geodesic statement of [Theorem 1.4](#) should be compared to [Theorem 1.2](#) of Tang-Webb [[TW15](#)], in which Teichmüller disks give rise to quasi-convex sets in curve complexes. While the results of Tang-Webb are more general, they are coarse, and it is interesting that in our setting a tighter statement holds. Finally, we note that work by several authors has focused on geometric aspects of the veering triangulation, including [[HRST11](#), [FG13](#), [HIS16](#)].

Summary of the paper. In [Section 2](#) we set some notation and give Guéritaud’s construction of the veering triangulation. We also recall basic facts about curve and arc complexes, subsurface projections and Thurston’s norm on homology. We spend some time in this section describing the flat geometry of a punctured surface with an integrable holomorphic quadratic differential, and in particular giving an explicit description of the circle at infinity of its universal cover ([Proposition 2.2](#)). While this is a fairly familiar picture, some delicate issues arise because of the incompleteness of the metric at the punctures.

In [Section 3](#) we study *sections* of the veering triangulations, which are simplicial surfaces isotopic to X in the cover $X \times \mathbb{R}$, and transverse to the suspension flow of the monodromy. These can be thought of as triangulations of the surface X using only edges coming from the veering triangulation. We prove [Lemma 3.2](#) which says that a partial triangulation of X using only edges from τ can always be extended to a full section, and [Proposition 3.3](#) which says that any two extensions of a partial triangulation are connected by a sequence of “tetrahedron moves”. This is what allows us to define and study the “pockets” that arise between any two sections.

In [Section 4](#) we define a simple but useful construction, rectangle and triangle hulls, which map saddle connections in our surface to unions of edges of the veering triangulation. An immediate consequence of the properties of these hulls is a proof of [Theorem 1.4](#).

In [Section 5](#) we apply the flat geometry developed in [Section 2](#) to control the convex hulls of subsurfaces of the fiber, and then use [Section 4](#) to construct what we call τ -hulls, which are representatives of the homotopy class of a subsurface that are simplicial with respect to the veering triangulations. [Theorem 5.3](#) states that quite

mild assumptions on $d_Y(\lambda^+, \lambda^-)$ imply that the τ -hull of Y has embedded interior. The idea here is that any pinching point of the τ -hull is crossed by leaves of λ^+ and λ^- that intersect each other very little. The main results of both [Section 4](#) and [Section 5](#) apply in a general setting and do not require that the surface X be fully-punctured.

In [Section 6](#) we put these ideas together to prove our main theorems for fibered manifolds with a fully-punctured fibered face. In [Proposition 6.2](#) we describe the maximal pocket associated to a subsurface Y with $d_Y(\Lambda^+, \Lambda^-)$ sufficiently large (greater than 2, for nonannular Y). We then introduce the notion of an *isolated pocket*, which is a subpocket of the maximal pocket that has good embedding properties in the manifold M . The existence and embedding properties of these pockets are established in [Lemma 6.4](#) and [Proposition 6.5](#), which together allow us to prove [Theorem 1.3](#).

From here, a simple counting argument gives [Theorem 1.1](#): the size of the embedded isolated pockets is bounded from below in terms of $d_Y(\Lambda^+, \Lambda^-)$ and $\chi(Y)$, and from above by the total number of veering tetrahedra.

To obtain [Theorem 1.2](#), we use the pocket embedding results to show that, if Y is a subsurface of one fiber F and Y essentially intersects another fiber S , then S must cross every level surface of the isolated pocket of Y , and hence the complexity of S gives an upper bound for $d_Y(\Lambda^+, \Lambda^-)$. To complete the proof we need to show that, if Y does not essentially cross S , it must be isotopic to an embedded (and not merely immersed) subsurface of S . This is handled by [Lemma 6.6](#), which may be of independent interest. It gives a uniform upper bound for $d_Y(\Lambda^+, \Lambda^-)$ when Y corresponds to a finitely generated subgroup of $\pi_1(S)$, unless Y covers an embedded subsurface.

Acknowledgments. The authors are grateful to Ian Agol and François Guéritaud for explaining their work to us. We also thank Tarik Aougab, Jeff Brock, and Dave Futer for helpful conversations and William Worden for pointing out some typos in an earlier draft. Finally, we thank the referee for a thorough reading of our paper and comments which improved its readability.

2. BACKGROUND

The following notation will hold throughout the paper. Let \bar{X} be a closed Riemann surface with an integrable meromorphic quadratic differential q . We remind the reader that q may have poles of order 1. We denote the vertical and horizontal foliations of q by λ^+ and λ^- respectively. Let \mathcal{P} be a finite subset of \bar{X} that includes the poles of q if any, and let $X = \bar{X} \setminus \mathcal{P}$. Let $\text{sing}(q)$ denote the union of \mathcal{P} with the set of zeros of q . We require further that q has no horizontal or vertical saddle connections, that is no leaves of λ^\pm that connect two points of $\text{sing}(q)$. This situation holds in particular if λ^\pm are the stable/unstable foliations of a pseudo-Anosov map $f : X \rightarrow X$, which will often be the case for us. If $\mathcal{P} = \text{sing}(q)$ (i.e. \mathcal{P} contains all zeros of q) we say X is *fully-punctured*.

Let \hat{X} denote the metric completion of the universal cover \tilde{X} of X , and note that there is an infinite branched covering $\hat{X} \rightarrow \tilde{X}$, infinitely branched over the points of \mathcal{P} . The preimage $\hat{\mathcal{P}}$ of \mathcal{P} is the set of completion points. The space \hat{X} is a complete CAT(0) space with the metric induced by q .

2.1. Veering triangulations. In this section let $\mathcal{P} = \text{sing}(q)$. The veering triangulation, originally defined by Agol in [Ago11] in the case where q corresponds to a pseudo-Anosov $f : X \rightarrow X$, is an ideal layered triangulation of $X \times \mathbb{R}$ which projects to a triangulation of the mapping torus M of f . The definition we give here is due to Guéritaud [Gué15]. (Agol’s “veering” property itself will not actually play a role in this paper, so we will not give its definition).

A *singularity-free rectangle* in \hat{X} is an embedded rectangle whose edges consist of leaf segments of the lifts of λ^\pm and whose interior contains no singularities of \hat{X} . If R is a *maximal* singularity-free rectangle in \hat{X} then it must contain a singularity on each edge. Note that there cannot be more than one singularity on an edge since λ^\pm have no saddle connections. We associate to R an ideal tetrahedron whose vertices are $\partial R \cap \hat{\mathcal{P}}$, as in Figure 1. This tetrahedron comes equipped with a “flattening” map into \hat{X} as pictured.

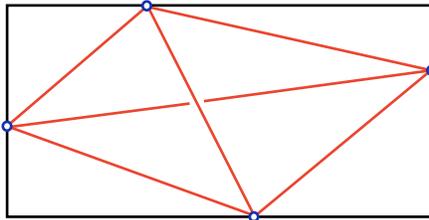


FIGURE 1. A maximal singularity-free rectangle R defines a tetrahedron equipped with a map into R .

The tetrahedron comes with a natural orientation, inherited from the orientation of \hat{X} using the convention that the edge connecting the horizontal boundaries of the rectangle lies *above* the edge connecting the vertical boundaries. This orientation is indicated in Figure 1.

The union of all these ideal tetrahedra, with faces identified whenever they map to the same triangle in \hat{X} , is Guéritaud’s construction of the veering triangulation of $\tilde{X} \times \mathbb{R}$.

Theorem 2.1. [Gué15] *Suppose that X is fully-punctured. The complex of tetrahedra associated to maximal rectangles of q is an ideal triangulation $\tilde{\tau}$ of $\tilde{X} \times \mathbb{R}$, and the maps of tetrahedra to their defining rectangles piece together to a fibration $\pi : \tilde{X} \times \mathbb{R} \rightarrow \tilde{X}$. The action of $\pi_1(X)$ on $(\tilde{X}, \tilde{\tau})$ lifts simplicially to $\tilde{\tau}$, and equivariantly with respect to π . The quotient is a triangulation of $X \times \mathbb{R}$.*

If q corresponds to a pseudo-Anosov $f : X \rightarrow X$ then the action of f on (X, q) lifts simplicially and π -equivariantly to $\Phi : X \times \mathbb{R} \rightarrow X \times \mathbb{R}$. The quotient is a

triangulation τ of the mapping torus M . The fibers of π descend to flow lines for the suspension flow of f .

We will frequently abuse notation and use τ to refer to the triangulation both in M and in its covers.

We note that a saddle connection σ of q is an edge of τ if and only if σ spans a singularity-free rectangle in X . See [Figure 2](#).

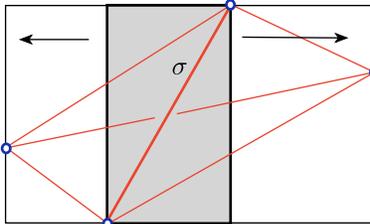


FIGURE 2. The singularity-free rectangle spanned by σ can be extended horizontally (or vertically) to a maximal one.

If e and f are two crossing τ -edges spanning rectangles R_e and R_f , note that R_e crosses R_f from top to bottom, or from left to right – any other configuration would contradict the singularity-free property of the rectangles ([Figure 3](#)). If $\text{slope}(e)$ denotes the absolute value of the slope of e with respect to q , we can see that R_e crosses R_f from top to bottom if and only if e crosses f and $\text{slope}(e) > \text{slope}(f)$. We say that e is *more vertical* than f and also write $e > f$. We will see that $e > f$ corresponds to e lying higher than f in the upward flow direction.

Indeed we can see already that the relation $>$ is transitive, since if $e > f$ and $f > g$ then the rectangle of g is forced to intersect the rectangle of e .

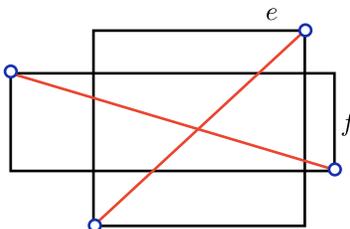


FIGURE 3. The rectangle of e crosses f from top to bottom and we write $e > f$.

We conclude with a brief description of the local structure of τ around an edge e : The rectangle spanned by e can be extended horizontally to define a tetrahedron lying below e in the flow direction ([Figure 2](#)), and vertically to define a tetrahedron lying above e in the flow direction. Call these Q_- and Q_+ as in [Figure 4](#). Between these, on each side of e , is a sequence of tetrahedra Q_1, \dots, Q_m ($m \geq 1$) so that two successive tetrahedra in the sequence $Q_-, Q_1, \dots, Q_m, Q_+$ share a triangular face

adjacent to e . We find this sequence by starting with one of the two top faces of Q_- , extending its spanning rectangle vertically until it hits a singularity, and calling Q_1 the tetrahedron whose projection is inscribed in the new rectangle. If the new singularity belongs to Q_+ we are done ($m = 1$), otherwise we repeat from the top face of Q_1 containing e to find Q_2 , and continue in this manner. Figure 4 illustrates this structure on one side of an edge e . Repeating on the other side, note that the link of the edge e is a circle, as expected.

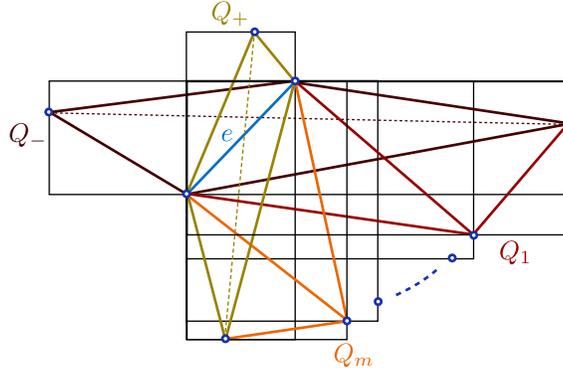


FIGURE 4. The tetrahedra adjacent to an edge e on one side form a sequence “swinging” around e

2.2. Arc and curve complexes. The arc and curve complex $\mathcal{A}(Y)$ for a compact surface Y is usually defined as follows: its vertices are essential homotopy classes of embedded circles and properly embedded arcs $([0, 1], \{0, 1\}) \rightarrow (Y, \partial Y)$, where “essential” means not homotopic to a point or into the boundary [MM00]. We must be clear about the meaning of homotopy classes here, for the case of arcs: If Y is not an annulus, homotopies of arcs are assumed to be homotopies of maps of pairs. When Y is an annulus the homotopies are also required to fix the endpoints. Simplices of $\mathcal{A}(Y)$, in all cases, correspond to tuples of vertices which can be simultaneously realized by maps that are disjoint on their interiors. We endow $\mathcal{A}(Y)$ with the simplicial distance on its 1-skeleton.

It will be useful, in the non-annular case, to observe that the following definition is equivalent: Instead of maps of closed intervals consider proper embeddings $\mathbb{R} \rightarrow \text{int}(Y)$ into the interior of Y , with equivalence arising from proper homotopy. This definition is independent of the compactification of $\text{int}(Y)$. The natural isomorphism between these two versions of $\mathcal{A}(Y)$ is induced by a straightening construction in a collar neighborhood of the boundary.

If $Y \subset S$ is an essential subsurface (meaning the inclusion of Y is π_1 -injective and is not homotopic to a point or to an end of S), we have subsurface projections $\pi_Y(\lambda)$ which are defined for simplices $\lambda \subset \mathcal{A}(S)$ that intersect Y essentially. Namely, after lifting λ to the cover S_Y associated to $\pi_1(Y)$ (i.e. the cover to which Y lifts homeomorphically and for which $S_Y \cong \text{int}(Y)$), we obtain a collection of properly

embedded disjoint essential arcs and curves, which determine a simplex of $\mathcal{A}(Y)$. We let $\pi_Y(\lambda)$ be the union of these vertices [MM00]. We make a similar definition for a lamination λ that intersects Y essentially, except that we include not just the leaves of λ but all leaves that one can add in the complement of λ which accumulate on λ . This is natural when we realize λ as a measured *foliation* (as we do in most of the paper), and need to include *generalized leaves*, which are leaves that are allowed to pass through singularities. Note that the diameter of $\pi_Y(\lambda)$ in $\mathcal{A}(Y)$ is at most 2.

Note that 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 $\lambda \cap Y$.

These definitions naturally extend to immersed surfaces arising from covers of S . Let Γ be a finitely generated subgroup of $\pi_1(S)$. Then the corresponding cover $S_\Gamma \rightarrow S$ has a compact core W – a compact subsurface $W \subset S_\Gamma$ such that $S_\Gamma \setminus W$ is a collection of boundary parallel annuli. For curves or laminations λ^\pm of S , we have lifts $\tilde{\lambda}^\pm$ to S_Γ and define $d_W(\lambda^-, \lambda^+) = d_{S_\Gamma}(\tilde{\lambda}^-, \tilde{\lambda}^+)$.

Throughout this paper, when λ, λ' are two laminations or arc/curve systems, we denote by $d_Y(\lambda, \lambda')$ the *minimal* distance between their images in $\mathcal{A}(Y)$, that is

$$d_Y(\lambda, \lambda') = \min\{d_Y(l, l') : l \in \pi_Y(\lambda), l' \in \pi_Y(\lambda')\}.$$

To denote the *maximal* distance between λ and λ' in $\mathcal{A}(Y)$ we write

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

2.3. Flat geometry. In this section we return to the singular Euclidean geometry of (X, q) and describe a circle at infinity for the flat metric induced by q on the universal cover \tilde{X} . We identify \tilde{X} with \mathbb{H}^2 after fixing a reference hyperbolic metric on X . Because of incompleteness of the flat metric at the punctures \mathcal{P} , the connection between the circle we will describe and the usual circle at infinity for \mathbb{H}^2 requires a bit of care. A related discussion appears in Guéritaud [Gué15], although he deals explicitly only with the fully-punctured case. With this picture of the circle at infinity we will be able to describe π_Y in terms of q -geodesic representatives, and to describe a q -convex hull for essential subsurfaces of X . In this section we do not assume that X is fully-punctured.

The completion points $\hat{\mathcal{P}}$ in \hat{X} correspond to parabolic fixed points for $\pi_1(X)$ in $\partial\mathbb{H}^2$, and we abuse notation slightly by identifying $\hat{\mathcal{P}}$ with this subset of $\partial\mathbb{H}^2$.

A *complete q -geodesic ray* is either a geodesic ray $r : [0, \infty) \rightarrow \hat{X}$ of infinite length, or a finite-length geodesic segment that terminates in $\hat{\mathcal{P}}$. A complete q -geodesic line is a geodesic which is split by any point into two complete q -geodesic rays. Our goal in this section is to describe a circle at infinity that corresponds to endpoints of these rays.

Proposition 2.2. *There is a compactification $\beta(\tilde{X})$ of \tilde{X} on which $\pi_1(X)$ acts by homeomorphisms, with the following properties:*

- (1) *There is a $\pi_1(X)$ -equivariant homeomorphism $\beta(\tilde{X}) \rightarrow \overline{\mathbb{H}^2}$, extending the identification of \tilde{X} with \mathbb{H}^2 and taking $\hat{\mathcal{P}}$ to the corresponding parabolic fixed points in $\partial\mathbb{H}^2$.*
- (2) *If l is a complete q -geodesic line in \hat{X} then its image in $\overline{\mathbb{H}^2}$ is an embedded arc with endpoints on $\partial\mathbb{H}^2$ and interior points in $\mathbb{H}^2 \cup \hat{\mathcal{P}}$. Conversely, every pair of distinct points x, y in $\partial\beta(\tilde{X}) = \beta(\tilde{X}) \setminus \tilde{X}$ are the endpoints of a complete q -geodesic line. The termination point in $\partial\mathbb{H}^2$ of a complete q -geodesic ray is in $\hat{\mathcal{P}}$ if and only if it has finite length.*
- (3) *The q -geodesic line connecting distinct $x, y \in \partial\beta(\tilde{X})$ is either unique, or there is a family of parallel geodesics making up an infinite Euclidean strip.*

One of the tricky points of this picture is that q -geodesic rays and lines may meet points of the boundary $\partial\beta(\tilde{X})$ not just at their endpoints.

Proof. When $\mathcal{P} = \emptyset$ and X is a closed surface, \tilde{X} is quasi-isometric to \mathbb{H}^2 and the proposition holds for the standard Gromov compactification. We assume from now on that $\mathcal{P} \neq \emptyset$.

We begin by setting $\hat{\mathbb{H}}^2 = \mathbb{H}^2 \cup \hat{\mathcal{P}}$ and endowing it with the topology obtained by taking, for each $p \in \hat{\mathcal{P}}$, horoballs based at p as a neighborhood basis for p .

Lemma 2.3. *The natural identification of \tilde{X} with \mathbb{H}^2 extends to a homeomorphism from \hat{X} to $\hat{\mathbb{H}}^2$.*

Proof. First note that $\hat{\mathcal{P}}$ is discrete as both a subspace of \hat{X} and of $\hat{\mathbb{H}}^2$. Hence, it suffices to show that a sequence of points x_i in $\tilde{X} = \mathbb{H}^2$ converges to a point $p \in \hat{\mathcal{P}}$ in \hat{X} if and only if it converges to p in $\hat{\mathbb{H}}^2$. This follows from the fact that the horoball neighborhoods of p descend to cusp neighborhoods in X which form a neighborhood basis for the puncture that is equivalent to the neighborhood basis of q -metric balls. \square

Our strategy now is to form the *Freudenthal space* of \hat{X} and equivalently $\hat{\mathbb{H}}^2$, which appends a space of *ends*. This space will be compact but not Hausdorff, and after a mild quotient we will obtain the desired compactification which can be identified with $\overline{\mathbb{H}^2}$. Simple properties of this construction will then allow us to obtain the geometric conclusions in part (2) of the proposition.

Let $\epsilon(\hat{X})$ be the space of ends of \hat{X} , that is the inverse limit of the system of path components of complements of compact sets in \hat{X} . The Freudenthal space $\text{Fr}(\hat{X})$ is the union $\hat{X} \cup \epsilon(\hat{X})$ endowed with the topology generated by using path components of complements of compacta to describe neighborhood bases for the ends. Because \hat{X} is not locally compact, $\text{Fr}(\hat{X})$ is not guaranteed to be compact, and we have to take a bit of care to describe it.

The construction can of course be repeated for $\hat{\mathbb{H}}^2$, and the homeomorphism of **Lemma 2.3** gives rise to a homeomorphism $\text{Fr}(\hat{X}) \rightarrow \text{Fr}(\hat{\mathbb{H}}^2)$. Let us work in $\hat{\mathbb{H}}^2$ now, where we can describe the ends concretely using the following observations:

Every compact set $K \subset \hat{\mathbb{H}}^2$ meets $\hat{\mathcal{P}}$ in a finite set A (since $\hat{\mathcal{P}}$ is discrete in $\hat{\mathbb{H}}^2$), and such a K is contained in an embedded closed disk D which also meets $\hat{\mathcal{P}}$ at A .

(This is not hard to see but does require attention to deal correctly with the horoball neighborhood bases). The components of $\hat{\mathbb{H}}^2 \setminus D$ determine a partition of $\epsilon(\hat{\mathbb{H}}^2)$, which in fact depends only on the set A and not on D (if D' is another disk meeting $\hat{\mathcal{P}}$ at A , then $D \cup D'$ is contained in a third disk D'' , and this common refinement of the neighborhoods gives the same partition). Thus we have a more manageable (countable) inverse system of neighborhoods in $\epsilon(\hat{\mathbb{H}}^2)$, and with this description it is not hard to see that $\epsilon(\hat{\mathbb{H}}^2)$ is a Cantor set.

For each $p \in \hat{\mathcal{P}}$ there are two distinguished ends $p^+, p^- \in \epsilon(\hat{\mathbb{H}}^2)$ defined as follows: For each finite subset $A \subset \hat{\mathcal{P}}$ with at least two points one of which is p , the two partition terms adjacent to p in the circle (or equivalently, in the boundary of any $D \subset \hat{\mathbb{H}}^2$ meeting $\hat{\mathcal{P}}$ in A) define neighborhoods in $\epsilon(\hat{\mathbb{H}}^2)$, and this pair of neighborhood systems determines p^+ and p^- respectively.

One can also see that p^+ (and p^-) and p do not admit disjoint neighborhoods, and this is why $\text{Fr}(\hat{\mathbb{H}}^2)$ is not Hausdorff. We are therefore led to define the quotient space

$$\beta(\hat{\mathbb{H}}^2) = \text{Fr}(\hat{\mathbb{H}}^2) / \sim,$$

where we make the identifications $p^- \sim p \sim p^+$, for each $p \in \hat{\mathcal{P}}$.

We can make the same definitions in \hat{X} , obtaining

$$\beta(\hat{X}) = \text{Fr}(\hat{X}) / \sim,$$

which we rename $\beta(\tilde{X})$. Since the definitions are purely in terms of the topology of the spaces $\hat{\mathbb{H}}^2$ and \hat{X} , the homeomorphism of [Lemma 2.3](#) extends to a homeomorphism $\beta(\tilde{X}) \rightarrow \beta(\hat{\mathbb{H}}^2)$.

Part (1) of [Proposition 2.2](#) follows once we establish that the identity map of \mathbb{H}^2 extends to a homeomorphism

$$\beta(\hat{\mathbb{H}}^2) \cong \overline{\mathbb{H}^2}.$$

This is not hard to see once we observe that the disks used above to define neighborhood systems can be chosen to be ideal hyperbolic polygons. Their halfspace complements serve as neighborhood systems for points of $\partial\mathbb{H}^2 \setminus \hat{\mathcal{P}}$. A sequence converges in $\overline{\mathbb{H}^2}$ to a point $p \in \hat{\mathcal{P}}$ if it is eventually contained in any union of a horoball centered at p and two half-planes adjacent to p on opposite sides. This is modeled exactly by the equivalence relation \sim .

For part (2), let D_0 be a fundamental domain for $\pi_1(X)$ in \hat{X} , which may be chosen to be a disk with vertices at points of $\hat{\mathcal{P}}$, and of finite q -diameter. Translates of D_0 can be glued to build a sequence of nested disks D_n exhausting \hat{X} , each of which meets $\hat{\mathcal{P}}$ in a finite set of vertices, and whose boundary is composed of arcs of bounded diameter between successive vertices.

A complete q -geodesic ray r either has finite length and terminates in a point of $\hat{\mathcal{P}}$, or has infinite length in which case it leaves every compact set of \hat{X} , and visits each point of $\hat{\mathcal{P}}$ at most once. Thus it must terminate in a point of $\epsilon(\hat{X})$ in the Freudenthal space. We claim that this point cannot be p^+ or p^- for $p \in \hat{\mathcal{P}}$. If r terminates in p^+ , then for each disk D_n (n large) it must pass through the edge of

∂D_n adjacent to p on the side associated to p^+ . Any two such consecutive edges meet in p at one of finitely many angles (images of corners of D_0), and hence the accumulated angle between edges goes to ∞ with n . If we replace these edges by their q -geodesic representatives, the angles still go to ∞ . This means that r contains infinitely many disjoint subsegments whose endpoints are a bounded distance from p , but this contradicts the assumption that r is a geodesic ray.

The image of r in the quotient $\beta(\tilde{X})$ therefore terminates in a point of $\hat{\mathcal{P}}$ when it has finite length, and a point in $\partial\beta(\tilde{X}) \setminus \hat{\mathcal{P}}$ otherwise. The same is true for both ends of a complete q -geodesic line l , and we note that both ends of l cannot land on the same point because then we would have a sequence of segments $l_n \subset l$ of length going to ∞ with both endpoints of l_n on the same edge or on two consecutive edges of ∂D_n , a contradiction to the fact that l_n is a geodesic and the arcs in ∂D_n have bounded q -length.

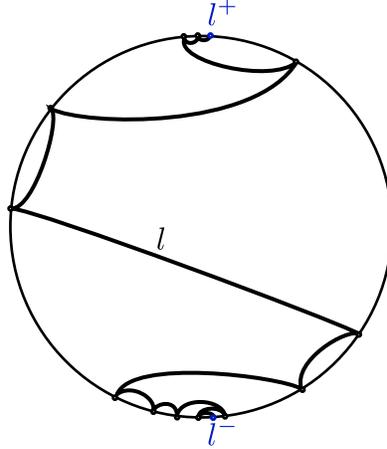
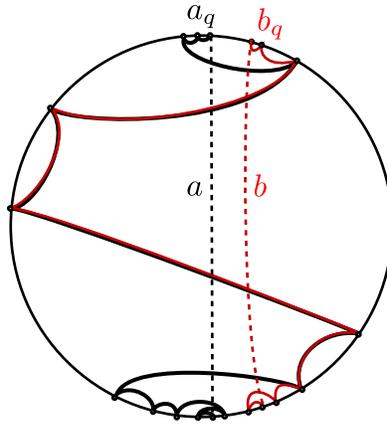
Now let x, y be two distinct points in $\partial\beta(\tilde{X})$. Assume first that both are not in $\hat{\mathcal{P}}$. Then for large enough n , they are in separate components of the complement of D_n . If we let $x_i \rightarrow x$ and $y_i \rightarrow y$ be sequences in $\beta(\tilde{X})$, then eventually x_i and y_i are in the same components of the complement of D_n as x and y , respectively. The geodesic from x_i to y_i must therefore pass through the corresponding boundary segments of D_n and in particular through D_n , so we can extract a convergent subsequence as $i \rightarrow \infty$. Letting $n \rightarrow \infty$ and diagonalizing we obtain a limiting geodesic which terminates in x, y as desired. If $x \in \hat{\mathcal{P}}$ or $y \in \hat{\mathcal{P}}$ the same argument works except that we can take $x_i \equiv x$ or $y_i \equiv y$. This establishes part (2).

Now let l and l' be two q -geodesics terminating in x and y . If x and y are in $\hat{\mathcal{P}}$ then $l = l'$ since the metric is CAT(0). If $x \notin \hat{\mathcal{P}}$ then both l and l' pass through infinitely many segments of ∂D_n on their way to x . Since these segments have uniformly bounded lengths, l and l' remain a bounded distance apart. If $y \in \hat{\mathcal{P}}$ then again CAT(0) implies that $l = l'$, and if $y \notin \hat{\mathcal{P}}$ then l and l' must cobound an infinite flat strip. This establishes part (3). \square

With [Proposition 2.2](#) in hand we can consider each complete q -geodesic line l in $\hat{X} = \hat{\mathbb{H}}^2$ as an arc in the closed disk $\overline{\mathbb{H}^2}$, which by the Jordan curve theorem separates the disk \mathbb{H}^2 into at least 2 components. Each component is an open disk whose closure meets $\partial\mathbb{H}^2$ in a subarc of one of the complementary arcs of the endpoints of l . We call the union of disks whose closures meet one of these complementary arcs of the endpoints of l an *open side* \mathcal{D}_l° of l . The closure of each open side in $\overline{\mathbb{H}^2}$ is then a connected union of closed disks, attached to each other along the points of $\hat{\mathcal{P}}$ that l meets on the circle. We call the closure of the open side \mathcal{D}_l° of l in $\overline{\mathbb{H}^2}$ the *side* \mathcal{D}_l . Note that $\mathcal{D}_l^\circ = \text{int}(\mathcal{D}_l \cap \mathbb{H}^2) = \mathcal{D}_l \setminus (\partial\mathbb{H}^2 \cup l)$, and if \mathcal{D}_l and \mathcal{D}_l' are the two sides of l , then $\mathcal{D}_l \cap \mathcal{D}_l' = l$. See [Figure 5](#).

With this picture we can state the following:

Corollary 2.4. *Let a, b be disjoint arcs in \mathbb{H}^2 with well-defined, distinct endpoints on $\partial\mathbb{H}^2$ and let a_q, b_q be q -geodesic lines with the same endpoints as a and b , respectively. Then b_q is contained in a single side of a_q .*


 FIGURE 5. A complete q -geodesic line l and its endpoints on $\partial\mathbb{H}^2$.

 FIGURE 6. Disjoint arcs with their q -geodesic representatives.

Proof. Letting L and R be the arcs of $\partial\mathbb{H}^2$ minus the endpoints of a , the endpoints of b must lie in one of them, say L , since a and b are disjoint.

Since a_q and b_q are geodesics in the CAT(0) space \tilde{X} , their intersection is connected. If their intersection is empty, then the corollary is clear. Otherwise, $b_q \setminus a_q$ is one or two arcs, each with one endpoint on a_q and the other on L . It follows that $b_q \setminus a_q$ is on one open side of a_q , and the corollary follows. \square

Subsurfaces and projections in the flat metric. Let $Y \subset X$ be an essential compact subsurface, and let $X_Y = \tilde{X}/\pi_1(Y)$ be the associated cover of X . (Here we have identified $\pi_1(X)$ with the deck transformations of $\tilde{X} \rightarrow X$ and fixed $\pi_1(Y)$ within its conjugacy class.) For any lamination λ in X , we want to show that the

projection $\pi_Y(\lambda)$ can be realized by subsegments of the q -geodesic representative of λ . Recall that X is not necessarily fully-punctured.

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 them 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 \tilde{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 $\Lambda \subset \partial\mathbb{H}^2$ be the limit set of $\pi_1(Y)$, $\Omega = \partial\mathbb{H}^2 \setminus \Lambda$, and $\hat{\mathcal{P}}_Y \subset \Lambda$ the set of parabolic fixed points of $\pi_1(Y)$. Let $C(X_Y)$ denote the compactification of X_Y given by $(\mathbb{H}^2 \cup \Omega \cup \hat{\mathcal{P}}_Y)/\pi_1(Y)$, adding a point for each puncture-parallel end of X_Y , and a circle for each of the other ends. Now given a lamination (or foliation) λ , realized geodesically in the hyperbolic metric on X , its lift to X_Y extends to properly embedded arcs in $C(X_Y)$, of which the ones that are essential give $\pi_Y(\lambda)$.

Proposition 2.2 allows us to perform the same construction with the q -geodesic representative of λ . Note that the leaves we obtain may meet points of $\tilde{\mathcal{P}}_Y$ in their interior, but a slight perturbation produces properly embedded lines in X_Y which are properly isotopic to the leaves coming from λ .

If Y is an annulus the same construction works, with the observation that the ends of Y cannot be puncture-parallel and hence $C(Y)$ is a closed annulus and the leaves have well-defined endpoints in its boundary. We have proved:

Lemma 2.5. *Let $Y \subset X$ be an essential subsurface. If λ is a proper arc or lamination in X then the lifts of its q -geodesic representatives to X_Y , after discarding inessential components, give representatives of $\pi_Y(\lambda)$.*

q -convex hulls. We will need a flat-geometry analogue of the hyperbolic convex hull. The main idea is simple – pull the boundary of the regular convex hull tight using q -geodesics. The only difficulty comes from the fact that these geodesics can pass through parabolic fixed points, and fail to be disjoint from each other, so the resulting object may fail to be an embedded surface. Our discussion is similar to Section 3 of Rafi [Raf05], but the discussion there requires adjustments to handle correctly the incompleteness at punctures.

As above, identify \tilde{X} with \mathbb{H}^2 . 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 . We define $\text{CH}_q(\Lambda)$ 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 . By part (3) of **Proposition 2.2**, l_q is unique unless it is part of a parallel family of geodesics, making a Euclidean strip.

The plane is divided by l_q into two sides as in the discussion before **Corollary 2.4**, and one of the sides, which we call \mathcal{D}_l , meets $\partial\mathbb{H}^2$ in a subset of the complement of Λ . Recall that \mathcal{D}_l is either a disk or a string of disks attached along puncture 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 $\partial\text{CH}(\Lambda)$ (which are disjoint by [Corollary 2.4](#)), we obtain $\text{CH}_q(\Lambda)$, the q -convex hull.

If Λ has 2 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 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 . Because $\text{CH}_q(X_Y)$ may not be homeomorphic to Y , we pay explicit attention to a marking map between Y and its hull.

Let $\hat{i} : Y \rightarrow X_Y$ be the lift of the inclusion map to the cover.

Lemma 2.6. *The lift $\hat{i} : Y \rightarrow X_Y$ is homotopic to a map $\hat{i}_q : Y \rightarrow \bar{X}_Y$ whose image is the q -hull $\text{CH}_q(X_Y)$ such that*

- (1) *The homotopy $(h_t)_{t \in [0,1]}$ from \hat{i} to \hat{i}_q has the property that $h_t(Y) \subset X_Y$ for all $t \in [0, 1]$.*
- (2) *Each component of $\partial_0 Y$ is taken by \hat{i}_q to the corresponding completion point of $\tilde{\mathcal{P}}_Y$.*
- (3) *If Y is an annulus then the image of \hat{i}_q is either a maximal flat cylinder in \bar{X}_Y or the unique geodesic representative of the core of Y in \bar{X}_Y .*
- (4) *If Y is not an annulus then each component γ of $\partial' Y$ is taken by \hat{i}_q to a q -geodesic representative in \bar{X}_Y . If there is a flat cylinder in the homotopy class of γ then the interior of the cylinder is disjoint from $\hat{i}_q(Y)$.*
- (5) *There is a deformation retraction $r : \bar{X}_Y \rightarrow \hat{i}_q(Y)$. For each component γ of $\partial' Y$, the preimage $r^{-1}(\hat{i}_q(\gamma))$ intersects X_Y in either an open annulus or a union of open disks joined in a cycle along points in their closures.*
- (6) *If the interior $\text{int}(\text{CH}_q(\Lambda_Y))$ is a disk then \hat{i}_q is a homeomorphism from $Y' = Y \setminus \partial_0 Y$ to its image.*

Proof. Let $\Gamma = \pi_1 Y$ and let $\Lambda = \Lambda_Y \subset \partial\mathbb{H}^2$ denote the limit set of Γ . As usual, $\text{CH}(\Lambda)/\Gamma$ can be identified with $Y' = Y \setminus \partial_0 Y$. After isotopy we may assume $\hat{i} : Y' \rightarrow \text{CH}(\Lambda)/\Gamma$ is this identification.

First assume that Y is not an annulus. Form $\text{CH}_q(\Lambda)$ as above, and for a boundary geodesic l of $\text{CH}(\Lambda)$ define l_q and its side \mathcal{D}_l as in the discussion above. The quotient of l_q is a geodesic representative of a component of ∂Y , and the quotient of the open side \mathcal{D}_l° in X_Y is either an open annulus or a union of open disks joined in a cycle along points in their completion. The q -geodesic may pass through points of $\hat{\mathcal{P}}$, so that there is a homotopy from l to l_q rel endpoints which stays in \mathbb{H}^2 until the last instant.

We may equivariantly deform the identity to a map $\text{CH}(\Lambda) \rightarrow \text{CH}_q(\Lambda)$, which takes each l to l_q : since $\text{CH}_q(\Lambda)$ is contractible, it suffices to give a Γ -invariant triangulation of $\text{CH}(\Lambda)$ and define the homotopy successively on the skeleta. This homotopy descends to a map from Y' to $\text{CH}_q(\Lambda)/\Gamma$, and can be chosen so that the puncture-parallel boundary components map to the corresponding points of \mathcal{P}_Y . This gives the desired map \hat{i}_q and establishes properties (1-4).

Using the description of the sides \mathcal{D}_l , we may equivariantly retract $\overline{\mathbb{H}^2}$ to $\text{CH}_q(\Lambda)$, giving rise to the retraction r of part (5).

Finally, if the interior of $\text{CH}_q(\Lambda)$ is a disk, then its quotient is a surface. Our homotopy yields a homotopy-equivalence of Y' to this surface which preserves peripheral structure and can therefore be deformed rel boundary to a homeomorphism. We let $\hat{\iota}_q$ be this homeomorphism, giving part (6).

When Y is a (nonperipheral) annulus, Λ_Y is a pair of points and we recall from above that $\text{CH}_q(\Lambda)$ is either a flat strip in \hat{X} which descends to a flat cylinder in \bar{X}_Y , or it is a single geodesic. The proof in the annular case now proceeds exactly as above. \square

Let $\iota_q : Y \rightarrow \bar{X}$ be the composition of $\hat{\iota}_q$ with the (branched) covering $\bar{X}_Y \rightarrow \bar{X}$ and set $\partial_q Y = \iota_q(\partial' Y)$. Note that this will be a 1-complex of saddle connections and not necessarily a homeomorphic image of $\partial' Y$.

2.4. Fibered faces of the Thurston norm. A fibration $\sigma : M \rightarrow S^1$ of a finite-volume hyperbolic 3-manifold M over the circle comes with the following structure: there is an 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 a representation of M as a quotient $F \times \mathbb{R}/\Phi$ where $\Phi(x, t) = (f(x), t - 1)$ and $f : F \rightarrow F$ is called the monodromy map. This map is pseudo-Anosov and has stable and unstable (singular) measured foliations λ^+ and λ^- on F . Finally there is the suspension flow inherited from the natural \mathbb{R} action on $F \times \mathbb{R}$, and suspensions Λ^\pm of λ^\pm which are flow-invariant 2-dimensional foliations of M . All these objects are defined up to isotopy.

The fibrations of M are organized by the *Thurston norm* $\|\cdot\|$ on $H^1(M; \mathbb{R})$ [Thu86] (see also [CC00]). This norm has a polyhedral unit ball B with the following properties:

- (1) Every cohomology class dual to a fiber is in the cone $\mathbb{R}_+ \mathcal{F}$ over a top-dimensional open face \mathcal{F} of B .
- (2) If $\mathbb{R}_+ \mathcal{F}$ contains a cohomology class dual to a fiber then *every* irreducible integral class in $\mathbb{R}_+ \mathcal{F}$ is dual to a fiber. \mathcal{F} is called a *fibered face* and its irreducible 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 a bit by saying that a fiber (rather than its Poincaré dual) is in $\mathbb{R}_+ \mathcal{F}$.

The fibered faces also organize the suspension flows and the stable/unstable foliations: If \mathcal{F} is a fibered face then there is a single flow ψ and a single pair Λ^\pm of foliations whose leaves are invariant by ψ , such that *every* fibration associated to $\mathbb{R}_+ \mathcal{F}$ may be isotoped so that its suspension flow is ψ up to a reparameterization, and the foliations λ^\pm for the monodromy of its fiber F are $\Lambda^\pm \cap F$. These results were proven by Fried [Fri82]; see also McMullen [McM00].

Veering triangulation of a fibered face. A key fact for us is that the veering triangulation of the manifold M depends only on the fibered face \mathcal{F} and not on a

particular fiber. This was known to Agol for his original construction (see sketch in [Ago12]), but Guéritaud's construction makes it almost immediate.

Proposition 2.7 (Invariance of τ). *Let M be a hyperbolic 3-manifold with fully-punctured fibered face \mathcal{F} . Let S_1 and S_2 be fibers of M each contained in $\mathbb{R}_+\mathcal{F}$ and let τ_1 and τ_2 be the corresponding veering triangulations of M . Then, after an isotopy preserving transversality to the suspension flow, $\tau_1 = \tau_2$.*

Proof. The suspension flow associated to \mathcal{F} lifts to the universal cover \widetilde{M} , and any fiber S in $\mathbb{R}_+\mathcal{F}$ is covered by a copy of its universal cover \widetilde{S} in \widetilde{M} which meets every flow line transversely, exactly once. Thus we may identify \widetilde{S} with the leaf space \mathcal{L} of this flow. The lifts $\widetilde{\Lambda}^\pm$ of the suspended laminations project to the leaf space where they are identified with the lifts $\widetilde{\lambda}^\pm$ of λ^\pm to \widetilde{S} .

The foliated rectangles used in the construction of τ from \tilde{q} on \widetilde{S} depend only on the (unmeasured) foliations $\widetilde{\lambda}^\pm$. Thus the abstract cell structure of τ depends only on the fibered face \mathcal{F} and not on the fiber. The map π from each tetrahedron to its rectangle does depend a bit on the fiber, as we choose q -geodesics for the edges (and the metric q depends on the fiber); but the edges are always mapped to arcs in the rectangle that are transverse to both foliations. It follows that there is a transversality-preserving isotopy between the triangulations associated to any two fibers. \square

Fibers and projections. We next turn to a few lemmas relating subsurface projections over the various fibers in a fixed face of the Thurston norm ball.

Lemma 2.8. *If \mathcal{F} is a fibered face for M and $Y \rightarrow S$ is an infinite covering where S is a fiber in $\mathbb{R}_+\mathcal{F}$ and $\pi_1(Y)$ is finitely generated, then the projection distance $d_Y(\lambda^-, \lambda^+)$ depends only on \mathcal{F} and the conjugacy class of the subgroup $\pi_1(Y) \leq \pi_1(M)$ (and not on S).*

Note that Y need not correspond to an embedded subsurface of S .

Proof. As in the proof of Proposition 2.7, \widetilde{S} can be identified with the leaf space \mathcal{L} of the flow in \widetilde{M} . The action of $\pi_1(M)$ on \widetilde{M} descends to \mathcal{L} , and thus the cover $Y = \widetilde{S}/\pi_1(Y)$ is identified with the quotient $\mathcal{L}/\pi_1(Y)$ and the lifts of λ^\pm to Y are identified with the images of $\widetilde{\lambda}^\pm$ in $\mathcal{L}/\pi_1(Y)$. Thus the projection $d_Y(\lambda^+, \lambda^-)$ can be obtained without reference to the fiber S . \square

This lemma justifies the notation $d_Y(\Lambda^+, \Lambda^-)$ used in the introduction.

We will also require the following lemma, where we allow maps homotopic to fibers which are not necessarily embeddings.

Lemma 2.9. *Let F be a fiber of M . Let $Y \subset M$ be a compact surface and let $h: F \rightarrow M$ be a map which is homotopic to the inclusion. Suppose that $h(F) \cap Y$ is inessential in Y , i.e. each component of the intersection is homotopic into the ends of Y . Then the image of $\pi_1(Y)$ is contained in $\pi_1(F) \triangleleft \pi_1(M)$.*

Proof. Let ζ be the cohomology class dual to F . Since $h(F)$ meets Y inessentially, every loop in Y can be pushed off of $h(F)$ so ζ vanishes on $\pi_1(Y)$. But the kernel of ζ in $\pi_1(M)$ is exactly $\pi_1(F)$, so the image of $\pi_1(Y)$ is in $\pi_1(F)$. \square

3. SECTIONS AND POCKETS OF THE VEERING TRIANGULATION

In this section the surface X is fully-punctured. A *section* of the veering triangulation τ is an embedding $(X, T) \rightarrow (X \times \mathbb{R}, \tau)$ which is simplicial with respect to an ideal triangulation T of X , and is a section of the fibration $\pi: X \times \mathbb{R} \rightarrow X$ (hence transverse to the vertical flow). By *simplicial* we mean that the map takes simplices to simplices. The edges of T are saddle connections of q that are also edges of τ (i.e. those which span singularity-free rectangles), and indeed any triangulation by τ -edges gives rise to a section. We will abuse terminology a bit by letting T denote both the triangulation and the section.

A *diagonal flip* $T \rightarrow T'$ between sections is an isotopy that pushes T through a single tetrahedron of τ , either above it or below it. Equivalently, if R is a maximal rectangle and Q its associated tetrahedron, the bottom two faces of Q might appear in T , in which case T' would be obtained by replacing these with the top two faces. This is an upward flip, and the opposite is a downward flip. We will refer to the transition as both a *diagonal flip/exchange* and a *tetrahedron move*, depending on the perspective.

An edge e of T can be flipped downward exactly when it is the tallest edge, with respect to q , among the edges in either of the two triangles adjacent to it. This makes e the top edge of a tetrahedron (i.e. the diagonal of a quadrilateral that connects the horizontal sides of the corresponding rectangle). Similarly it can be flipped upward when it is the widest edge among its neighbors. See [Figure 7](#).

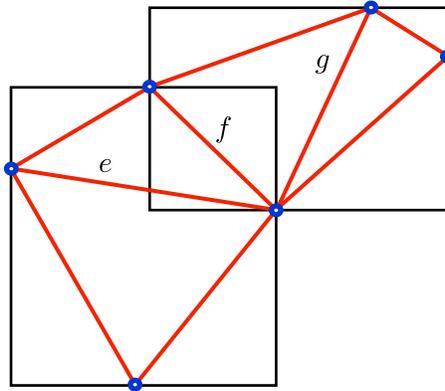


FIGURE 7. The edge e is upward flippable, g is downward flippable, and f is not flippable.

In particular it follows that every section has to admit both an upward and downward flip – simply find the tallest edge and the widest edge.

However it is not a priori obvious that a section even exists. Guéritaud gives an argument for this and more:

Lemma 3.1 ([Gué15]). *There is a sequence of sections $\cdots \rightarrow T_i \rightarrow T_{i+1} \rightarrow \cdots$ separated by upward diagonal flips, which sweeps through the entire manifold $(X \times \mathbb{R}, \tau)$. Moreover, when $(X \times \mathbb{R}, \tau)$ covers the manifold (M, τ) , this sequence is invariant by the deck translation Φ .*

We remark that Agol had previously proven a version of Lemma 3.1 with his original definition of the veering triangulation [Ago11].

For an alternative proof that sections exist, see the second proof of Lemma 3.2. We remark that Lemma 3.1 does not give a complete picture of all possible sections of τ . In this section we will establish a bit more structure.

For a subcomplex $K \leq \tau$, denote by $T(K)$ the collection of sections T of τ containing the edges of K . A necessary condition for $T(K)$ to be nonempty is that $\pi(K)$ is an embedded complex in X composed of τ -simplices. We will continue to blur the distinction between K and $\pi(K)$.

Our first result states that the necessary condition is sufficient:

Lemma 3.2 (Extension lemma). *Suppose that E is a collection of τ -edges in X with pairwise disjoint interiors. Then $T(E)$ is nonempty.*

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

Proposition 3.3 (Connectivity). *If K is a collection of τ -edges in X with pairwise disjoint interiors, then $T(K)$ is connected via tetrahedron moves.*

Finding flippable edges. Let T be a section and let σ be an edge of τ , which is not an edge of T . Any edge e of T crossing σ must do so from top to bottom ($e > \sigma$) or left to right ($e < \sigma$), as in Section 2.1, and we further note that all edges of T that cross σ do it consistently, all top-bottom or all left-right, since they are disjoint from each other.

Lemma 3.4. *Let T be a section and suppose that an edge σ of τ is crossed by an edge e of T . If $e > \sigma$, then there is an edge of T crossing σ which is downward flippable. Similarly if $e < \sigma$ then there is an edge of T crossing σ which is upward flippable.*

Proof. Assuming the crossings of σ are top to bottom, let e be the edge crossing σ that has largest height with respect to q . Let D be a triangle of T on either side of e and let f be its tallest edge. Drawing the rectangle M in which D is inscribed (Figure 8) one sees that R , the rectangle of σ , is forced to cross it from left to right. Hence, the edge f must also cross σ . Therefore, $f = e$ by choice of e . It follows that e is a downward flippable edge. \square

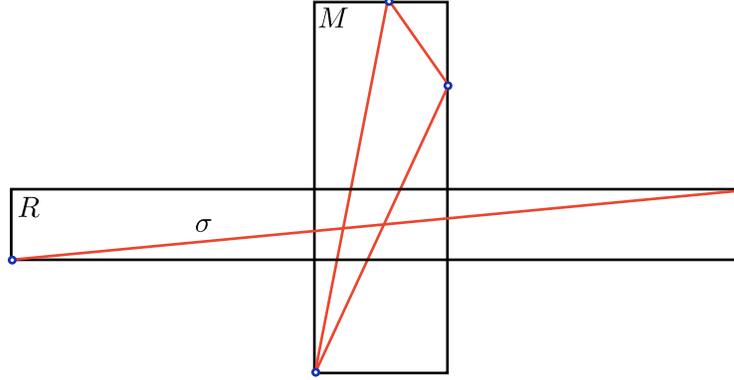


FIGURE 8. The tallest T -edge crossing σ must also be tallest in its own triangles.

Pockets. Let T and T' be two sections and K their intersection, as a subcomplex in $X \times \mathbb{R}$. Because both sections are embedded copies of X transverse to the suspension flow, their union $T \cup T'$ divides $X \times \mathbb{R}$ into two unbounded regions and some number of bounded regions. Each bounded region U is a union of tetrahedra bounded by two isotopic subsurfaces of T and T' , which correspond to a component W of the complement of $\pi(K)$ in X . The isotopy is obtained by following the flow, and if it takes the subsurface of T' upward to the subsurface of T we say that T lies above T' in U . We call U a *pocket over W* , and sometimes write U_W . We call W the *base* of the pocket U .

Lemma 3.5. *With notation as above, T lies above T' in the pocket U_W if and only if, for every edge e of T in W and edge e' of T' in W , if e and e' cross then $e > e'$.*

Note that, for each edge e of T in W there is in fact an edge e' of T' in W which crosses e , since both T and T' are triangulations, with no common edges in W .

Proof. Suppose that T lies above T' in U_W and let e be an edge of T in W ; hence, it is in the top boundary of U . Let Q be the tetrahedron of τ for which e is the top edge. Via the local picture around e (see Section 2.1 and Figure 4), we see that Q lies locally below T . Its interior is of course disjoint from T and T' (and the whole 2-skeleton), hence it is inside U . Let e_1 be the bottom edge of Q . Note $e > e_1$. If e_1 is in T' , stop (with $e' = e_1$). Otherwise it is in the interior of U , and we can repeat with the tetrahedron for which e_1 is the top edge. We get a sequence of steps terminating in some e' in T' , which must be in the boundary of U , and conclude $e > e'$ (by the transitivity of $>$ as in Section 2.1). Now from the paragraph before Lemma 3.4, the same slope relation holds for every edge of T' crossing e , hence giving the first implication of the lemma. For the other direction, exchange the roles of T and T' in the proof. \square

Connectedness of $T(K)$. We can now prove Proposition 3.3.

Proof. Let us consider T, T' in $T(K)$. Let U be one of the pockets, and suppose T lies above T' in U . [Lemma 3.5](#) together with [Lemma 3.4](#) implies that T has a downward flippable edge e which crosses an edge of T' that is in W . In particular e itself is in W . Performing this flip we reduce the number of tetrahedra contained in pockets. Thus a finite number of moves will take T to T' , without disturbing K . \square

As a consequence of [Proposition 3.3](#) and its proof we have:

Corollary 3.6. *If K is a nonempty subcomplex of τ and $T(K) \neq \emptyset$, then there are unique sections $T^+(K)$ and $T^-(K)$ in $T(K)$ such that every $T \in T(K)$ can be upward flipped to $T^+(K)$ and downward flipped to $T^-(K)$.*

Proof. First note that $T(K)$ is finite: because τ is locally finite at the edges, there are only finitely many choices for a triangle adjacent to K . We then enlarge K successively, noting that there is a bound on the number of triangles in a section. Thus there exists a section T^+ in $T(K)$ which is not upward flippable in $T(K)$. For any two sections $T_1, T_2 \in T(K)$ there is a $T_3 \in T(K)$ obtained as the union of the tops of the pockets of T_1 and T_2 and their intersection. Thus T_1 is upward flippable unless $T_1 = T_3$, and similarly for T_2 . This implies that T^+ is the unique section in $T(K)$ which is not upward flippable, and every other section is upward flippable to T^+ . We define T^- analogously. \square

The section $T^+(K)$ is called the *top of $T(K)$* and the section $T^-(K)$ is called the *bottom of $T(K)$* . Note that any section obtained from $T^+(K)$ by upward diagonal exchanges is not in $T(K)$.

Extension lemma. We conclude this section with two proofs of [Lemma 3.2](#).

Proof one. [Lemma 3.1](#) gives us, in particular, the existence of at least one section T_0 which is disjoint from E , which we may assume lies above every edge of E .

Then by [Lemma 3.4](#) there is a downward flippable edge e in T_0 . The tetrahedron involved in the move lies above E , so E still lies below (or is contained in) the new section T_1 . We repeat this process, and at each stage every edge of E is either contained in T_i or crosses an edge of T_i and lies below it. Thus by [Lemma 3.4](#), unless $E \subset T_i$ each T_i contains a downward flippable edge that is not contained in E .

Because τ is locally finite at each edge, *any* sequence of downward flips is a proper sweepout of the region below T_0 , and hence must eventually meet every edge of τ below T_0 . Thus we may continue until every edge of E lies in T_i . \square

Proof two. Our second proof does not use [Lemma 3.1](#), and in particular it gives an independent proof of the existence of sections.

Let D be a component of the complement of E which is not a triangle. Let e be an edge of ∂D and consider the collection of τ -tetrahedra adjacent to e . These contain a sequence $Q_-, Q_1, \dots, Q_m, Q_+$, as in [Figure 4](#), where Q_- is the tetrahedron with e as its top edge, Q_+ is the tetrahedron with e as its bottom edge, and the rest are adjacent to e on the same side as D (if D meets e on two sides we just choose one). Two successive tetrahedra in this sequence share a triangular face. We claim

that one of these faces must be contained in D . Equivalently we claim that one of the triangles is not crossed by any edge of E .

Since each tetrahedron Q is inscribed in a singularity free rectangle R , if an edge f of E crosses any edge of Q its rectangle crosses all of R . It follows immediately, since the edges of E have disjoint interiors, that they consistently cross R all vertically, or all horizontally. Because successive tetrahedra in the sequence share a face it follows inductively that, if all the faces are crossed by E , then they are all consistently crossed horizontally, or all vertically.

However, Q_- can only be crossed vertically by E (since E does not cross e). Similarly Q_+ can only be crossed horizontally. It follows that there must be a triangular face F that is *not* crossed by E . Thus F is contained in D . Since D is not a triangle, at least one edge of F passes through the interior of D . We add this edge to E and proceed inductively. \square

4. RECTANGLE AND TRIANGLE HULLS

In this section we discuss a number of constructions that associate a configuration of τ -edges to a saddle connection of the quadratic differential q . These will be used later to show that subsurfaces with large projection are compatible with the veering triangulation in the appropriate sense. As a byproduct of our investigation, we prove the (to us) unexpected result ([Theorem 1.4](#)) that the edges of the veering triangulation form a totally geodesic subgraph of the curve and arc graph of X .

We emphasize that in [Section 4.1](#) and [Section 4.2](#), the surface X is not necessarily fully-punctured. Thus by τ we mean the veering triangulation associated to the fully-punctured surface $X \setminus \text{sing}(q)$. We will say that a saddle connection of X is a τ -edge if its interior is an edge of this veering triangulation. In particular this means that its lift to \tilde{X} spans a singularity-free rectangle.

4.1. Maximal rectangles along a saddle connection. Let σ be a saddle connection, for the moment in the completed universal cover \hat{X} . Consider the set $\mathcal{R}(\sigma)$ of all rectangles which are *maximal with respect to the property that σ passes through a diagonal*. Thus each $R \in \mathcal{R}(\sigma)$ contains singularities in at least two edges. Let $h(R)$ be the convex hull in R of the singularities in the boundary of R and let $h^{(1)}(R)$ denote its 1-skeleton (see [Figure 9](#)).

Let

$$\mathbf{r}(\sigma) = \bigcup \{h^{(1)}(R) : R \in \mathcal{R}(\sigma)\}.$$

See [Figure 10](#) for an example. Note that all the saddle connections in $\mathbf{r}(\sigma)$ are edges of τ — each of these arcs spans a singularity-free rectangle by construction. Moreover, $\mathbf{r}(\sigma) = \{\sigma\}$ if σ is itself a τ -edge.

The following lemma will play an important role throughout this paper.

Lemma 4.1. *If saddle connections σ_1 and σ_2 have no transversal intersections then neither do $\mathbf{r}(\sigma_1)$ and $\mathbf{r}(\sigma_2)$.*

Proof. Say that two rectangles meet *crosswise* if their interiors intersect, and no corners of one are in the interior of the other. Note that when two distinct rectangles

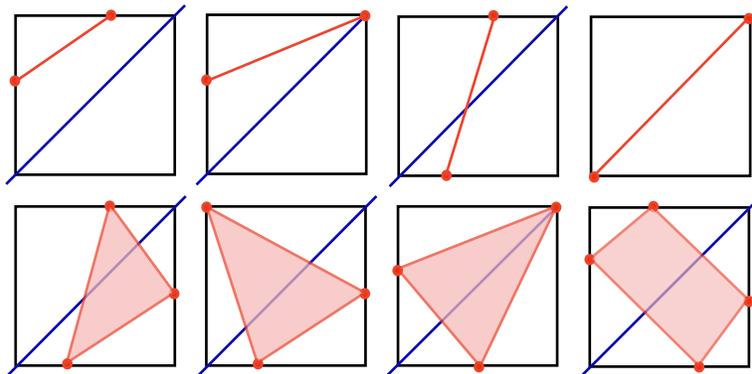


FIGURE 9. The eight possible (up to symmetry) convex hulls $h(R)$, assuming at most one singularity per leaf of λ^\pm . The saddle connection σ is in blue.

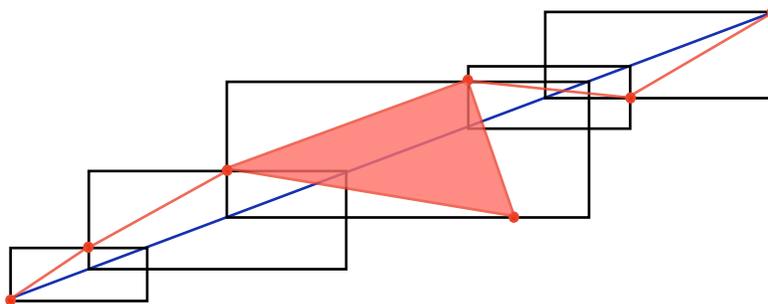


FIGURE 10. Example of $\mathbf{r}(\sigma)$ (in red)

meet crosswise, any two of their diagonals intersect. We say that the rectangles meet *properly crosswise* if they also do not share any corners, in which case any two diagonals intersect in the interior.

Let τ_1 and τ_2 be saddle connections in $\mathbf{r}(\sigma_1)$ and $\mathbf{r}(\sigma_2)$, respectively, and suppose that they intersect transversely. Hence their spanning rectangles Q_1 and Q_2 must cross as in [Figure 3](#). Assume that Q_1 is the taller and Q_2 the wider.

Now let R_1 and R_2 be the rectangles of $\mathcal{R}(\sigma_1)$ and $\mathcal{R}(\sigma_2)$ containing Q_1 and Q_2 , respectively. Because of the singularities in the corners of Q_1 and Q_2 , R_2 is no taller than Q_1 and R_1 is no wider than Q_2 . Hence R_1 and R_2 meet crosswise. (See [Figure 11](#)).

If they met properly crosswise then σ_1 and σ_2 would have an interior intersection, which is a contradiction. Hence R_1 and R_2 share a corner c . But the edges meeting at c would have to pass through boundary edges of Q_1 and Q_2 . Those edges already have the singularities of τ_1 and τ_2 , and so c cannot be a singularity. Thus if c is the intersection of the diagonals contained in σ_1 and σ_2 it would be in the interior of both saddle connections, again a contradiction.

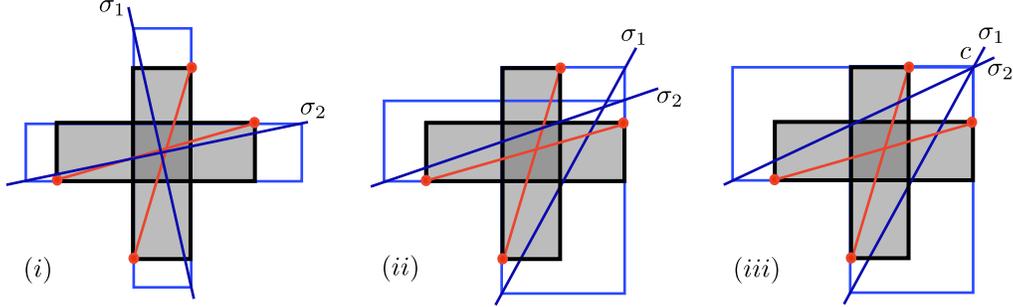


FIGURE 11. Three examples of the crossing pattern. The rectangles R_1 and R_2 are in blue, τ_1 and τ_2 are in red, and Q_1 and Q_2 are shaded. In (i) and (ii) the crossing is proper. In (iii) the corner c is shared.

We conclude that τ_1 and τ_2 cannot cross. \square

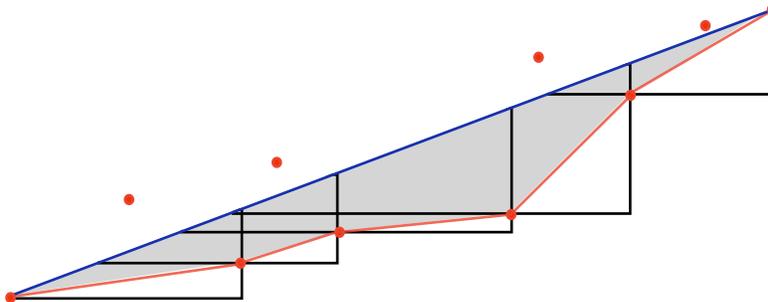
An immediate consequence of [Lemma 4.1](#) is that we can carry out the construction downstairs: If σ is a saddle connection in \bar{X} we can construct $\mathbf{r}(\hat{\sigma})$ for each of its lifts $\hat{\sigma}$ to \hat{X} , and the lemma tells us none of them intersect transversally. Thus the construction projects downstairs to give a collection of τ -edges with disjoint interior. Moreover if K is *any* collection of saddle connections with disjoint interiors then $\mathbf{r}(K)$ makes sense as a subcomplex of τ supported on some section by [Lemma 3.2](#). Hence, we will continue to use $\mathbf{r}(\cdot)$ to denote the corresponding map on saddle connections of \bar{X} . We remark that although $\mathbf{r}(\cdot)$ takes collections of saddle connections with disjoint interiors to collections of τ -edges with disjoint interiors, it may do so with multiplicity.

4.2. Triangle hulls. Now let us consider a similar operation that uses right triangles instead of rectangles, and associates to a transversely oriented saddle connection in the universal cover a homotopic path of saddle connections.

If σ is a saddle connection in \hat{X} equipped with a transverse orientation, let $\mathcal{T}(\sigma)$ denote the collection of Euclidean right triangles which are *maximal with respect to the property that they are attached along the hypotenuse to σ along the side given by its transverse orientation*. A triangle t in $\mathcal{T}(\sigma)$ must have exactly one singularity in each of its legs, and so their convex hull $h(t)$ is a single saddle connection. The set $\mathcal{T}(\sigma)$ must be finite, and its hypotenuses cover σ in a sequence of non-nested intervals, ordered by their left (or right) endpoints. See [Figure 12](#). Let $\mathbf{t}(\sigma)$ be the union of segments $h(t)$ for $t \in \mathcal{T}(\sigma)$.

Lemma 4.2. *Either $\mathbf{t}(\sigma) = \sigma$ or $\sigma \cup \mathbf{t}(\sigma)$ is the boundary of an embedded Euclidean polygon $P(\sigma)$ in \hat{X} which is foliated by arcs of λ^\pm .*

Proof. Suppose that t and t' are triangles of $\mathcal{T}(\sigma)$ and $p \in t \cap t'$ is in the interior of t . Let l and l' be the vertical line segments in t and t' , respectively, joining p to the respective hypotenuses (l' could be a single point). If l and l' leave p in


 FIGURE 12. An example of $\mathbf{t}(\sigma)$ and $P(\sigma)$.

opposite directions then $l \cup l'$ is a vertical geodesic connecting two points of σ , which contradicts the uniqueness of geodesics in \hat{X} . If they leave p in the same direction but are not equal, then their difference is a vertical geodesic with endpoints on σ , again a contradiction.

We conclude that if t and t' intersect they do so on a common subarc of their hypotenuses. This subarc spans a (nonmaximal) right triangle which is exactly $t \cap t'$.

Now given $t \in \mathcal{T}(\sigma)$, the vertical and horizontal legs of t each contain a single singularity of \hat{X} ; denote these singularities by v_t and h_t , respectively. By construction of $\mathcal{T}(\sigma)$, there is a unique triangle $t' \in \mathcal{T}(\sigma)$ such that $h_{t'} = v_t$, unless v_t is an endpoint of σ . Hence, given an orientation on σ , the edges of $\mathbf{t}(\sigma)$ come with a natural ordering induced by moving along σ . By our observations above, we see that $\mathbf{t}(\sigma)$ is an embedded arc and meets σ only at its endpoints. Since \hat{X} is contractible, σ and $\mathbf{t}(\sigma)$ must be homotopic and hence cobound a disk $P(\sigma)$. In fact this disk is foliated by both λ^+ and λ^- , as we can see by noting that each edge of $\mathbf{t}(\sigma)$ cobounds a vertical (similarly a horizontal) strip with a segment in σ . Hence $P(\sigma)$ admits an isometry to a polygon in \mathbb{R}^2 . \square

Let us define a map $\mathbf{t}_\sigma^+ : \sigma \rightarrow \mathbf{t}(\sigma)$ (resp. \mathbf{t}_σ^-) which is the result of pushing the points of σ along the vertical (resp. horizontal) foliation to the other side of $P(\sigma)$.

If $f : I \rightarrow \hat{X}$ is an embedding of an oriented 1-manifold I that parametrizes some union of saddle connections, we let

$$(1) \quad \mathbf{t}^+ f : I \rightarrow \hat{X}$$

be the map that sends each $p \in I$ to $\mathbf{t}_\sigma^+(f(p))$, where σ is the saddle connection containing $f(p)$ with transverse orientation induced by the orientation on I . By composing with covering maps we can use the same notation for the resulting operation in quotients \hat{X}_Y or \bar{X} .

Unlike the rectangle hulls, the edges of $\mathbf{t}(\sigma)$ are not necessarily τ -edges. (See the upper-right red saddle connection in Figure 12.) Moreover, the \mathbf{t} -version of Lemma 4.1 is in general not true. That is, the image of \mathbf{t} may not project to an embedded complex in \bar{X} since σ_1 and σ_2 can be disjoint while $\mathbf{t}(\sigma_1)$ and $\mathbf{t}(\sigma_2)$ cross. However, we do have the following:

Lemma 4.3. *Let σ, σ' be saddle connections in \hat{X} with disjoint interiors. Let l be an arc of λ^+ with endpoints on σ and σ' , and give σ and σ' the transverse orientation pointing toward the interior of l . Then the polygons $P(\sigma)$ and $P(\sigma')$ of \hat{X} (from Lemma 4.2) have disjoint interiors.*

Proof. Suppose towards a contradiction that there is a point p which is in the interior of each of the polygons $P = P(\sigma)$ and $P' = P(\sigma')$. Since P and P' are foliated by λ^+ , let m and m' be the arcs of λ^+ which are properly embedded in P and P' respectively, and pass through p . Orient m so that it begins in σ , and m' so that it terminates in σ' . These orientations agree at p : if they did not we would obtain a contradiction by applying Gauss–Bonnet to the circuit passing through m, σ, l, σ' and m' .

Thus, the union $J = m \cup m'$ is an interval in a leaf of λ^+ with endpoints on σ and σ' , with p in the interior of $m \cap m'$. (If p were in l already then we would have $J = l$.) Orienting J as $[y, y']$ where $y \in \sigma$ and $y' \in \sigma'$, we can write $m = [y, x]$ and $m' = [x', y']$, where $x = J \cap \mathbf{t}(\sigma)$ and $x' = J \cap \mathbf{t}(\sigma')$. These points appear, in order along J , as y, x', p, x, y' .

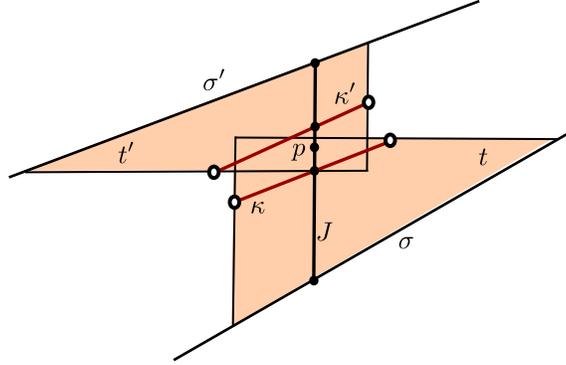


FIGURE 13. The point p cannot lie in the interior of both $P(\sigma)$ and $P(\sigma')$.

Let t and t' be the triangles of $\mathcal{T}(\sigma)$ and $\mathcal{T}(\sigma')$ containing x and x' , respectively. Then $p \in t \cap t'$. Let κ and κ' be the saddle connections of $\mathbf{t}(\sigma)$ and $\mathbf{t}(\sigma')$ spanning t and t' , respectively (See Figure 13). The fact that the endpoints of κ and κ' are disjoint from the intersection of t and t' implies that $\kappa \cap J$, which is x , lies below $\kappa' \cap J$, which is x' . This contradicts the ordering of the points in J . \square

4.3. Retractions in \mathcal{A} . In this subsection, X is fully-punctured. Let $\mathcal{A}(\tau) \subset \mathcal{A}(X)$ be the span of the vertices of $\mathcal{A}(X)$ which are represented by edges of τ . We will construct a *coarse 1-Lipschitz retraction* from $\mathcal{A}(X)$ to $\mathcal{A}(\tau)$. By this, we mean a coarse map which takes diameter ≤ 1 sets to diameter ≤ 1 sets and restricts to the identity on the 0-skeleton of $\mathcal{A}(\tau) \subset \mathcal{A}(X)$.

First, let $\mathcal{SC}(q) \subset \mathcal{A}(X)$ be the arcs of X which can be realized by saddle connections of q . Hence, $\mathcal{A}(\tau) \subset \mathcal{SC}(q) \subset \mathcal{A}(X)$. For any $a \in \mathcal{A}(X)$ define $\mathbf{s}(a) \subset \mathcal{SC}(q)$

as follows: If a_q is the q -geodesic representative of a in \bar{X} , then let $\mathbf{s}(a)$ be the set of saddle connections of q composing a_q . If a is a cylinder curve of q , then we take $\mathbf{s}(a)$ to be the set of saddle connections appearing in the boundary of the maximal cylinder of a . Note that if $a \in \mathcal{A}(X)$ is itself represented by a saddle connection of q , then $\mathbf{s}(a) = \{a\}$.

The following lemma shows that \mathbf{s} is well-defined and is a coarse 1-Lipschitz retraction, in the above sense.

Lemma 4.4. *For adjacent vertices $a, b \in \mathcal{A}(X)$, the vertices of $\mathbf{s}(a)$ and $\mathbf{s}(b)$ are pairwise adjacent or equal.*

Proof. Recall that adjacency of vertices in $\mathcal{A}(X)$ corresponds to disjointness of their hyperbolic geodesic representative, and for vertices realized by saddle connections, this corresponds to the lack of transverse intersection of their interiors. But if any arcs of $\mathbf{s}(a)$ and $\mathbf{s}(b)$ have crossing interiors, [Corollary 2.4](#) implies that the hyperbolic geodesic representatives of a and b must cross as well. The lemma follows. \square

Combining this lemma with [Lemma 4.1](#) gives us the proof of [Theorem 1.4](#), which we restate here in somewhat more precise language:

Theorem 1.4 (Geodesically connected theorem). *Let (X, q) be fully punctured with associated veering triangulation τ . The composition $\mathbf{r} \circ \mathbf{s}: \mathcal{A}(X) \rightarrow \mathcal{A}(\tau)$ is a coarse 1-Lipschitz retraction in the sense that it takes diameter ≤ 1 sets to diameter ≤ 1 sets, and is the identity on the 0-skeleton of $\mathcal{A}(\tau)$. Hence, any two vertices in $\mathcal{A}(\tau)$ are joined by a geodesic of $\mathcal{A}(X)$ that lies in $\mathcal{A}(\tau)$.*

Proof. [Lemma 4.4](#) says that $\mathbf{s}: \mathcal{A}(X) \rightarrow \mathcal{SC}(q)$ is a coarse 1-Lipschitz retraction. [Lemma 4.1](#), interpreted as a statement about the arc and curve complexes, says the same for $\mathbf{r}: \mathcal{SC}(q) \rightarrow \mathcal{A}(\tau)$. The theorem follows. \square

5. PROJECTIONS AND COMPATIBLE SUBSURFACES

In this section we show that if $Y \subset X$ is a compact essential subsurface of large projection distance $d_Y(\lambda^+, \lambda^-)$, then Y has particularly nice representations with respect to, first, the quadratic differential q and, second, the veering triangulation τ . We emphasize that in this section, the surface X is not necessarily fully-punctured.

5.1. Projection and q -compatibility. Recall the q -convex hull map $\hat{\iota}_q: Y \rightarrow \bar{X}_Y$ constructed in [Lemma 2.6](#). We say that Y is q -compatible if $\hat{\iota}_q$ is an embedding of $Y' = Y \setminus \partial_0 Y$, as in part (6) of [Lemma 2.6](#). (Recall that $\partial_0 Y$ maps to completion points of $\tilde{\mathcal{P}}_Y$). This condition implies a little more:

Lemma 5.1. *If $Y \subset X$ is q -compatible, then*

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

Recall that $\partial' Y = \partial Y \setminus \partial_0 Y$.

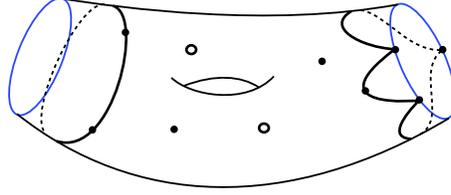


FIGURE 14. The image of a q -compatible subsurface Y in \bar{X}_Y under $\hat{\iota}_q$. Open circles are points of $\tilde{\mathcal{P}}_Y$ (corresponding to the image of $\partial_0 Y$) and dots are singularities not contained in $\tilde{\mathcal{P}}_Y$. The ideal boundary of X_Y is in blue.

Proof. Recall from [Lemma 2.6](#) that q -compatibility of Y is equivalent to the statement that the interior of the q -hull $\text{CH}_q(\Lambda) \subset \hat{X}$ is a disk (i.e. it is not pinched along singularities or saddle connections).

If $\iota_q: \text{int}(Y) \rightarrow X$ fails to be an embedding, then it must be that for some deck transformation g of the universal covering $\tilde{X} \rightarrow X$ the interiors of $\text{CH}_q(\Lambda)$ and $g \cdot \text{CH}_q(\Lambda)$ are distinct and overlap. But then it follows immediately from [Corollary 2.4](#) that the distinct hyperbolic convex hulls $\text{CH}(\Lambda)$ and $g \cdot \text{CH}(\Lambda)$ overlap, contradicting that Y is a subsurface of X . This proves part (1).

For part (2), let β be a component of $\partial_0 Y$. Since $\hat{\iota}_q$ embeds $Y \setminus \partial_0 Y$, a collar neighborhood U of β in Y maps to a neighborhood V of the puncture $p = \hat{\iota}_q(\beta)$. Now if γ is a component of $\partial' Y$, q -compatibility again implies its image must avoid $V \setminus p$. Since $\hat{\iota}_q(\gamma)$ cannot equal p , it must be disjoint from it. \square

Note that Y is a q -compatible annulus if and only if the core of Y is a cylinder curve in X . In this case, the corresponding open flat cylinder in X is $\iota_q(\text{int}(Y))$. In general, if Y is q -compatible then one component of $X \setminus \partial_q Y$ is an open subsurface isotopic to the interior of Y ; this is the image $\iota_q(\text{int}(Y))$ and is denoted $\text{int}_q(Y)$.

The following proposition shows that mild assumptions on $d_Y(\lambda^+, \lambda^-)$ imply that Y is q -compatible.

Proposition 5.2 (q -Compatibility). *Let $Y \subset X$ be an essential subsurface.*

If Y is non-annular and $d_Y(\lambda^+, \lambda^-) > 0$, then Y is q -compatible.

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

Proof. We treat the non-annular case first. Suppose that $d_Y(\lambda^+, \lambda^-) > 0$.

Recall from [Section 2.3](#) that we have identified \tilde{X} with \mathbb{H}^2 , set $\Lambda \subset \partial\mathbb{H}^2$ to be the limit set of $\Gamma = \pi_1(Y)$, set $\Omega = \partial\mathbb{H}^2 \setminus \Lambda$, and defined $\hat{\mathcal{P}}_Y \subset \Lambda$ to be the set of parabolic fixed points of $\pi_1(Y)$. Note that $\hat{\mathcal{P}}_Y = \Lambda \cap \hat{\mathcal{P}}$. Further recall from part (6) of [Lemma 2.6](#) that the map from Y' to $\text{CH}_q(X_Y)$ is an embedding, provided the interior of $\text{CH}_q(\Lambda)$ is a disk. Since $\text{CH}_q(\Lambda)$ is the result of deleting the interior of the side \mathcal{D}_l from \hat{X} for each hyperbolic geodesic line l in $\partial\text{CH}(\Lambda)$, it suffices to show that

- (1) for each geodesic line l in $\partial\text{CH}(\Lambda)$, the interior of the corresponding q -geodesic l_q does not meet $\partial\mathbb{H}^2 \setminus \mathcal{D}_l$, and
- (2) if l and l' are distinct geodesic lines in $\partial\text{CH}(\Lambda)$ then l_q and l'_q do not meet in \tilde{X} .

First suppose that condition (1) is violated for some geodesic line l in $\partial\text{CH}(\Lambda)$ and point $\hat{p} \in \partial\mathbb{H}^2 \setminus \mathcal{D}_l$. Set p to be the image of \hat{p} in \tilde{X}_Y . Letting γ be the boundary component of $\partial'Y$ that is the image of l in X_Y , we see that the image of l_q in \tilde{X}_Y , which equals $\gamma_q = \hat{l}_q(\gamma)$, passes through the point p .

Since l_q is a geodesic in \tilde{X} , we see that \hat{p} is a completion point and so either $\hat{p} \in \hat{\mathcal{P}}_Y$ or $\hat{p} \in \hat{\mathcal{P}} \setminus \hat{\mathcal{P}}_Y$.

Assume that $\hat{p} \in \hat{\mathcal{P}}_Y$. Then $p \in \tilde{\mathcal{P}}_Y$ corresponds to a puncture of Y . Recall that by [Lemma 2.6](#), the image of the open side $\mathcal{D}_l^\circ = \text{int}(\mathcal{D}_l \cap \tilde{X})$ in X_Y is either an open annulus or a disjoint union of open disks; in either case, set A_γ equal to the component which contains p in its boundary. The angle at p in A_γ between the incoming and outgoing edges of γ_q is at least π , which implies that A_γ contains a horizontal and a vertical ray l^-, l^+ emanating from p . ([Figure 15](#).)

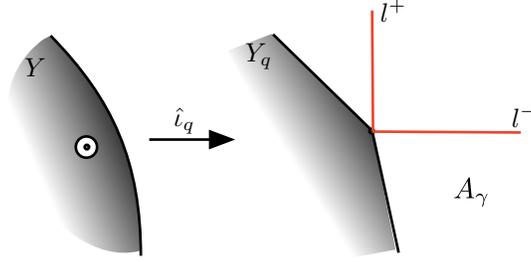


FIGURE 15. When $\hat{l}_q(\partial'Y)$ passes through a point of $\tilde{\mathcal{P}}_Y$, $d_Y(\lambda^+, \lambda^-) = 0$.

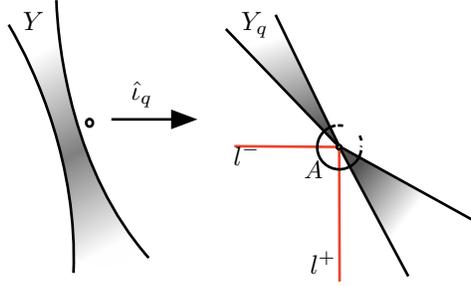
These rays are proper q -geodesic lines in X_Y (because p is a puncture, not a point of X_Y), and hence by [Lemma 2.5](#) represent vertices of $\pi_Y(\lambda^-)$ and $\pi_Y(\lambda^+)$, respectively. Further, since the rays only intersect within the annulus or disk A_γ and Y is itself nonannular, we see that l^- and l^+ in fact represent the same point in $\mathcal{A}(Y)$. (Actually, if A does not contain a flat cylinder, then the interiors of l^- and l^+ are disjoint as we show below). Either way, it follows that

$$d_Y(\lambda^+, \lambda^-) = 0,$$

a contradiction.

Next assume that $\hat{p} \in \hat{\mathcal{P}} \setminus \hat{\mathcal{P}}_Y$. Since $\hat{p} \notin \mathcal{D}_l \cap \partial\mathbb{H}^2$ we may set A to be the component of the image of \mathcal{D}_l° in X_Y which contains $p \in \tilde{X}_Y$ in its boundary. As before, the angle subtended by γ_q at x in the boundary of A is at least π (see [Figure 16](#)). A pair of rays l^\pm emanating from x into A are properly embedded lines and again represent the same vertex of $\mathcal{A}(Y)$, giving us $d_Y(\lambda^+, \lambda^-) = 0$.

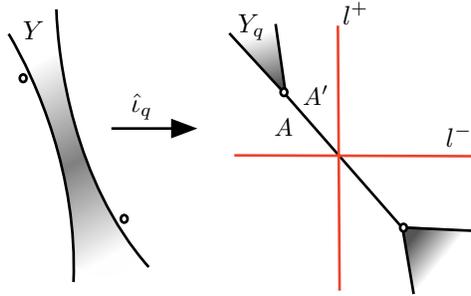
We conclude that condition (1) is satisfied.

FIGURE 16. Y_q is pinched at a completion point.

Next suppose that geodesics l and l' in the boundary of $\text{CH}(\Lambda)$ violate (2), i.e. l_q and l'_q meet in \tilde{X} . Let $\tilde{I} = l_q \cap l'_q \subset \hat{X}$ which, since \hat{X} is $\text{CAT}(0)$, is a connected subset of each of l_q, l'_q . In general, the intersection in \hat{X} of two q -geodesic lines is either a single singularity (possibly a completion point) or a union of saddle connections. Because l_q and l'_q meet in \tilde{X} , \tilde{I} contains either a saddle connection or a singularity which is not a completion point. Let $\gamma, \gamma', \gamma_q, \gamma'_q, I$, be the images in \bar{X}_Y of $l, l', l_q, l'_q, \tilde{I}$, respectively.

Suppose first that I contains a saddle connection σ . In this case, let A be the component of the image of the open side \mathcal{D}_i° in X_Y which contains σ in its boundary, and define A' similarly. (Note that it is possible that $A = A'$ and that A and A' meet along other saddle connections and singularities besides σ , but this will not change the discussion.)

Any point of σ is crossed by a pair l^+, l^- of leaves of λ^+, λ^- , which as proper arcs of X_Y determine the same vertex of $\mathcal{A}(Y)$. Hence, we conclude once again that $d_Y(\lambda^+, \lambda^-) = 0$.

FIGURE 17. Y_q is pinched along a saddle connection.

Finally, suppose that I contains a singularity x in X_Y (i.e. x is not a completion point). Again, set A to be the component of the image of \mathcal{D}_i° in X_Y which contains x in its boundary and A' to be the component of the image of \mathcal{D}_i° in X_Y which

contains x in its boundary. As before, there is an angle of at least π on the A side of γ_q and on the A' side of γ'_q , so we can find pairs of rays r_0^\pm emanating from x on the A side, and r_1^\pm emanating on the A' side (see [Figure 18](#)). The unions $l^+ = r_0^+ \cup_x r_1^+$ and $l^- = r_0^- \cup_x r_1^-$ are generalized leaves of λ^+ and λ^- , respectively, and again determine the same point in $\mathcal{A}(Y)$ so we conclude that $d_Y(\lambda^+, \lambda^-) = 0$.

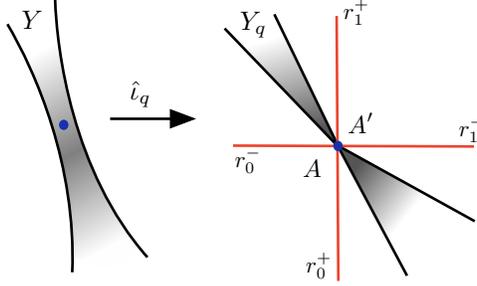


FIGURE 18. Y_q is pinched at a singularity which is not a completion point.

We conclude that if Y is nonannular and $d_Y(\lambda^+, \lambda^-) > 0$, then Y is q -compatible.

When Y is an annulus, almost the same argument applies. The difference is that the arcs l^\pm we obtain are not homotopic with fixed endpoints, and so do not determine the same vertex of $\mathcal{A}(Y)$. However, in each case we will show they have disjoint interiors, concluding $d_Y(l^+, l^-) \leq 1$, and so

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

To see this, let γ denote the core of Y and let γ_q be a geodesic representative in \bar{X}_Y . Supposing that $\text{int}_q(Y)$ is not a flat annulus, we first claim the following: For any singular point p crossed by γ_q , if l^+ and l^- are rays of λ^+ and λ^- , respectively, meeting with angle $\pi/2$ at p , then the interiors of l^+ and l^- do not meet.

To establish the claim, assume that the interiors of l^\pm meet and refer to [Figure 19](#). Let A_γ be the complementary region of γ_q in \bar{X}_Y containing p' , the interior intersection of l^\pm . If A_γ is a disk, then the claim follows immediately from the uniqueness of geodesics in a CAT(0) space. Hence, we may assume that A_γ is an annulus. Let l_ϵ^+ be leaf of λ^+ parallel to l^+ and slightly displaced to the interior of A_γ , so that the region R bounded by γ_q and the segments of l^- and l_ϵ^+ is an annulus. The total curvature of the $l^-l_\epsilon^+$ boundary of R is 0 since it is straight except for two right turns of opposite signs, and the total curvature of γ_q as measured from inside R is nonpositive (since each singularity on γ_q subtends at least angle π within R). Since $\chi(R) = 0$ and the Gaussian curvature in R (including singularities) is nonpositive, the Gauss–Bonnet theorem implies that the total curvature of ∂R is nonnegative. This implies that the total curvature of γ_q is 0, which means that γ_q bounds a flat cylinder, contradicting our assumption. This establishes the claim.

We now return to the proof of the proposition. First suppose that γ_q passes through a completion point x of \bar{X}_Y . Then, just as in [Figure 16](#), we can find a pair

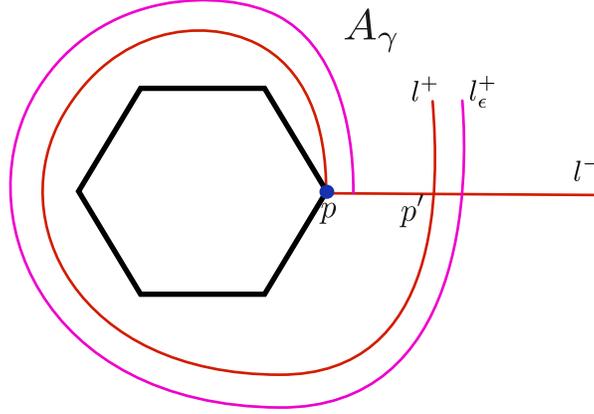


FIGURE 19. The q -geodesic γ_q is the black hexagon. An interior intersection between l^+ and l^- contradicts the Gauss–Bonnet theorem.

of rays l^\pm emanating from x into A_γ . By the claim above, the interiors of these rays do not meet and so $d_Y(\lambda^+, \lambda^-) \leq 1$ as desired.

Finally, suppose that γ_q remains in X_Y , i.e. it does not pass through any completion points. It must still pass through a singularity x , and we note that the total angle at x is at least 3π . Recall that γ_q subtends at least angle π at x to either of its sides and we note that some side of γ sees angle at least $3\pi/2$ at x . Let A denote this side of γ_q and let A' denote the other side. Note that $A \neq A'$ since X_Y is an annulus which γ_q separates. The angle of $3\pi/2$ tells us there are at least 3 rays of λ^\pm emanating into A . Now choose rays r_0^\pm of λ^\pm emanating from x on the A' side. Because the 3 (or more) rays of λ^\pm emanating from x into A alternate between λ^+ and λ^- , we can choose from them two rays r_1^\pm of λ^\pm such that $r_0^+, r_1^+, r_1^-, r_0^-$ are listed in the cyclic ordering of directions at x (either clockwise or counterclockwise). The generalized leaves $l^+ = r_0^+ \cup_x r_1^+$ and $l^- = r_0^- \cup_x r_1^-$ then represent arcs in the projections $\pi_Y(\lambda^+)$ and $\pi_Y(\lambda^-)$ and after a slight perturbation these leaves have disjoint interiors. Hence, again we see that $d_Y(\lambda^+, \lambda^-) \leq 1$.

We conclude that if Y is an annulus with $d_Y(\lambda^+, \lambda^-) \geq 1$ then Y is q -compatible. \square

5.2. Projections and τ -compatibility. We now show how to associate to a subsurface Y of large projection a representative of Y which is “simplicial” with respect to the veering triangulation. This will later be used to prove that such a subsurface induces a pocket of the veering triangulation τ .

Informally, we start with a q -compatible subsurface $Y \subset X$ and homotope $\hat{\iota}_q$ by pushing $\partial_q Y$ onto τ -edges (this process is depicted locally in Figure 21). Formally, this is done in two steps using the map $\mathbf{t}(\cdot)$ described in Section 4.2, although some care must be taken in order to ensure that the resulting object gives an embedded representative of $\text{int}(Y)$.

Call a 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 and
- (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 .

We will show that when $d_Y(\lambda^-, \lambda^+)$ is sufficiently large, the subsurface Y is τ -compatible and in this case we set $\partial_\tau Y = \iota_\tau(\partial' Y)$ 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. 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 $\text{int}_\tau(Y)$.

Theorem 5.3 (τ -Compatibility). *Let $Y \subset X$ be an essential subsurface.*

- (1) *If Y is nonannular and $d_Y(\lambda^+, \lambda^-) > 0$, then Y is τ -compatible.*
- (2) *If Y is an annulus and $d_Y(\lambda^+, \lambda^-) > 1$, then Y is τ -compatible.*

Proof. Suppose that $d_Y(\lambda^+, \lambda^-) > 0$ if Y is nonannular and $d_Y(\lambda^+, \lambda^-) > 1$ otherwise. By [Proposition 5.2](#), Y is q -compatible and so $\hat{\iota}_q : Y \rightarrow \bar{X}_Y$ is an embedding on Y' . Let Y_q denote its image. We first suppose that Y is not an annulus.

Give $\partial' Y$ the transverse orientation pointing into Y . For any saddle connection σ in $\hat{\iota}_q(\partial' Y)$ and any triangle $t \in \mathcal{T}(\sigma)$ pointing into Y (see [Section 4.2](#) for definitions), note that the singularities of \bar{X}_Y in ∂t are *not* completion points of \bar{X}_Y , that is they do not correspond to punctures of X . This is because any completion point lying in t is the endpoint of leaves l^\pm of λ^\pm whose initial segments lie in t . These leaves correspond to essential proper arcs of X_Y which are homotopic giving $d_Y(\lambda^-, \lambda^+) = 0$, a contradiction.

Similarly, we can conclude that for each saddle connection σ in $\hat{\iota}_q(\partial' Y)$ and any $t \in \mathcal{T}(\sigma)$ pointing into Y , the triangle t is entirely contained in Y_q . Otherwise, similar to the proof of [Proposition 5.2](#), we find leaves l^+ and l^- in \bar{X}_Y whose intersection with Y_q is contained in t and hence whose projections to $\mathcal{A}(Y)$ are equal. See the left side of [Figure 20](#). Since $d_Y(\lambda^+, \lambda^-) > 0$ this is impossible.

Hence, the map $\mathbf{t}^+(\hat{\iota}_q|_{\partial' Y})$ (as defined in (1) in [Section 4.2](#)) is homotopic to $\hat{\iota}_q|_{\partial' Y}$ in $\bar{X}_Y \setminus \tilde{\mathcal{P}}_Y$ by pushing across the polygonal regions given by [Lemma 4.2](#) along leaves of λ^+ . This extends to a homotopy of $\hat{\iota}_q$ to a map $\hat{\iota}' : Y \rightarrow \bar{X}_Y$ which we claim is still an embedding. (Note that, in the case that X is fully-punctured, $\hat{\iota}' = \hat{\iota}_q$, since all singularities of fully-punctured surfaces are completion points.)

To prove that $\hat{\iota}'$ is an embedding, let C be a component of the preimage of $\hat{\iota}_q(Y')$ in \hat{X} (using the notation of [Section 2.3](#), C is a translate of $\text{CH}_q(\Lambda)$). If α is a geodesic segment in ∂C , the triangles used in the hull construction are attached to α and are contained in C . If such a triangle t intersects a triangle t' from a different segment α' , they overlap as in the right side of [Figure 20](#). Any two arcs l^+, l^- of λ^+ and λ^- passing through a point in the overlap must intersect both α and α' . These arcs are at distance 0 in $\mathcal{A}(Y)$, since they can be isotoped to each other rel

∂Y . Hence $d_Y(\lambda^-, \lambda^+) = 0$, contradicting the hypothesis. Therefore, t, t' cannot overlap.

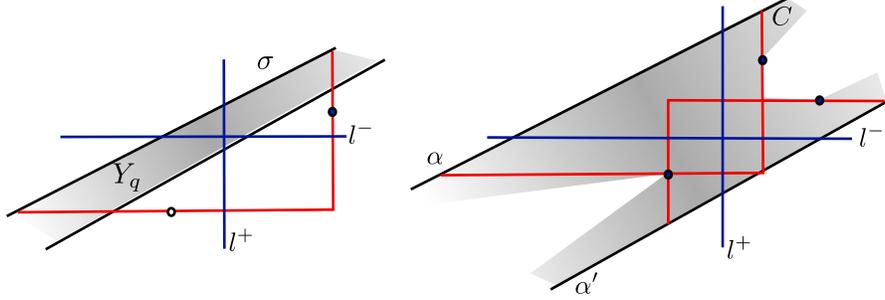


FIGURE 20. Left: If $t \in \mathcal{T}(\sigma)$ (in red) is not contained in Y_q then $d_Y(\lambda^+, \lambda^-) = 0$. Right: An overlap of two hull triangles. Any completion point in the boundary of a hull triangle does not correspond to a puncture in $\tilde{\mathcal{P}}_Y$.

We conclude that the polygonal regions of our homotopy are embedded and disjoint, and thus the homotopy can be chosen so that \hat{i}' is an embedding. Since the image of \hat{i}' is contained in the image of \hat{i}_q , we apply [Lemma 5.1](#) to get that the projection $\iota': Y \rightarrow X$ restricts to an embedding on $\text{int}(Y)$.

Now orient $\partial'Y$ in the opposite direction, pointing out of the surface, and apply \mathbf{t} again, this time to $\hat{i}'(\partial'Y)$. The triangles in the construction now extend outside the surface, and the result of the operation is the rectangle hull $\mathbf{r}(\mathbf{t}(\hat{i}_q(\partial'Y)))$, which is therefore composed of τ -edges. Using the homotopy pushing $\hat{i}'|_{\partial'Y}$ outward along leaves of λ^+ to $\mathbf{t}^+(\hat{i}'|_{\partial'Y})$ (again using [Lemma 4.2](#)) we obtain our final map \hat{i}_τ . See [Figure 21](#). It remains to show that $\hat{i}_\tau: Y \rightarrow X_Y$ has the required properties. To prove this, let us recapitulate the construction in the universal cover.

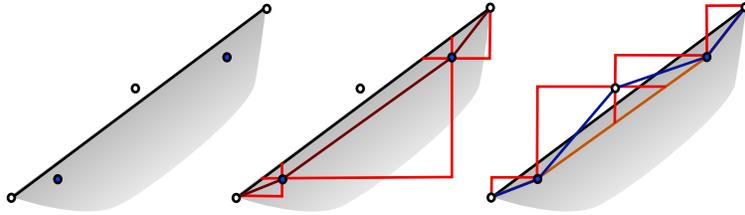


FIGURE 21. An inner \mathbf{t} followed by outer \mathbf{t} yields τ -edges. This locally depicts the homotopy from \hat{i}_q to \hat{i}_τ .

As before, let $C = \text{CH}_q(\Lambda)$. The map \hat{i}' lifts to a $\pi_1(Y)$ -equivariant homeomorphism $C \rightarrow C'$, where C' is obtained by giving each saddle connection κ in the

boundary of C the transverse orientation pointing into C and removing the polygons $P(\kappa)$ given in [Lemma 4.2](#). This map is equivariantly homotopic to the identity by pushing along leaves of the vertical foliation.

The outward step of our construction then pushes back along leaves of the vertical foliation to obtain a $\pi_1(Y)$ -equivariant map $C' \rightarrow C_\tau \subset \hat{X}$, so that the composition $C \rightarrow C' \rightarrow C_\tau$ is a lift of the map $\hat{\iota}_\tau: Y' \rightarrow \bar{X}_Y$. To show that $\hat{\iota}_\tau: Y' \rightarrow \bar{X}_Y$ is an embedding, it suffices to show that the composition $C \rightarrow C_\tau$ is a homeomorphism.

For every non-singular point $p \in \partial C$ there is an arc n_p in λ^+ such that the deformation of C to C_τ is supported on the union $\bigcup n_p$, and preserves each n_p . Thus to show that $C \rightarrow C_\tau$ is a homeomorphism it suffices to show that $n_p \cap n_{p'} = \emptyset$ for each $p \neq p'$ in ∂C . The interior pieces, $n_p \cap C$, are already disjoint for distinct points, by our construction. Thus if n_p intersects $n_{p'}$ their union is an interval J in a leaf of λ^+ with some subinterval between p and p' lying outside C . This contradicts the convexity of C .

To show that ι_τ is an embedding when restricted to $\text{int}(Y)$, it suffices to check that the interior of C_τ is disjoint from all its translations under the entire deck group $\pi_1(X)$. To see this, take $g \in \pi_1(X)$ so that C_τ and $g \cdot C_\tau$ are distinct and intersect. Since $\iota': \text{int}(Y) \rightarrow X$ is an embedding, C' and $g \cdot C'$ meet only along their boundary. Further, if σ is a saddle connection in $\partial C' \cap \partial(g \cdot C')$, then σ is the hypotenuse of a singularity-free triangle pointing into C' as well as one pointing into $g \cdot C'$. Hence, σ is a τ -edge and so is fixed under the map $C' \rightarrow C_\tau$.

Now if the interiors of C_τ and $g \cdot C_\tau$ intersect there must be saddle connections $\sigma, \kappa \subset \partial C'$ such that $P(\sigma)$ and $P(g \cdot \kappa)$ have intersecting interiors. (Here, σ, κ are oriented out of C' .) By the previous paragraph, σ and $g \cdot \kappa$ are distinct. As C' and $g \cdot C'$ meet only along their boundary, σ and $g \cdot \kappa$ have disjoint interiors and any arc l of λ^+ joining σ to $g \cdot \kappa$ within $P(\sigma) \cup P(g \cdot \kappa)$ lives outside of C' and $g \cdot C'$. In particular, the chosen transverse orientations on σ and $g \cdot \kappa$ point to the interior of l . However, by [Lemma 4.3](#), in this situation, the interiors of $P(\sigma)$ and $P(g \cdot \kappa)$ do not intersect. It follows that $\iota_\tau: \text{int}(Y) \rightarrow X$ is an embedding.

It only remains to prove property (1) of the definition of τ -compatible. Since $\hat{\iota}_\tau: Y' \rightarrow \bar{X}_Y$ is an embedding, it follows that $\hat{\iota}_\tau|_{\partial Y}$ is an embedding, and its image does not meet \tilde{P}_Y by the same argument used to prove item (2) of [Lemma 5.1](#). By construction the image $\hat{\iota}_\tau(\partial Y)$ is composed of τ -edges.

Now suppose that Y is an annulus. Then $\hat{\iota}_q(Y)$ is the (nondegenerate) maximal flat cylinder of \bar{X}_Y by [Proposition 5.2](#). Choosing the inward-pointing orientation for ∂Y , we claim that $\mathbf{t}^+(\hat{\iota}_q|_{\partial Y}) = \hat{\iota}_q|_{\partial Y}$: Otherwise, there must be a saddle connection σ on the boundary of the flat annulus $\hat{\iota}_q(Y)$, and a triangle t pointing into the annulus with hypotenuse on σ , which encounters a singularity or puncture x on the other side of the annulus. The picture is similar to the left side of [Figure 20](#). A variation on the Gauss–Bonnet argument in the annulus case of [Proposition 5.2](#) then produces vertical and horizontal leaves passing through x which have disjoint representatives, and hence $d_Y(\lambda^+, \lambda^-) \leq 1$. Thus the inward step of the process is the identity, and the outward step and the rest of the proof proceed just as in the nonannular case. \square

Remark 5.4. From the proof of [Theorem 5.3](#), we record the fact that if X is fully-punctured and Y satisfies the hypotheses of [Theorem 5.3](#), then $\mathbf{t}^+(\hat{\iota}_q|_{\partial Y}) = \hat{\iota}_q|_{\partial Y}$ and $\iota' = \hat{\iota}_q$. Hence, in this case we have that $\partial_\tau Y = \mathbf{r}(\partial_q Y)$.

6. EMBEDDED POCKETS OF THE VEERING TRIANGULATION AND BOUNDED PROJECTIONS

In this section, let X be fully-punctured with respect to the foliations λ^\pm of a pseudo-Anosov $f : X \rightarrow X$, and let M be the mapping torus. Recall that every fiber associated to the fibered face \mathcal{F} of X must also be fully-punctured because they are transverse to the same suspension flow, and hence that \mathcal{F} is a *fully-punctured fibered face*.

We now prove our two main theorems on the structure of subsurface projections in a fully-punctured fibered face, [Theorem 1.1](#) and [Theorem 1.2](#). The main tools in the proof are the structure and embedding theorems for pockets associated with high-distance subsurfaces, which we develop below. Recall that $\text{diam}_Z(\cdot)$ denotes the diameter of $\pi_Z(\cdot)$ in $\mathcal{A}(Z)$ and that subsurfaces Y and Z *overlap* if, up to isotopy, they are neither disjoint nor nested.

6.1. Projections and τ -compatible subsurfaces. We begin by discussing projection to τ -compatible subsurfaces.

Lemma 6.1. *Let Y and Z be τ -compatible subsurfaces of X and let $K \subset X$ be a disjoint collection of saddle connections which correspond to edges from τ . Then*

- (1) *If K meets $\text{int}_\tau(Y)$, then $\pi_Y(K) \neq \emptyset$, and $\text{diam}_Y(\pi_Y(K)) \leq 1$.*
- (2) *If Y and Z are disjoint, then so are $\text{int}_\tau(Y)$ and $\text{int}_\tau(Z)$.*
- (3) *If Y and Z overlap, then $\text{diam}_Z(\partial Y \cup \partial_\tau Y) \leq 1$.*
- (4) *The subsurface $\text{int}_\tau(Y)$ is in minimal position with the foliations λ^\pm . In particular, the arcs of $\text{int}_\tau(Y) \cap \lambda^\pm$ agree with the arcs of $\pi_Y(\lambda^\pm)$.*

Proof. For item (1), the main point is to show that an edge of K that meets $\text{int}_\tau(Y)$ lifts to an *essential edge* in \bar{X}_Y . This is true for edges meeting $\text{int}_q(Y)$, using the local CAT(0) geometry of \bar{X}_Y and the fact that $\hat{\iota}_q(Y')$ is a locally convex embedding.

Thus it will suffice to show that any τ -edge e meeting $\text{int}_\tau(Y)$ must also meet $\text{int}_q(Y)$. Suppose, on the contrary, that e meets $\text{int}_\tau(Y)$ but not $\text{int}_q(Y)$. Then e meets the interior of a polygon $P(\sigma)$ where σ is an outward-oriented saddle connection in $\partial_q Y$ (recall from [Remark 5.4](#) that, since X is fully-punctured, the inner \mathbf{t} step in the construction of $\hat{\iota}_\tau$ is the identity, and the outer \mathbf{t} is in fact a rectangle hull). Let R be the singularity-free rectangle spanned by e . If e is contained in $P(\sigma)$ then R can be extended to a rectangle whose diagonal lies in σ , and hence e is one of the edges of $\mathbf{r}(\sigma)$; but this contradicts the assumption that e meets $\text{int}_\tau(Y)$. Thus e crosses some edge f of $\mathbf{r}(\sigma)$. However, f is contained in a singularity-free triangle whose hypotenuse lies along σ and so σ must cross the rectangle R either top-to-bottom or side-to-side. In either case, we see that e crosses $\sigma \subset \partial_q Y$, a contradiction. We conclude that if a τ -edge meets $\text{int}_\tau(Y)$, then it also meets $\text{int}_q(Y)$ and hence has a well-defined projection to Y . The diameter bound in item (1) is then immediate since K is a disjoint collection of essential arcs of $\mathcal{A}(X)$.

For item (2), first note that when Y and Z are disjoint subsurfaces of X , the interiors $\text{int}_q(Y)$ and $\text{int}_q(Z)$ are also disjoint. This follows from [Corollary 2.4](#) and the q -hulls construction in [Lemma 2.6](#). More precisely, let Λ_Y and Λ_Z be the limit sets of Y and Z in $\partial\mathbb{H}^2$ (using our identifications from [Section 2.3](#)). Since Y and Z do not intersect, Λ_Y and Λ_Z do not link in $\partial\mathbb{H}^2$ and so $\text{CH}_q(\Lambda_Y)$ and $\text{CH}_q(\Lambda_Z)$ have disjoint interiors by [Corollary 2.4](#). This implies that $\text{int}_q(Y)$ and $\text{int}_q(Z)$ are disjoint in X .

To obtain $\text{int}_\tau(Y)$ from $\text{int}_q(Y)$ we append to each saddle connection σ in $\partial_q Y$ the (open) polygon $P(\sigma)$, where σ is oriented out of Y . We obtain $\text{int}_\tau(Z)$ from $\text{int}_q(Z)$ by the same construction. Since $\text{int}_q(Y)$ and $\text{int}_q(Z)$ are disjoint in X , it suffices to show that $P(\sigma)$ and $P(\kappa)$ have disjoint interiors, where $\sigma \subset \partial_q Y$ and $\kappa \subset \partial_q Z$. If $\sigma = \kappa$, then this saddle connection spans a singularity-free rectangle and $P(\sigma) = \sigma = \kappa = P(\kappa)$. Otherwise, σ and κ have disjoint interiors and [Lemma 4.3](#) implies that $P(\sigma)$ and $P(\kappa)$ have disjoint interiors, as required. This proves item (2).

Since $\text{int}_\tau(Y)$ is an embedded representative of the interior of Y , ∂Y has a representative disjoint from the collection of saddle connections in $\partial_\tau Y$. Hence $\text{diam}_Z(\partial Y \cup \partial_\tau Y) \leq 1$, proving item (3). For item (4), first note that the subsurface $\text{int}_q(Y)$ is in minimal position with the foliations λ^\pm . This is immediate from the local CAT(0) geometry in \bar{X}_Y and the fact that λ^\pm are geodesic: any bigon in \bar{X}_Y between $\hat{i}_q(\partial' Y)$ and a leaf of λ^\pm would lift to a bigon in \hat{X} bounded by two geodesic segments, a contradiction to uniqueness of geodesics in \hat{X} . The statement for $\text{int}_\tau(Y)$ then follows from the fact that the homotopy from $\partial_q Y$ to $\partial_\tau Y$ can be taken to move either along vertical or along horizontal leaves, using either \mathbf{t}^+ or \mathbf{t}^- as in the proof of [Theorem 5.3](#). \square

6.2. Pockets for a τ -compatible subsurface. Suppose that $Y \subset X$ is τ -compatible. By [Corollary 3.6](#), the set $T(\partial_\tau Y)$ of sections containing $\partial_\tau Y$ contains a top and a bottom section, denoted $T^+ = T^+(\partial_\tau Y)$ and $T^- = T^-(\partial_\tau Y)$, which between them bound a number of pockets. See [Section 3](#) for terminology related to sections and pockets. Our assumption on $d_Y(\lambda^-, \lambda^+)$ will imply that one of these pockets is isotopic to a thickening of Y , as explained in the following proposition:

Proposition 6.2 (Pockets in τ). *Let (X, q) be fully-punctured and $Y \subset X$ an essential nonannular subsurface.*

- (1) *If $d_Y(\lambda^-, \lambda^+) > 0$ then $d_Y(T^+, \lambda^+) = d_Y(T^-, \lambda^-) = 0$.*
- (2) *If $d_Y(\lambda^-, \lambda^+) > 2$ then T^+ and T^- bound a pocket U_Y whose interior is isotopic to a thickening of $\text{int}(Y)$.*

When Y is an annulus,

- (1) *If $d_Y(\lambda^-, \lambda^+) > 1$ then $d_Y(T^+, \lambda^+) = d_Y(T^-, \lambda^-) = 1$.*
- (2) *If $d_Y(\lambda^-, \lambda^+) > 4$ then T^+ and T^- bound a pocket U_Y whose interior is isotopic to a thickening of $\text{int}(Y)$.*

Proof. Begin with the following lemma:

Lemma 6.3. *Suppose that $Y \subset X$ is τ -compatible, let e be an edge of $\partial_\tau Y$ and let f be a τ -edge crossing e with $f > e$. Then $d_Y(f, \lambda^+) \leq 1$ if Y is an annulus and $d_Y(f, \lambda^+) = 0$ otherwise. Similarly if $f < e$ then the same statement holds for $d_Y(f, \lambda^-)$.*

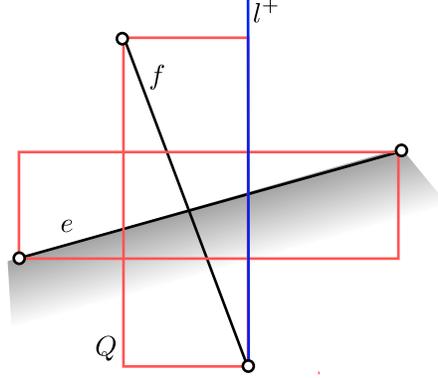


FIGURE 22. Local picture near the τ -edge e of $\partial_\tau Y$ with $\text{int}_\tau(Y) \subset X$ shaded. When $f > e$, the edge l^+ of Q represents $\pi_Y(\lambda^+)$ and is disjoint from f . Note that Q is immersed in X .

The key idea of the proof is pictured in [Figure 22](#). Here it is shown that if f crosses $e \subset \partial_\tau Y$ with $f > e$, then some component of the intersection of f with $\text{int}_\tau(Y)$ is disjoint from some arc in $\pi_Y(\lambda^+)$. However, the spanning rectangle Q for f is immersed in X (rather than necessarily embedded). To handle this issue, we work in the cover \tilde{X} .

Proof. Let C° be a component of the preimage of $\text{int}_\tau(Y)$ under $\tilde{X} \rightarrow X$ and choose a saddle connection \tilde{e} in the boundary of C° which projects to e . Further, let \tilde{f} be any lift of f which crosses \tilde{e} . Since f is a τ -edge, \tilde{f} spans a singularity-free rectangle \tilde{Q} whose immersed image in X we denote by Q .

Every τ -edge which crosses \tilde{Q} does so either top to bottom or side to side. Since $f > e$, \tilde{e} must cross \tilde{Q} from side to side (see [Section 3](#)). Since all τ -edges in ∂C° are disjoint, they all must cross \tilde{Q} from side to side.

Since $\text{int}_\tau(Y)$ is in minimal position with λ^+ ([Lemma 6.1](#)), C° intersects each leaf of the vertical foliation in a connected set. Together these observations imply that $\tilde{Q} \cap C^\circ$ is a single polygon \tilde{B} , bounded by at least one edge crossing \tilde{Q} from side to side (which we have called \tilde{e}). See [Figure 23](#).

Claim. \tilde{B} embeds in $\text{int}_\tau(Y)$ under the covering $\tilde{X} \rightarrow X$.

Proof of claim. Since $\tilde{B} \subset C^\circ$, the image of \tilde{B} is contained in $\text{int}_\tau(Y)$. Suppose that $x, y \in \tilde{B}$ map to the same point in $\text{int}_\tau(Y)$, and denote by l_x and l_y the vertical leaf segments in \tilde{X} starting at x and y , respectively, and continuing to \tilde{e} . Since \tilde{B} is

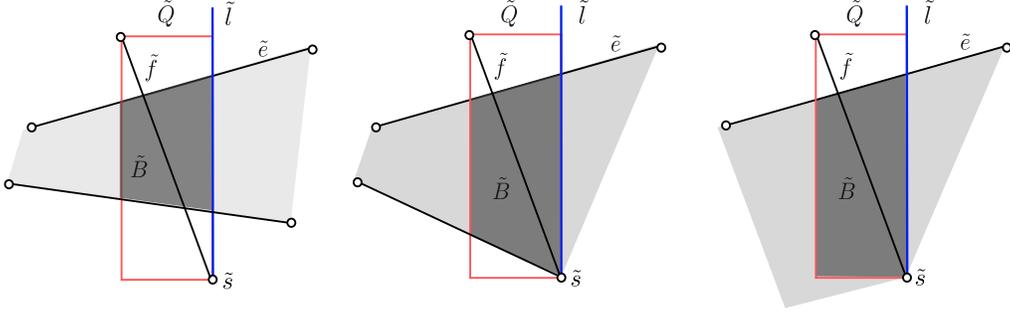


FIGURE 23. The 3 possibilities for \tilde{B} . The lightly shaded region is part of C^o in \tilde{X} .

convex, $l_x, l_y \subset \tilde{B}$. Suppose that l_x is no longer than l_y and let l'_y be the subsegment of l_y with length equal to that of l_x . Then l_x and l'_y are identified under the map $\tilde{X} \rightarrow X$. But the identification of $\partial l_x \setminus \{x\} \subset \tilde{e}$ and $\partial l'_y \setminus \{y\} \subset \tilde{B} \cup \tilde{e} \subset C^o \cup \tilde{e}$ gives a contradiction, unless $x = y$: the edge \tilde{e} is mapped injectively into X with image $e \subset \partial_\tau Y$ disjoint from the image of C^o , which is $\text{int}_\tau(Y)$. \square

Let \tilde{s} be the vertex of \tilde{f} which is on the same side of \tilde{e} as \tilde{B} . Let \tilde{l} be the vertical side of \tilde{Q} starting at \tilde{s} .

Let B be the image of \tilde{B} in X . By the claim, B is a singularity-free quadrilateral in X whose interior is contained in $\text{int}_\tau(Y)$. The images in X of $\tilde{f} \cap \tilde{B}$ and $\tilde{l} \cap \tilde{B}$ are therefore disjoint proper arcs in $\text{int}_\tau(Y)$, which by [Lemma 6.1](#) are representatives of $\pi_Y(f)$ and $\pi_Y(\lambda^+)$, respectively. Moreover, these arcs are properly homotopic in $\text{int}_\tau(Y)$ by a homotopy supported in B .

Hence, when Y is nonannular, we conclude that $d_Y(f, \lambda^+) = 0$. If Y is an annulus, we project the picture to the annular cover X_Y , where we note that the image l of \tilde{l} , continued to infinity, cannot intersect f without meeting Q , and hence e , again. Since l can only meet $\partial_q Y$ once in the annular cover, we conclude it is disjoint from f and so $d_Y(f, \lambda^+) = 1$.

The case $f < e$ is similar, so [Lemma 6.3](#) is proved. \square

We return to the proof of [Proposition 6.2](#). Let Y be nonannular. Note that by definition the only upward-flippable edges in T^+ must lie in $\partial_\tau Y$. Let e be such an edge and consider the single flip move that replaces e with an edge f . Then $f > e$, so by [Lemma 6.3](#), $d_Y(f, \lambda^+) = 0$. On the other hand f and e are diagonals of a quadrilateral made of edges of T^+ , at least one of which, e' , gives the same element of $\mathcal{A}(Y)$ as f . Hence $d_Y(T^+, \lambda^+) = 0$.

If Y is an annulus, we note that e' and the vertical leaf in the proof of [Lemma 6.3](#) give adjacent vertices of $\mathcal{A}(Y)$, so $d_Y(T^+, \lambda^+) \leq 1$. Note that $d_Y(T^+, \lambda^+) \neq 0$ because no leaf of the foliation λ^+ has both its endpoints terminating at completion points.

To prove the statements about pockets, let K be the common edges of T^+ and T^- , viewed as a subcomplex of X . If $\text{int}_\tau(Y)$ contains an edge of K then from the triangle inequality, together with the first part of the proposition, we obtain $d_Y(\lambda^+, \lambda^-) \leq 2$ when Y is nonannular, and $d_Y(\lambda^+, \lambda^-) \leq 4$ when Y is an annulus. By our hypotheses this does not happen, so we conclude that $T^+, T^- \in T(\partial_\tau Y)$ have no common edges contained in $\text{int}_\tau(Y)$. Hence T^+ and T^- bound a pocket U_Y whose base is $\text{int}_\tau(Y)$. This completes the proof. \square

6.3. Isolated pockets and projection bounds. Let X be a fiber in $\mathbb{R}_+ \mathcal{F}$, and let Y be a τ -compatible subsurface of X such that $d_Y(\lambda^-, \lambda^+) > 4$. An *isolated pocket* for Y in $(X \times \mathbb{R}, \tau)$ is a subpocket $V = V_Y$ of U_Y with base $\text{int}_\tau(Y)$ such that

- (1) For each edge e of V which is not contained in $\partial_\tau Y$,

$$d_Y(e, \lambda^+) \geq 3 \quad \text{and} \quad d_Y(e, \lambda^-) \geq 3$$

if Y is nonannular, and

$$d_Y(e, \lambda^+) \geq 4 \quad \text{and} \quad d_Y(e, \lambda^-) \geq 4$$

if Y is an annulus.

- (2) Denoting by V^\pm the top and bottom of V with their induced triangulations,

$$d_Y(V^-, V^+) \geq 1.$$

Note that condition (2) guarantees that $\text{int}(V_Y) \cong \text{int}_\tau(Y) \times (0, 1)$ is still a pocket just as in [Proposition 6.2](#). The next lemma shows that for Y with $d_Y(\lambda^-, \lambda^+)$ sufficiently large, Y has an isolated pocket with $d_Y(V^-, V^+)$ roughly $d_Y(\lambda^-, \lambda^+)$.

Lemma 6.4. *Suppose that Y is a nonannular subsurface of X with $d_Y(\lambda^-, \lambda^+) > 8$. Then Y has an isolated pocket V with $d_Y(V^-, V^+) \geq d_Y(\lambda^-, \lambda^+) - 8$.*

If Y is an annulus with $d_Y(\lambda^-, \lambda^+) > 10$, then Y has an isolated pocket V with $d_Y(V^-, V^+) \geq d_Y(\lambda^-, \lambda^+) - 10$.

Proof. Let $c = 4$ if Y is an annulus and $c = 3$ otherwise, and assume that $d_Y(\lambda^+, \lambda^-) > 2c + 2$. Since the pocket $U = U_Y$ is connected ([Proposition 3.3](#)), there is a sequence of sections $T^- = T_0, T_1, \dots, T_N = T^+$ in $T(\partial_\tau Y)$ such that T_{i+1} differs from T_i by an upward diagonal exchange. From [Proposition 6.2](#), we know that $d_Y(T^-, \lambda^-) \leq 1$ and $d_Y(T^+, \lambda^+) \leq 1$. Let $0 < a < N$ be largest integer such that $d_Y(T_{a-1}, \lambda^-) < c$; hence $d_Y(T_i, \lambda^-) \geq c$ for all $i \geq a$. Now let $b < N$ be the smallest integer greater than a such that $d_Y(T_{b+1}, \lambda^+) < c$; then $d_Y(T_i, \lambda^+) \geq c$ for all $a \leq i \leq b$.

Note that these indices exist since $d_Y(\lambda^-, \lambda^+) \geq 2c + 1$.

Now let V be the pocket between T_a and T_b with base contained in $\text{int}_\tau(Y)$ and note that V is a subpocket of U . Any edge e of V not contained in $\partial_\tau Y$ is contained in a section $T_i \in T(\partial_\tau Y)$ for $a \leq i \leq b$. Since we have $d_Y(T_i, \lambda^\pm) \geq c$, we have $d_Y(e, \lambda^\pm) \geq c$. Thus it only remains to get a lower bound on $d_Y(V^+, V^-)$.

The triangle inequality (and diameter bound on T_a and T_b) gives

$$d_Y(V^-, V^+) = d_Y(T_a, T_b) \geq d_Y(\lambda^-, \lambda^+) - 2c - 2 \geq 1.$$

This implies that $\text{int}_\tau(Y)$ is the base of V and completes the proof. \square

The following proposition shows that isolated pockets coming from either disjoint or overlapping subsurfaces of X have interiors which do not meet.

Proposition 6.5 (Disjoint pockets). *Suppose that Y and Z are subsurfaces of X with isolated pockets V_Y and V_Z . Then, up to switching Y and Z , either Y is nested in Z , or the isolated pockets V_Y and V_Z have disjoint interiors in $X \times \mathbb{R}$.*

Proof. If the subsurfaces Y and Z are disjoint, then $\text{int}_\tau(Y)$ and $\text{int}_\tau(Z)$ are also disjoint by [Lemma 6.1](#). Hence, the maximal pockets U_Y and U_Z have disjoint interiors by definition.

Now suppose that Y is not an annulus. We claim that if Y and Z overlap then either

$$d_Y(\partial_\tau Z, \lambda^+) \leq 1 \quad \text{or} \quad d_Y(\partial_\tau Z, \lambda^-) \leq 1.$$

To see this, first note that there is some edge f contained in $\text{int}_\tau(Z)$ such that f crosses some edges of $\partial_\tau Y$. Otherwise, every triangulation of $\text{int}_\tau(Z)$ by τ -edges would contain edges from $\partial_\tau Y$. But then applying this to $T^\pm(\partial_\tau Z)$ and using [Proposition 6.2](#), we would have that

$$d_Z(\lambda^-, \lambda^+) \leq 2 + \text{diam}_Z(\partial_\tau Y) \leq 3,$$

contradicting our assumption on the subsurface Z . Now if f intersects an edge e of $\partial_\tau Y$ and $f > e$, then by [Lemma 6.3](#), $d_Y(\partial_\tau Z, \lambda^+) \leq d_Y(\partial_\tau Z, f) \leq 1$. If $f < e$ then [Lemma 6.3](#) gives $d_Y(\partial_\tau Z, \lambda^-) \leq d_Y(\partial_\tau Z, f) \leq 1$.

Now suppose that e is an edge of $U_Y \cap U_Z$ which is not contained in $\partial_\tau Y \cup \partial_\tau Z$. Then e , as a τ -edge in X , is disjoint from $\partial_\tau Z$ and so $d_Y(e, \lambda^+) \leq 2$ or $d_Y(e, \lambda^-) \leq 2$. Hence e cannot be contained in V_Y . We conclude that $V_Y \cap V_Z \subset \partial_\tau Y \cup \partial_\tau Z$. This completes the proof when Y is not an annulus.

When Y is an annulus, then a similar argument using the annular case of [Lemma 6.3](#) shows that if Y and Z overlap then either

$$d_Y(\partial_\tau Z, \lambda^+) \leq 2 \quad \text{or} \quad d_Y(\partial_\tau Z, \lambda^-) \leq 2.$$

Hence, if e is an edge of $U_Y \cap U_Z$ which is not contained in $\partial_\tau Y \cup \partial_\tau Z$, then $d_Y(e, \lambda^\pm) \leq 3$. So again e cannot be contained in V_Y and we conclude that $V_Y \cap V_Z \subset \partial_\tau Y \cup \partial_\tau Z$ as required. \square

We next prove that isolated pockets embed into the fibered manifold M . This is [Theorem 1.3](#), which we restate here in more precise language.

Theorem 1.3 (Embedding the pocket). *Suppose Y is a subsurface of a fully-punctured fiber X with $d_Y(\lambda^-, \lambda^+) > \beta$, where $\beta = 8$ if Y is nonannular and $\beta = 10$ if Y is an annulus. Then Y has an isolated pocket V_Y in $X \times \mathbb{R}$, and the covering map $X \times \mathbb{R} \rightarrow M$ restricts to an embedding of the subcomplex $V_Y \rightarrow M$.*

Proof. Let Φ be the simplicial isomorphism of $X \times \mathbb{R}$ induced by f as in [Theorem 2.1](#). Note that if T is a section of τ , then $\Phi(T)$ is the section of τ whose corresponding triangulation of X is $f(T)$. Hence, $\Phi(T(\partial_\tau Y)) = T(\partial_\tau f(Y))$.

By [Lemma 6.4](#), Y has an isolated pocket $V = V_Y$. Note that V embeds into M if and only if it is disjoint from its translates $V_i = \Phi^i(V)$ for each $i \neq 0$. By the remark

above, each V_i is itself an isolated pocket for the subsurface $Y_i = f^i(Y)$, and any two of these subsurfaces are either disjoint or overlap in X . Hence, by [Proposition 6.5](#) the isolated pockets V_i are disjoint as required. \square

We will now prove [Theorem 1.1](#), whose statement we recall here:

Theorem 1.1 *Let M be a hyperbolic 3-manifold with fully-punctured fibered face \mathcal{F} and veering triangulation τ . For any subsurface W of any fiber of \mathcal{F} ,*

$$\alpha \cdot (d_W(\lambda^-, \lambda^+) - \beta) < |\tau|,$$

where $|\tau|$ is the number of tetrahedra in τ , $\alpha = 1$ and $\beta = 10$ when W is an annulus and $\alpha = 3|\chi(W)|$ and $\beta = 8$ when W is not an annulus.

Proof. Suppose that W is any nonannular subsurface of any fiber F in $\mathbb{R}_+\mathcal{F}$. We may assume that $d_W(\lambda^-, \lambda^+) > 8$. Then [Lemma 6.4](#) implies that W has an isolated pocket V_W in $(F \times \mathbb{R}, \tau)$ such that $d_W(V_W^-, V_W^+) \geq d_Y(\lambda^-, \lambda^+) - 8$. By [Theorem 1.3](#), the isolated pocket $V_W \subset (F \times \mathbb{R}, \tau)$ embeds into (M, τ) . Hence $|V_W| \leq |\tau|$, where $|V_W|$ denotes the number of tetrahedra of V_W . Now each tetrahedron of V_W corresponds to a diagonal exchange between the triangulations V_W^- and V_W^+ of W_τ and each diagonal exchange replaces a single edge of the triangulation. There are at least $3|\chi(W)| + 1$ non-boundary edges to each triangulation of W , and the diameter in $\mathcal{A}(W)$ of an ideal triangulation is 1, so we conclude

$$\begin{aligned} (2) \quad |\tau| &\geq |V_W| = \#\{\text{diagonal exchanges from } V_W^- \text{ to } V_W^+\} \\ &> 3|\chi(W)| \cdot d_W(V^-, V^+) \\ &\geq 3|\chi(W)| \cdot (d_W(\lambda^-, \lambda^+) - 8). \end{aligned}$$

This completes the proof when W is nonannular.

When W is an annulus, we use the annular case of [Lemma 6.4](#) to obtain an isolated pocket V_W in $(F \times \mathbb{R}, \tau)$ such that $d_W(V_W^-, V_W^+) \geq d_Y(\lambda^-, \lambda^+) - 10$. Noting that a triangulation of the annulus contains at least 2 (non-boundary) edges, the same argument implies that

$$\begin{aligned} |\tau| &\geq |V_W| = \#\{\text{diagonal exchanges from } V_W^- \text{ to } V_W^+\} \\ &> d_W(V^-, V^+) \\ &\geq d_W(\lambda^-, \lambda^+) - 10, \end{aligned}$$

as required. \square

6.4. Sweeping through embedded pockets. We are now ready to prove [Theorem 1.2](#), whose statement we reproduce below. This theorem relates subsurfaces of large projections among different fibers of a fixed face.

Theorem 1.2 *Let M be a hyperbolic 3-manifold with fully-punctured fibered face \mathcal{F} and suppose that S and F are each fibers in $\mathbb{R}_+\mathcal{F}$. If W is a subsurface of F , then either W is isotopic along the flow to a subsurface of S , or*

$$3|\chi(S)| \geq d_W(\lambda^-, \lambda^+) - \beta,$$

where $\beta = 10$ if W is an annulus and $\beta = 8$ otherwise.

Recall from [Lemma 2.8](#) that we can identify $d_W(\lambda^+, \lambda^-)$ with $d_W(\Lambda^+, \Lambda^-)$, agreeing with the statement given in the introduction.

We will require the following lemma, which essentially states that immersed sub-surfaces with large projection are necessarily covers of subsurfaces. Recall that in [Section 2.2](#) we defined the distance $d_W(\lambda^+, \lambda^-)$ when W is a compact core of a cover $X_\Gamma \rightarrow X$ corresponding to a finitely generated subgroup $\Gamma \leq \pi_1(X)$.

Lemma 6.6 (Immersion to cover). *Suppose that (X, q) is a fully-punctured surface. Let Γ be a finitely generated subgroup of $\pi_1(X)$ and let W be a compact core of the cover $X_\Gamma \rightarrow X$. If W is nonannular and $d_W(\lambda^-, \lambda^+) > 4$ or if W is an annulus and $d_W(\lambda^-, \lambda^+) > 6$, then there is a subsurface Y of X such that $W \rightarrow X$ is homotopic to a finite cover $W \rightarrow Y \subset X$.*

In particular, Γ is a finite index subgroup of $\pi_1(Y)$.

Proof. Suppose that $d_W(\lambda^-, \lambda^+) > 4$ if W is nonannular and $d_W(\lambda^-, \lambda^+) > 6$ if W is an annulus. Let $p: \check{X} \rightarrow X$ be a finite cover to which $W \rightarrow X$ lifts to an embedding $W \rightarrow \check{X}$ (this exists since surface groups are LERF [[Sco78](#)]), and identify W with its image in \check{X} . Lift q along with the veering triangulation to $(\check{X} \times \mathbb{R}, \tau)$. By [Theorem 5.3](#), W is a τ -compatible subsurface of \check{X} , and by [Theorem 5.3](#) and [Proposition 3.3](#), $T_{\check{X}}(\partial_\tau W)$ is nonempty and connected. To prove the lemma, we show that $\text{int}_\tau(W) \rightarrow X$ covers a subsurface of X . For this, it suffices to prove that each edge of $p^{-1}(p(\partial_\tau W))$ is disjoint from $\text{int}_\tau(W)$. Indeed, since W is τ -compatible, one component of $\check{X} \setminus \partial_\tau W$ is $\text{int}_\tau(W)$. If $p^{-1}(p(\partial_\tau W))$ is disjoint from $\text{int}_\tau(W)$, then $\text{int}_\tau(W)$ is also a component of $\check{X} \setminus p^{-1}(p(\partial_\tau W))$. As components of $\check{X} \setminus p^{-1}(p(\partial_\tau W))$ cover components of $X \setminus p(\partial_\tau W)$, this will show that $\text{int}_\tau(W) \rightarrow X$ covers a subsurface of X .

Hence, we must show that each edge of $p^{-1}(p(\partial_\tau W))$ is disjoint from $\text{int}_\tau(W)$. This is equivalent to the statement that no edge of $p^{-1}(p(\partial_\tau W))$ crosses $\partial_\tau W$ nor is contained in $\text{int}_\tau(W)$.

First suppose that W is not an annulus. If \check{T} is a section of $(\check{X} \times \mathbb{R}, \tau)$ with an edge f such that $f > e$ for an edge e of $\partial_\tau W$, then [Lemma 6.3](#) implies that $d_W(\check{T}, \lambda^+) = 0$. Similarly if $f < e$ then $d_W(\check{T}, \lambda^-) = 0$. Hence, if T is any section of $(X \times \mathbb{R}, \tau)$ such that $d_W(T, \lambda^\pm) \geq 1$, then its lift $\check{T} = p^{-1}(T)$ to \check{X} must contain the edges of $\partial_\tau W$ and so $\check{T} \in T_{\check{X}}(\partial_\tau W)$. Moreover, such a section T of $(X \times \mathbb{R}, \tau)$ with $d_W(T, \lambda^\pm) \geq 1$ must exist. This is because by [Lemma 3.1](#), we may sweep through $X \times \mathbb{R}$ with sections going from near λ^- to near λ^+ . If all sections were to have d_W -distance 0 from either λ^- or λ^+ , then there would be a pair T, T' differing by a single diagonal exchange such that $d_W(T, \lambda^-) = d_W(T', \lambda^+) = 0$. But this would imply that $d_W(\lambda^-, \lambda^+) \leq 2$, contradicting our assumption on distance.

Putting these facts together, we conclude that there exists a section T of $(X \times \mathbb{R}, \tau)$ with $d_W(T, \lambda^\pm) \geq 1$, and that for each such section

$$p^{-1}(T) \in T_{\check{X}}(p^{-1}(p(\partial_\tau W))).$$

Note that this in particular implies that no edge of $p^{-1}(p(\partial_\tau W))$ crosses an edge of $\partial_\tau W$.

We claim now that no edge e in $p^{-1}(p(\partial_\tau W))$ can be contained in $\text{int}_\tau(W)$. Such an edge would have a well-defined projection to $\mathcal{A}(W)$ and would necessarily appear in each section of $T_{\check{X}}(p^{-1}(p(\partial_\tau W)))$ (by definition of $T_{\check{X}}(\cdot)$). Using our conclusion from above, this would imply that $d_W(p^{-1}(T), e) = 0$ whenever $d_W(T, \lambda^\pm) \geq 1$. But just as before, by sweeping through $X \times \mathbb{R}$ with sections going from near λ^- to near λ^+ , we produce sections T_1, T_2 with $d_W(T_1, \lambda^-) = d_W(T_2, \lambda^+) = 1$. Since each of these sections' preimage in \check{X} contains the edge e , we get that $d_W(\lambda^\pm, e) \leq 2$, which contradicts our hypothesis that $d_W(\lambda^+, \lambda^-) > 4$.

This shows that no edge of $p^{-1}(p(\partial_\tau W))$ can meet $\text{int}_\tau(W)$ and completes the proof when W is nonannular. When W is an annulus, one proceeds exactly as above using the annular version of [Lemma 6.3](#). \square

Proof of [Theorem 1.2](#). We may assume that W is a subsurface of F such that $d_W(\lambda^-, \lambda^+) > \beta$.

First suppose that $\pi_1(W)$ is contained in $\pi_1(S)$. Then by [Lemma 6.6](#), there is a subsurface Y of S such that, up to conjugation in $\pi_1(S)$, $\pi_1(W) \leq \pi_1(Y)$ is a finite index subgroup; let $n \geq 1$ denote this index. If $\eta_F: \pi_1(M) \rightarrow \mathbb{Z}$ is the homomorphism representing the cohomology class dual to F , then $\eta_F|_{\pi_1(Y)}$ vanishes on the index n subgroup $\pi_1(W)$. Since \mathbb{Z} is torsion-free we must have that η vanishes on $\pi_1(Y)$ and hence $\pi_1(Y)$ is contained in $\pi_1(F)$. However, since the fundamental group of an embedded subsurface, in this case $W \subset F$, can not be nontrivially finite-index inside another subgroup of $\pi_1(F)$, we see that $n = 1$ and $\pi_1(W) = \pi_1(Y)$. That W is isotopic along the flow in M to $Y \subset S$ can be seen by lifting W and Y to the cover $S \times \mathbb{R} \rightarrow M$.

Hence, we may suppose by [Lemma 2.9](#) that the image of any $S \rightarrow M$ homotopic to the fiber S intersects any isotope of $W \subset F$ essentially. Since $d_W(\lambda^-, \lambda^+) > \beta$, W has a nonempty isolated pocket $V_W \subset F \times \mathbb{R}$ which simplicially embeds into (M, τ) by [Theorem 1.3](#). Let $\{W_i\}$ denote a sequence of sections of V_W from V_W^- to V_W^+ with W_{i+1} differing from W_i by an upward diagonal flip. Also, fix a simplicial map $f: S \rightarrow (M, \tau)$ which is obtained by composing a section of $(S \times \mathbb{R}, \tau)$ with the covering map $S \times \mathbb{R} \rightarrow M$.

Note that for each i , $f(S)$ meets at least one edge of the interior of W_i . Otherwise, the image of S in M misses the interior of W_i contradicting our assumption. In fact, even more is true: Call a component c of $f(S) \cap W_i$ *removable* if the triangles of $f(S)$ incident to the edges of c lie locally to one side of W_i in M . If c is removable, then there is an isotopy of W_i supported in a neighborhood of c which removes c from the intersection $f(S) \cap W_i$. Hence, if we denote by E_i the edges of $f(S) \cap W_i$ which do not lie in removable components, then E_i must be nonempty for each i .

We claim that for each i , E_i shares an edge with E_{i+1} . Otherwise, both E_i and E_{i+1} consist of a single edge and the tetrahedron corresponding to the diagonal exchange from W_i to W_{i+1} has E_i as its bottom edge and E_{i+1} as its top edge. But then both of these edges must be removable since pushing the bottom two faces of the tetrahedron slightly upward makes that intersection disappear, and similarly for

the top. This contradicts our above observation and establishes that E_i and E_{i+1} have a common edge.

We obtain a sequence in $\mathcal{A}(W)$,

$$V_W^- \supset E_0, E_1, \dots, E_n \subset V_W^+,$$

having the property that for each edge e_i of E_i there is an edge e_{i+1} of E_{i+1} such that e_i and e_{i+1} are disjoint. We conclude that the number of distinct edges in the sequence E_0, E_1, \dots, E_n is at least $d_W(V_W^-, V_W^+)$. Combining this with the fact that the number of edges in an ideal triangulation of S is $3|\chi(S)|$ and [Lemma 6.4](#), we see that

$$3|\chi(S)| \geq d_W(V_W^-, V_W^+) \geq d_W(\lambda^-, \lambda^+) - \beta,$$

as required. \square

We conclude the paper by recording the following corollary of [Lemma 6.6](#) and the proof of [Theorem 1.2](#).

Corollary 6.7. *Let M be a hyperbolic manifold with fully-punctured fibered face \mathcal{F} . Let W be a subsurface of a fiber $F \in \mathbb{R}_+\mathcal{F}$ such that $d_W(\Lambda^+, \Lambda^-) > 4$ if W is nonannular and $d_W(\Lambda^+, \Lambda^-) > 6$ if W is an annulus. If S is any fiber in $\mathbb{R}_+\mathcal{F}$ such that $\pi_1(W) < \pi_1(S)$, then W is isotopic to a subsurface of S .*

REFERENCES

- [Ago11] Ian Agol, *Ideal triangulations of pseudo-Anosov mapping tori*, *Topology and geometry in dimension three* **560** (2011), 1–17.
- [Ago12] ———, *Comparing layered triangulations of 3-manifolds which fiber over the circle*, *MathOverflow discussion*, <http://mathoverflow.net/questions/106426>, 2012.
- [BCM12] J. Brock, R. Canary, and Y. Minsky, *The classification of Kleinian surface groups, II: The ending lamination conjecture*, *Ann. of Math.* **176** (2012), no. 1, 1–149. MR 2925381
- [BKMM12] Jason Behrstock, Bruce Kleiner, Yair Minsky, and Lee Mosher, *Geometry and rigidity of mapping class groups*, *Geom. Topol.* **16** (2012), 781–888.
- [BS05] David Bachman and Saul Schleimer, *Surface bundles versus Heegaard splittings*, *Communications in Analysis and Geometry* (2005), no. 5, 903 – 928.
- [CC00] Alberto Candel and Lawrence Conlon, *Foliations II*, American Mathematical Society Providence, 2000.
- [FG13] David Futer and François Guéritaud, *Explicit angle structures for veering triangulations*, *Algebraic & Geometric Topology* **13** (2013), no. 1, 205–235.
- [FLM11] B. Farb, C. J. Leininger, and D. Margalit, *Small dilatation pseudo-Anosov homeomorphisms and 3-manifolds*, *Adv. Math.* **228** (2011), no. 3, 1466–1502. MR 2824561
- [Fri82] David Fried, *The geometry of cross sections to flows*, *Topology* **21** (1982), no. 4, 353–371.
- [Gué15] François Guéritaud, *Veering triangulations and the Cannon-Thurston map*, *ArXiv:1506.03387*, 2015.
- [Har02] K. Hartshorn, *Heegaard splittings of Haken manifolds have bounded distance*, *Pacific J. Math.* **204** (2002), no. 1, 61–75. MR MR1905192 (2003a:57037)
- [HIS16] Craig D Hodgson, Ahmad Issa, and Henry Segerman, *Non-geometric veering triangulations*, *Experimental Mathematics* **25** (2016), no. 1, 17–45.
- [HRST11] Craig D Hodgson, J Hyam Rubinstein, Henry Segerman, and Stephan Tillmann, *Veering triangulations admit strict angle structures*, *Geometry & Topology* **15** (2011), no. 4, 2073–2089.

- [JMM10] J. Johnson, Y. Minsky, and Y. Moriah, *Heegaard splittings with large subsurface distances*, *Algebr. Geom. Topol.* **10** (2010), no. 4, 2251–2275. MR 2745671 (2012g:57038)
- [McM00] Curtis T McMullen, *Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations*, *Annales scientifiques de l’Ecole normale supérieure* **33** (2000), no. 4, 519–560.
- [Min10] Yair Minsky, *The classification of Kleinian surface groups, I: Models and bounds*, *Ann. of Math.* (2010), 1–107.
- [MM00] Howard A. Masur and Yair N. Minsky, *Geometry of the complex of curves. II. Hierarchical structure*, *Geom. Funct. Anal.* **10** (2000), no. 4, 902–974.
- [MS13] Howard Masur and Saul Schleimer, *The geometry of the disk complex*, *Journal of the American Mathematical Society* **26** (2013), no. 1, 1–62.
- [Raf05] Kasra Rafi, *A characterization of short curves of a Teichmüller geodesic*, *Geometry & Topology* **9** (2005), no. 1, 179–202.
- [Sco78] Peter Scott, *Subgroups of surface groups are almost geometric*, *J. London Math. Soc.* (2) **17** (1978), no. 3, 555–565. MR 0494062
- [ST06] Martin Scharlemann and Maggy Tomova, *Alternate Heegaard genus bounds distance*, *Geom. Topol.* **10** (2006), 593–617 (electronic). MR 2224466
- [Thu86] William P Thurston, *A norm for the homology of 3-manifolds*, *Mem. Amer. Math. Soc.* **59** (1986), no. 339, 99–130.
- [TW15] Robert Tang and Richard Webb, *Shadows of Teichmüller disks in the curve graph*, arXiv:1510.04259, 2015.

DEPARTMENT OF MATHEMATICS, YALE UNIVERSITY
E-mail address: yair.minsky@yale.edu

DEPARTMENT OF MATHEMATICS, TEMPLE UNIVERSITY
E-mail address: samuel.taylor@temple.edu