From loom spaces to veering triangulations

Saul Schleimer and Henry Segerman

Abstract. We introduce *loom spaces*, a generalisation of both the *leaf spaces* associated to pseudo-Anosov flows and the *link spaces* associated to veering triangulations. Following work of Guéritaud, we prove that there is a locally veering triangulation canonically associated to every loom space, and that the realisation of this triangulation is homeomorphic to \mathbb{R}^3 .

1. Introduction

One-dimensional foliations, for example orbits of a flow, appeared early in the history of dynamical systems. More delicate applications, such as foliations stable for, or transverse to, a flow, arrived in due course. We refer to [17] for a readable and well-illustrated introduction to this area.

Suppose that *S* is a closed, connected, oriented surface. Suppose that $f: S \to S$ is a surface homeomorphism. The *mapping torus* for *f* is the manifold M(f) obtained from $S \times [0, 1]$ by gluing, for every $x \in S$, the point (x, 1) to the point (f(x), 0). Then M(f) is equipped with a *suspension flow* $\Phi(f)$ along the intervals; this flow has a transverse foliation given by the copies of *S*. For an example in genus one, see Figure 2.18. Suspension flows associated to surface homeomorphisms are particularly important, for example due to the work of Thurston [29, Theorem 5.6].

When f is (pseudo-)Anosov we also have the *stable* and *unstable* (singular) foliations of M(f) associated to $\Phi(f)$; the two-dimensional leaves of these foliations are obtained by taking suitable unions of flow-lines. See Examples 2.15 and 2.17; for more detail, we refer to [7, Chapter 1]. Unfortunately, the *leaf space* of $\Phi(f)$ is highly non-Hausdorff. To obtain a somewhat calmer object, we define $\mathcal{L}(f)$ to be the leaf space of the lift of $\Phi(f)$ to the universal cover of M(f). Since S intersects every orbit of $\Phi(f)$, it follows that $\mathcal{L}(f)$ is homeomorphic to \tilde{S} , the universal cover of S. Furthermore, the stable and unstable foliations for $\Phi(f)$ induce singular transverse foliations F^f and F_f of $\mathcal{L}(f)$.

Following Agol and Guéritaud [1, 16], we form S° by removing the singular points from S. Let f° be the restriction of f to S° . Thus $M(f^{\circ})$ is obtained from M(f) by *drilling*. The transverse foliations in $\mathcal{L}(f^{\circ})$ are now non-singular.

Mathematics Subject Classification 2020: 37C86 (primary); 57M50, 57Q15 (secondary).

Keywords: veering triangulations, pseudo-Anosov flows, loom spaces, astroid lemma.



Figure 1.1. The green and purple threads give transverse foliations.

The goal of this paper is to generalise leaf spaces to what we call *loom spaces*. To explain the name *loom space*, we recall that the warp and weft of a fabric, as produced by a loom, give a pair of transverse foliations. See Figure 1.1. We provide axioms, discuss a number of examples, and draw out connections to *veering triangulations* [1].

1.2. This paper

A loom space \mathcal{L} is a copy of \mathbb{R}^2 equipped with transverse (non-singular) foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$, satisfying two further axioms (Definition 2.11).

In Section 2, we list several families of examples of loom spaces. We also discuss elementary relationships between the various *skeletal rectangles* appearing in a loom space. In Section 3, we formalise the notion of a *cusp* of a loom space. These play an important combinatorial role in the rest of the work.

In Section 4, we prove a key finiteness result: the *astroid lemma* (Lemma 4.10). This places a strong restriction on the projections of certain cusps to certain leaves of the two foliations of \mathcal{L} . See Remark 4.11 for several versions of the astroid lemma appearing in previous work.

In Section 5, we review the basics of ideal triangulations and introduce *locally veering triangulations*; we show in Proposition 5.7 that these are a mild generalisation of veering triangulations. In Section 5.8, we give our version of Guéritaud's construction [16]. We then prove the following.

Proposition 5.19. *Guéritaud's construction is a functor from the category of loom spaces to the category of locally veering triangulations.*

In Section 6, we define notions of *geodesics* and *convexity* in loom spaces. Using these we prove the following.

Theorem 6.50. For any loom space, the topological realisation of its induced triangulation is homeomorphic to \mathbb{R}^3 .

1.3. Related work

Let V be the functor given by Guéritaud's construction. In our joint work with Steven Frankel, we show that V is an *equivalence* from $Loom(\mathbb{R}^2)$, the category of loom spaces, to $Veer(\mathbb{R}^3)$, the category of veering triangulations of \mathbb{R}^3 [14, Theorem 11.39]. That is, there is another functor

L: Veer(
$$\mathbb{R}^3$$
) \rightarrow Loom(\mathbb{R}^2)

so that the $L \circ V$ and $V \circ L$ admit natural transformations to the identities on $Loom(\mathbb{R}^2)$ and $Veer(\mathbb{R}^3)$, respectively. As a corollary, Guéritaud's construction gives isomorphisms of automorphism groups. We prove the above by building, from a veering triangulation \mathcal{V} of \mathbb{R}^3 , a *veering circle*, a pair of laminations in that circle, and thus a *link space* $L(\mathcal{V})$. After proving that $L(\mathcal{V})$ is a loom space, we check naturality.

In other work, joint with Jason Manning, we will show how the veering circle for \mathcal{V} compactifies the link space $L(\mathcal{V})$ to give the *veering disk* $D(\mathcal{V})$. We will then use the astroid lemma (Lemma 4.10) to give a careful description of various Hausdorff limits in $D(\mathcal{V})$. Further work will prove, when \mathcal{V} gives a finite-volume cusped hyperbolic three-manifold M, that the *veering two-sphere* is equivariantly homeomorphic to the Bowditch boundary of $\pi_1(M)$: that is, to $\partial \mathbb{H}^3$. We will then use naturality to obtain new examples of Cannon–Thurston maps.

1.4. Constructions of veering triangulations

The definition and first construction of veering triangulations in the fibred case are due to Agol [1]. The second author and collaborators generalised the definition [18]; they also answered a question of Agol, using a computer search to find the first non-fibred examples. Guéritaud [16] gave an alternative construction in the fibred case, which has inspired much later work, including this paper. We [26] announced a procedure to perform Dehn surgery along horizontal annuli or Möbius strips in veering triangulations. We gave an implementation of a special case of this in the file veering_dehn_surgery.py in our codebase [25]. Shortly afterwards, Agol and Guéritaud [2] announced an extension of Guéritaud's construction to drillings of manifolds admitting pseudo-Anosov flows without perfect fits.

A computer-generated census of all transverse veering triangulations with up to 16 tetrahedra was found by Giannopoulos and ourselves [15]. Tsang [31, Section 4] announced a procedure very similar to our veering Dehn surgery, which he calls *horizontal surgery*. He also introduced *vertical surgery* along strictly ascending loops in the stable branched surface. Landry, Minsky, and Taylor [22, Section 4] gave an exposition of the Agol– Guéritaud construction. Moreover, they proved that the veering triangulation can be made smoothly transverse to the pseudo-Anosov flow.

2. Loom spaces

2.1. Rectangles

Suppose that \mathcal{L} is a copy of \mathbb{R}^2 , equipped with two transverse foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$. We call these the *upper* and *lower* foliations, respectively.

Remark 2.2. The foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ have no singularities in \mathcal{L} . Thus, by the Poincaré–Hopf theorem [23, p. 35], any two leaves are equal, disjoint, or intersect in exactly one point. We deduce that every leaf is properly embedded in \mathcal{L} . Thus, by the Jordan curve theorem [32, p. 94], every leaf separates \mathcal{L} .

Definition 2.3. A rectangle R in \mathcal{L} is an open subset equipped with a homeomorphism $f_R: (0, 1)^2 \to R$. We require that f_R sends intervals parallel to the x-axis to arcs of $F_{\mathcal{L}}$ and sends intervals parallel to the y-axis to arcs of $F^{\mathcal{L}}$.

Lemma 2.4. With \mathcal{L} as above, the rectangles give a basis for the topology of \mathcal{L} .

Remark 2.5. Since \mathcal{L} is simply connected, we may choose orientations for the foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$. When we do this, we also assume that all rectangle maps f_R preserve these orientations. This allows us to refer to the directions south, east, north, and west in \mathcal{L} .

Definition 2.6. Suppose that $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ are oriented. Suppose that *R* is a rectangle in \mathcal{L} . Let $\gamma_t: (0, 1) \to (0, 1)^2$ be given by $\gamma_t(s) = (t, s)$. The *west side* of *R* is the set of accumulation points of the sequence of arcs $(f_R(\gamma_t))_{t\to 0}$. We define the south, east, and north sides of *R* similarly. Intersections of sides, when they exist, are called *material corners* of *R*.

Definition 2.7. Suppose that $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ are oriented. A rectangle *R* in \mathcal{L} is a *south-west* cusp rectangle if there is a continuous extension of f_R to a homeomorphism

$$\overline{f}_R: [0,1]^2 - \{(0,0)\} \to \overline{R}.$$

We define *south-east*, *north-east*, and *north-west cusp rectangles* similarly. Suppose that R is a south-west cusp rectangle. Note that the north and east sides of R are closed intervals and that the south and west sides of R are half-open intervals. We call the south and west sides of R cusp sides. We orient the south side to the east and the west side to the north. We make similar definitions for the other types of cusp rectangle.

See Figure 2.8 (a) for an example of a south-west cusp rectangle.

Definition 2.9. A rectangle *R* in \mathcal{L} is a *tetrahedron rectangle* if there are $a, b, c, d \in (0, 1)$ and a continuous extension of f_R to a homeomorphism

$$\overline{f_R}$$
: $[0,1]^2 - \{(a,0), (1,b), (c,1), (0,d)\} \to \overline{R}$.

See Figure 2.8 (d) for an example of a tetrahedron rectangle.



Figure 2.8. Examples of cusp (a), edge (b), face (c), and tetrahedron (d) rectangles. Here we indicate a point missing from the closure of a rectangle with a black dot. In Section 5.8, we transform skeletal rectangles into cells of a triangulation.

2.10. Loom spaces

We are now ready to state our main definition.

Definition 2.11. A *loom space* \mathcal{L} is a copy of \mathbb{R}^2 equipped with two transverse foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ satisfying the following axioms:

- (1) For every cusp side s of every cusp rectangle, some initial open interval of s is contained in some rectangle. See Figure 2.12.
- (2) Every rectangle is contained in some tetrahedron rectangle.



Figure 2.12. A cusp rectangle *R*. An initial segment of its southern side is contained in the shaded rectangle.

Definition 2.13. Suppose that \mathcal{L} and \mathcal{M} are loom spaces. We say that $f: \mathcal{L} \to \mathcal{M}$ is a *loom isomorphism* if

- f is a homeomorphism and
- *f* sends leaves to leaves.

Note that a loom isomorphism f may send leaves of $F^{\mathcal{L}}$ to leaves of $F^{\mathcal{M}}$ or to leaves of $F_{\mathcal{M}}$. We use $\text{Isom}(\mathcal{L}, \mathcal{M})$ to denote the set of loom isomorphisms from \mathcal{L} to \mathcal{M} . Note

that loom isomorphisms compose in the usual way. Thus loom spaces, together with their loom isomorphisms, form a category; we denote this by $Loom(\mathbb{R}^2)$. Finally, since loom isomorphisms have inverses, the set $Aut(\mathcal{L}) = Isom(\mathcal{L}, \mathcal{L})$ is a group with respect to composition.

2.14. Examples of loom spaces

Our first loom space comes from a well-known example in dynamics. The earliest exposition that we are aware of is due to Smale [27, p. 757].

Example 2.15. Suppose that $A_0 \in SL_2(\mathbb{Z})$ is an *Anosov matrix*: that is, trace $(A_0)^2 > 4$. As an example, in Figure 2.16 we take

$$A_0 = \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$



Figure 2.16. The action of A_0 on its eigenfoliations. The dots are placed at integer lattice points. The rectangles containing letters (with their usual aspect ratio) are mapped by A_0 to the corresponding shaded rectangles. These descend to T to give a *Markov partition*.

Let $T = \mathbb{R}^2/\mathbb{Z}^2$ be the two-torus; let A be the homeomorphism of T induced by A_0 . Let F^A and F_A be the resulting eigenfoliations in T. Let $x \in T$ be the image of the origin. Let \mathcal{L} be the universal cover of $T^\circ = T - \{x\}$. Define $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ by lifting the eigenfoliations. We claim that \mathcal{L} , equipped with these foliations, is a loom space. To prove this, one uses any Markov partition of $T - \{x\}$, compatible with A, to verify the axioms of Definition 2.11.

The deck transformations of the covering give examples of loom isomorphisms. We obtain two more isomorphisms by lifting the actions of the matrices

$$R = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \text{ and } G = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$$

on \mathbb{R}^2 to \mathcal{L} . It is an exercise to show that these (and the deck transformations) generate Aut(\mathcal{L}).

Our next family of examples comes from work of Thurston [30, Theorem 4(ii)]; as their name indicates, these generalise Example 2.15 to surfaces of higher genus.

Example 2.17. Suppose that *S* is a closed, connected, oriented surface with genus two or more. Suppose that $f: S \to S$ is a *pseudo-Anosov map*: that is, there are transverse measured singular foliations F^f and F_f , each preserved leafwise by f, whose measures are, by f, respectively expanded and contracted by a common factor $\lambda_f > 1$. Let $Z \subset S$ be the set of singularities of F^f and F_f . Let $S^\circ = S - Z$. Form \mathcal{L} by taking the universal cover of S° and lifting the foliations. We claim that \mathcal{L} , equipped with these foliations, is a loom space. As before, this is proved by building a Markov partition.

The element f and the deck transformations of \widetilde{S}° generate a free-by-cyclic group. In fact, this group is a finite index subgroup of Aut(\mathcal{L}) [14, Theorem 11.39]. For a discussion, see Section 1.3.

We can generalise Example 2.17 by instead taking q to be a *quadratic differential* on S. (Excellent introductions to abelian and quadratic differentials include [33, 34].) We must assume that the vertical and horizontal foliations F^q and F_q have no compact leaves. If we take Z to be the set of zeros of q, the rest of the construction is the same as Example 2.17. This gives uncountably many examples of loom spaces.

Returning to the topological theme, suppose that $f: S \to S$ is a surface homeomorphism. As discussed in Section 1, from f we form the mapping torus M(f) and its suspension flow $\Phi(f)$. For an example, see Figure 2.18.



Figure 2.18. A flow box decomposition of a torus bundle. The associated flow is isomorphic to the suspension of Example 2.15 and thus is Anosov. The manifold is homeomorphic to the longitudinal filling of the figure-eight knot complement. Taking branched covers of the bundle, branched along the suspension of the origin, we obtain examples of pseudo-Anosov flows.

Equipped with the results of this paper, it is a (difficult) exercise to show that $\Phi(f)$ is a *pseudo-Anosov flow* if and only if f is a pseudo-Anosov homeomorphism. For definitions, see [7, Section 6.6]. The exercise can also be deduced from various more general results in the literature such as [10, Main Theorem] or [11, Theorem B].

Our next example generalises the exercise to other three-manifolds; we refer to Fenley's work, in particular [9, Definition 3.2], for an overview of pseudo-Anosov flows *without perfect fits*.

Example 2.19. Suppose that M is a closed, connected, oriented three-manifold. Suppose that $\Phi: M \times \mathbb{R} \to M$ is a *topological pseudo-Anosov flow* (without perfect fits). (For the definition, see [24, p. 80].) Let Σ^{Φ} and Σ_{Φ} be the stable and unstable foliations of M. We remove from M all singular flow loops to obtain the *drilled space* M° . We restrict Φ to M° to obtain Φ° as well as the drilled foliations Σ_{\circ}^{Φ} and Σ_{Φ}° . We form the universal cover $\widetilde{M^{\circ}}$ and lift Φ° as well as Σ_{\circ}^{Φ} and Σ_{\circ}° . The *leaf space* $\mathcal{L}^{\circ} = \mathcal{L}(\widetilde{\Phi^{\circ}})$ is the quotient of $\widetilde{M^{\circ}}$ by the flow $\widetilde{\Phi^{\circ}}$. The lifts of Σ_{\circ}^{Φ} and Σ_{Φ}° descend to give non-singular foliations of \mathcal{L}° .

Alternatively, we could form the universal cover \tilde{M} and the lifted flow $\tilde{\Phi}$. Here the leaf space $\mathcal{L}(\tilde{\Phi})$ is homeomorphic to the plane \mathbb{R}^2 , equipped with a pair of singular foliations [13, Proposition 4.2]. See also [7, Lemma 6.53]. Furthermore, the singularities are a discrete subset of $\mathcal{L}(\tilde{\Phi})$. So we may remove these and take a second universal cover. This again gives \mathcal{L}° and proves that it is homeomorphic to \mathbb{R}^2 .

In [14, Theorem 10.51], we give a combinatorial proof that \mathcal{L}° is a loom space. See also recent work of Landry, Minsky, and Taylor [22, Section 4]. In addition, we show [14, Theorem 11.39] that $\pi_1(M^{\circ})$ lies in Aut(\mathcal{L}°) as a finite index subgroup.

Remark 2.20. Pseudo-Anosov flows and maps are closely related to *expansive flows* and maps, respectively. These are defined by Bowen and Walters [4]. They give suspensions as a particular example in [4, Section 4]. In [4, Theorem 6], they prove that $\Phi(f)$ is expansive if and only if f is expansive. For further discussion of the subtle connections between expansive and pseudo-Anosov flows, we refer the reader to [3,6].

Removing the singular orbits of a pseudo-Anosov flow yields a manifold with torus boundary components. Our final examples are closely related to these *drilled* flows but are completely combinatorial. These rely on Agol's notion of a *veering triangulation* [1, Definition 4.1]; see Section 5 for precise definitions and Figure 2.21 for a concrete example.



Figure 2.21. The veering triangulation for the figure-eight knot complement.

Example 2.22. Suppose that M is a compact, connected, oriented three-manifold with boundary ∂M a non-empty collection of tori. Suppose that \mathcal{V} is a *veering* triangulation of M: that is, an ideal triangulation of the interior of M equipped with a *taut* structure and a veering colouring. (We give the definitions in Section 5.1.) In [14, Theorems 10.51 and 11.39], we show that

- there is a canonical *link space* \mathcal{L} associated to the lift of \mathcal{V} to the universal cover of M,
- \mathcal{L} is a loom space and
- $\pi_1(M)$ is finite index in Aut(\mathcal{L}).

For a discussion, see Section 1.3.

The overall goal of this paper is to provide the converse to Example 2.22. In Proposition 5.19, from a given loom space, we build a *locally veering* triangulation. In Theorem 6.50, we prove that the realisation of this triangulation is homeomorphic to \mathbb{R}^3 .

Remark 2.23. All of the examples of loom spaces given above have large automorphism groups. It is interesting to contemplate how one might obtain a finitely described, yet aperiodic, loom space.

2.24. Cusps repel

From now on, we will assume that \mathcal{L} , equipped with the foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$, is a loom space in the sense of Definition 2.11. We begin our analysis of \mathcal{L} by proving a version of a condition introduced by Keane [19, Section 2]. In the language of foliations, Keane's *minimality condition* on a singular foliation is that there are no leaves connecting singularities. Our foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ do not have singularities; we instead use the "missing points" from the sides of tetrahedron rectangles.

Lemma 2.25. Suppose that R is a tetrahedron rectangle with associated parameters a, b, c, and d. Then $a \neq c$ and $b \neq d$.

Proof. For a contradiction, suppose that d = b. See Figure 2.26. Let *m* be the leaf of $F^{\mathcal{X}}$ inside of *R* running from (0, b) to (1, b). Fix a small positive number $\varepsilon < \frac{1}{2}$ such that $b + \varepsilon < 1$. We now build two open rectangles in \mathbb{R}^2 as products of open intervals:





Figure 2.26. A tetrahedron rectangle with b = d. We construct the rectangle C (shaded) in the proof.

Thus the images $f_R(A)$ and $f_R(B)$ are cusp rectangles in \mathcal{L} with south sides contained in the leaf m.

Applying Definition 2.11 (1) twice, we obtain small rectangles A' and B' containing initial segments of the west and east sides of A and B, respectively. Cutting these with a leaf of $F_{\mathcal{L}}$, we obtain a single rectangle $C \subset \mathcal{L}$ which has four material corners and which contains m in the interior of its south side. Thus C is not contained in any tetrahedron rectangle. This contradicts the axiom Definition 2.11 (2), as desired.

A similar argument deals with the case that a = c.

2.27. Skeletal rectangles

For the next two definitions, we choose orientations as in Remark 2.5. See Figure 2.8 for the following definitions and lemmas.

Definition 2.28. A rectangle *R* in \mathcal{L} is a *red edge rectangle* if there is a continuous extension of f_R to a homeomorphism

$$\overline{f}_R: [0,1]^2 - \{(0,0), (1,1)\} \to \overline{R}.$$

An edge rectangle R is *blue* if the missing points are instead (0, 1) and (1, 0).

Definition 2.29. A rectangle *R* in \mathcal{L} is a *south-west face rectangle* if there are $a, b \in (0, 1)$ and a continuous extension of f_R to a homeomorphism

$$\bar{f}_R: [0,1]^2 - \{(0,0), (1,a), (b,1)\} \to \bar{R}.$$

We define the three other types of face rectangle similarly.

We call edge, face, and tetrahedron rectangles *skeletal* rectangles. We do not include cusp rectangles among the skeletal rectangles because they are not uniquely determined by their cusps.

Lemma 2.30. We have the following:

- Every tetrahedron rectangle contains exactly four face rectangles.
- Every tetrahedron rectangle contains exactly six edge rectangles.
- Every face rectangle contains exactly three edge rectangles.

Proof. Suppose that R is a tetrahedron rectangle, and let f_R be the given parametrisation. There are at most four face rectangles in R; the sides of each intersect three of the four sides of R.

Let (a, 0) and (c, 1) be the missing points on the southern and northern sides. Appealing to Lemma 2.25 and breaking symmetry, suppose that a < c. Then

$$F = f_R(\{(x, y) \in (0, 1)^2 \mid x > a\})$$

is one of the desired face rectangles. The remaining three are formed similarly.

The other two statements are proved similarly.

Lemma 2.31. Every face rectangle is contained in exactly two tetrahedron rectangles.

Proof. Breaking symmetry, we may assume that F is a north-west face rectangle. See Figure 5.13 (a). Let δ_F be the northern side of F. Suppose that C is a small rectangle contained in F; suppose that the north and west sides of C are contained in the north and west sides of F. Let $\delta_C \subset \delta_F$ be the northern side of C.

Since $C \subset F$, we deduce that *C* is a cusp rectangle and that δ_C is a cusp side. By Definition 2.11 (1), we have that an initial segment of δ_C is contained in a rectangle, say *D*. Note that $\varepsilon = \delta_F - D$ is compact in \mathcal{X} . Applying Lemma 2.4, we cover ε by a finite collection of rectangles. We deduce that there is a rectangle *F'* so that *F'* contains both *F* and δ_F . We appeal to Definition 2.11 (2) to obtain a tetrahedron rectangle *P* containing *F'*.

Repeating the argument with the western side of F gives another tetrahedron rectangle Q containing F. Note that δ_F lies in the interior of P and lies in the northern side of Q. Hence P and Q are distinct.

We now show that there are at most two tetrahedron rectangles containing F. Suppose that R is any tetrahedron rectangle containing F. Then the southern and eastern sides of R contain those of F. By Definition 2.29, the closure of F is missing three points. The missing point at the north-western corner is necessarily the missing point from either the western or northern side of R. We deduce that R is thus equal to P or to Q.

It is more difficult to prove that an edge rectangle E is contained in only finitely many face rectangles. This is deferred to Corollary 4.18.

3. Cusps and corners

The *cusps* of a loom space \mathcal{L} provide the beginnings of a boundary at infinity for \mathcal{L} ; we construct the boundary in a future work, see Section 1.3. This section provides the background needed for the statement of the astroid lemma.

Definition 3.1. Suppose that R and Q are cusp rectangles in \mathcal{L} . We say that R is *equivalent* to Q if there is a finite sequence of cusp rectangles

$$(R = R_0, R_1, \ldots, R_n = Q)$$

such that for each pair (R_i, R_{i+1}) , some cusp side of one is contained in some cusp side of the other.

The equivalence relation is illustrated in Figure 3.2.

Definition 3.3. A cusp is an equivalence class of cusp rectangles.

For any cusp rectangle R, if c = [R], then we refer to R as a *cusp rectangle for* c.



Figure 3.2. In each subfigure, the labelled cusp rectangles are equivalent.

Definition 3.4. Suppose that Q is a subset of \mathcal{L} . We say that c is a *cusp of* Q if there is a sequence of rectangles $(R_i)_{i=0}^{\infty}$ with the following properties:

- (1) R_i is a cusp rectangle for c, for all i.
- (2) $\overline{R_i} \cap Q$ is non-empty, for all *i*.
- (3) $R_{i+1} \subset R_i$, for all *i*.
- (4) The intersection $\bigcap_i R_i$ is empty.

We define $\Delta(Q)$ to be the set of cusps of Q.

Recall from Definition 2.6 that a rectangle may have as many as four material corners. Cusps provide any remaining corners, as follows.

Definition 3.5. Suppose that $Q \subset \mathcal{L}$ is a rectangle. Suppose that Q contains a cusp rectangle R where the cusp sides of R are contained in sides of Q. Then we call c = [R] an *ideal corner* of Q.

Lemma 3.6. Every rectangle has four corners. At most two of these are ideal.

Proof. Suppose that *R* is the given rectangle. By Definition 2.11 (2), we have that *R* is contained in a tetrahedron rectangle. The result now follows from Lemma 2.25.

Definition 3.7. Fix a rectangle $R \subset \mathcal{L}$. Suppose that x and y are corners (material or ideal) of R. We say that x and y are *adjacent* if they are incident to a single side of R. If x and y are not adjacent, then they are *opposite*.

Definition 3.8. Suppose that ℓ is a leaf of $F^{\mathcal{L}}$ or $F_{\mathcal{L}}$. Suppose that *R* is a cusp rectangle in \mathcal{L} with a cusp side δ . If δ is contained in ℓ , then we call ℓ a *cusp leaf for c*.

It follows that δ , the cusp side of *R*, contains an end of the leaf ℓ .

Remark 3.9. Suppose that c is a cusp. As in Remark 2.2, the Poincaré–Hopf theorem implies that any two cusp leaves for c are disjoint. It follows that any leaf intersecting the interior of a rectangle R is not a cusp leaf for an ideal corner of R.

Lemma 3.10. Any leaf ℓ of $F^{\mathcal{X}}$ (or of $F_{\mathcal{X}}$) is a cusp leaf for at most one cusp.

Proof. Suppose that c and d are distinct cusps of \mathcal{L} . Suppose for a contradiction that ℓ is a cusp leaf for both c and d. Let R and Q be cusp rectangles at c and d with cusp sides γ and δ , both contained in ℓ . If γ and δ contain the same end of ℓ , then c = d, contrary to the assumption. Thus γ and δ contain the two ends of ℓ . By Definition 2.11(1), there are rectangles R' and Q' that contain initial segments of γ and δ , respectively, in their interiors.

Applying Lemma 2.4, we cover the compact interval $\ell - (R' \cup Q')$ by finitely many rectangles. We deduce that all of ℓ is contained in a single rectangle. According to Definition 2.11 (2), this rectangle is contained in a tetrahedron rectangle. Appealing to Lemma 2.25, we arrive at the desired contradiction.

Lemma 3.11. Suppose that R is an edge rectangle. Then R has two cusps.

Proof. Let x be an interior point of R. Let ℓ^x be the leaf of $F^{\mathscr{X}}$ containing x. Let P and Q be the two components of $R - \ell^x$. These are both cusp rectangles. We must prove that $[P] \neq [Q]$.

Set c = [P]. Let ℓ^c and m_c be the cusp leaves containing the cusp sides of P. Suppose that P' is any cusp rectangle equivalent to, and disjoint from, P. Let

$$(P = P_0, P_1, \ldots, P_n = P')$$

be a minimal sequence of cusp rectangles satisfying Definition 3.1. Minimality implies that P_i is disjoint from P_{i+1} . See Figure 3.2. By Remark 2.2, one of the leaves ℓ^c or m_c separates P_1 from R, and thus from Q. By minimality and induction, the same leaf separates P_k from Q for all k > 0.

Lemma 3.12. Suppose that R and R' are edge rectangles. Then R = R' if and only if $\Delta(R) = \Delta(R')$.

Proof. The forward direction follows from Definition 3.4. Now suppose that $R \neq R'$. Breaking symmetry, suppose that x lies in R and does not lie in R'. Let ℓ^x and m_x be the leaves of $F^{\mathcal{X}}$ and $F_{\mathcal{X}}$, respectively, that contain x. By Remark 2.2, the intersection $\ell^x \cup m_x$ is the singleton set $\{x\}$. By Remark 3.9, the leaves ℓ^x and m_x are not cusp leaves for either cusp of R. If ℓ_x and m_x both intersect R', then by Definition 2.3 they meet in a point $y \in R'$. Since $x \notin R'$, we have that $x \neq y$, and we have reached a contradiction.

Breaking symmetry, we have that ℓ_x does not intersect R'. Let A and B be the two components of $\mathcal{L} - \ell^x$. Again appealing to Remark 2.2, the rectangle R' is in one of A or B, but not both. Using Remark 2.2 once more, we deduce that the sets $\Delta(A) \cap \Delta(R)$ and $\Delta(B) \cap \Delta(R)$ are both singletons. Therefore, $\Delta(R) \neq \Delta(R')$.

4. The astroid lemma

Here we prove the *astroid lemma* (Lemma 4.10). This controls the projection of certain cusps to certain leaves of $F^{\mathcal{X}}$ and $F_{\mathcal{X}}$.

4.1. Staircases

Suppose that x is a point or a cusp of \mathcal{L} . Fix any rectangle R with a corner at x. Following Guéritaud [16, Section 4.3], we make the following definition.

Definition 4.2. The *staircase* $\Gamma(x, R)$ is the closure of the union of all rectangles $Q \subset \mathcal{L}$, where

- x is a corner of Q and
- $Q \cap R$ is non-empty.

We often write $\Gamma(x) = \Gamma(x, R)$, suppressing the choice of *R*.

Definition 4.3. We take $m = m(x, R) \subset F_{\mathcal{X}}$ to be the union of the arcs s in $F_{\mathcal{X}}$ so that there is a rectangle Q so that

- $Q \subset \Gamma(x, R)$,
- s is a side of Q, and
- *s* has an endpoint at *x*.

We define $\ell = \ell(x, R) \subset F^{\mathcal{X}}$ similarly. We call *m* and ℓ the *lower* and *upper axis rays*, respectively, for $\Gamma(x, R)$.

Definition 4.4. Suppose that $\Gamma(x, R)$ is a staircase. Suppose that *m* and ℓ are its lower and upper axis rays. Any cusp c_m of *m* is called an *axis cusp* for *m*. We make the same definition for a cusp c_ℓ of ℓ .

By Lemma 3.10, each axis ray has at most one axis cusp.

Lemma 4.5. Suppose that $\Gamma(x) = \Gamma(x, R)$ is a staircase. Breaking symmetry, suppose that x is the south-west corner of R. Suppose that m' is an initial segment of the lower axis ray m. Then there is a rectangle $R' \subset \Gamma(x)$ so that

- x is the south-west corner of R' and
- *m'* is the south side of *R'*.

Proof. Note that x is either a point or a cusp of \mathcal{L} . Thus using either Lemma 2.4 or Definition 2.11 (1), respectively, we deduce that there is an initial segment m'' of m contained in a rectangle Q. If m'' contains m', then we cut Q using the axis rays to obtain R'.

If not, there are two cases. Either the axis cusp c_m is contained or not contained in m'. If c_m is not contained in m', then the interval m' - Q is compact and so covered by a rectangle Q'. Cutting $Q \cup Q'$ by the axis rays and reducing their height gives the desired rectangle R'. If c_m is contained in m', then we apply Definition 2.11(1) three times, and the remaining argument is as above.

Lemma 4.6. Suppose that x is the south-west corner of R. Then the lower axis ray m has the following properties:

- Suppose x is a cusp. Then m is a cusp leaf.
- Suppose x is not a cusp. Let m_x be the leaf of $F_{\mathcal{X}}$ containing x.
 - If m_x is a non-cusp leaf (or a cusp leaf emanating from the east of its cusp), then m is the eastern component of $m_x x$.
 - If m_x is a cusp leaf emanating from the west of its cusp, then its cusp c_m is the axis cusp for m. In this case, $m m_x$ is a cusp leaf emanating from the east of c_m .

A similar statement holds when x is one of the other corners of R. Similar statements also hold for the upper axis ray ℓ .

Proof. Suppose that x is a cusp. Shrinking R slightly, we may assume that R is a cusp rectangle for x. Let m' be the cusp leaf to the east of x which contains the southern side of R. We must show that the lower axis ray m equals m'.

Note *m* is the union of connected sets (the southern sides of rectangles), all intersecting a connected set (the southern side of *R*). Thus *m* is connected. We deduce that *m* is contained in m'.

It remains to prove that m' is contained in m. Let m'' be any closed initial segment of m'. By Lemma 4.5, the segment m'' is the southern side of some rectangle R', showing that m'' lies in m. We deduce that m' is contained in m.

The remaining cases are similar.

Definition 4.7. A cusp *c* is an *exterior cusp* of $\Gamma(x)$ if there is a rectangle in $\Gamma(x)$ having *c* and *x* as opposite corners. We define $\Delta(\Gamma(x))$ to be the set of exterior cusps of $\Gamma(x)$.

Note that $\Delta(\Gamma(x)) \subset \Delta(\Gamma(x))$. When x is a cusp, or when axis cusps exist, the containment $\Delta(\Gamma(x)) \subset \Delta(\Gamma(x))$ is proper. See Figure 4.8.

Let ∂m denote the end of the lower axis ray m which is not at x (or at the axis cusp c_m , if it exists). Note that ∂m is not at a cusp by Lemmas 3.10 and 4.6. We define $\partial \ell$ similarly.

We define a pair of projections

$$\pi_m: \Delta(\Gamma(x)) \to m \text{ and } \pi^\ell: \Delta(\Gamma(x)) \to \ell$$

as follows. Suppose that $c \in \Delta(\Gamma(x))$ is an exterior cusp. Then there is a rectangle $Q \subset \Gamma(x)$ with opposite corners at x and c. We define $\pi_m(c)$ and $\pi^{\ell}(c)$ to be the corners of Q, other than x, lying on m and ℓ , respectively. See Figure 4.8.



Figure 4.8. A staircase. Cusps are indicated with black dots. Labelled material points are indicated by yellow dots. In this example, x is a material point of \mathcal{L} to the west of a cusp c_m , the axis cusp of the lower axis ray m.

4.9. Statement and proof

We now control the images of the projections π_m and π^{ℓ} .

Lemma 4.10 (Astroid lemma). Suppose that $\Gamma(x)$ is a staircase in \mathcal{L} ; suppose that m and ℓ are its axis rays.

- (1) The image of π_m does not accumulate at any interior point of m (nor does it accumulate at the axis cusp c_m , if present).
- (2) The image of π_m accumulates at x and at ∂m .

Similar statements hold for π^{ℓ} .

Remark 4.11. Loom spaces associated to pseudo-Anosov homeomorphisms have various natural non-complete euclidean metrics. Each such metric has a definite injectivity radius; the astroid lemma is immediate in these cases. See [8, Lemma 14], [16, Figure 12], and [21, Figure 12].

Loom spaces associated to pseudo-Anosov flows (without perfect fits) on finite volume hyperbolic three-manifolds need not have a natural choice of metric. However, in this setting the action of the fundamental group still gives a local finiteness that can replace the lower bound on injectivity radius. This, in slightly different language, is carried out in [22, Section 4]. See, in particular, Figure 18 therein.

Proof of Lemma 4.10. Breaking symmetry, suppose that x is south-west of $\Gamma(x)$. Suppose that, in contradiction to (1), there is a sequence of distinct exterior cusps $c_i \in \Delta(\Gamma(x))$ so that $r_i = \pi_m(c_i)$ accumulates at r_{∞} , an interior point of m (possibly the axis cusp c_m). Define $s_i = \pi^{\ell}(c_i)$. Let $R_i \subset \Gamma(x)$ be the rectangle with corners at x, r_i, c_i , and s_i . Since the c_i are all distinct, by Lemma 2.25 the points r_i and s_i are also all distinct.



Figure 4.12. The two possibilities as $r_i = \pi_m(c_i)$ is increasing or decreasing along *m*.

We orient *m* and ℓ away from *x*. We pass to a subsequence of the c_i to ensure that the sequence (r_i) is strictly monotonic in *m*. Note that the rectangles R_i cannot nest; we deduce that the sequence (s_i) is strictly monotonic in ℓ . Likewise, exactly one of the sequences (r_i) and (s_i) is increasing while the other is decreasing. See Figure 4.12.

We break the proof into two cases.

Case 1. Suppose that (r_i) is increasing in m. Thus (s_i) is decreasing along ℓ . See Figure 4.12 (a). Since m can be realised as an increasing union of southern sides of rectangles, by Lemma 4.6 we have a rectangle Q in the staircase with corners at x and r_{∞} . Thus the points s_i do not enter the interior of the western side of Q. That is, the sequence (s_i) is bounded away from x in ℓ . Thus there is some s_{∞} where they accumulate. Note that s_{∞} may be c_{ℓ} , the axis cusp of ℓ .

Let m_{∞} be the ray of $F_{\mathcal{X}}$ emanating from s_{∞} and entering $\Gamma(x)$. Recall that R_i is the rectangle with opposite corners at x and c_i . Define R'_i to be the component of $R_i - m_{\infty}$ with a corner at x. Define R'_{∞} to be the increasing union of the R'_i . Thus R'_{∞} is a rectangle. Let c_{∞} be its northeastern corner.

By Lemma 3.6, we have that c_{∞} is a point or a cusp of \mathcal{L} . Appealing to Lemma 2.4 or Definition 2.11 (1), respectively, we get that there is a small rectangle *P* so that

- c_{∞} lies in the interior of the eastern side of P and
- the interior of P intersects m_{∞} .

However, the projections of (c_i) accumulate on $s_{\infty} \in \ell$ and $r_{\infty} \in m$, respectively. Thus the c_i enter P, a contradiction. Again, see Figure 4.12 (a).

Case 2. Suppose that (r_i) is decreasing along *m*. Thus (s_i) is increasing along ℓ . See Figure 4.12 (b). Let ℓ_{∞} be the ray of $F^{\mathcal{X}}$ emanating from r_{∞} and entering $\Gamma(x)$. Define R'_i

to be the component of $R_i - \ell_{\infty}$ with a corner at x. Define R'_{∞} to be the union of the R'_i . Again, R'_{∞} is a rectangle. Note that the west side of R'_{∞} is contained in ℓ . By Definition 2.11 (2), there is a tetrahedron rectangle Q containing R'_{∞} . So the west side of R'_{∞} is either contained in the interior of Q or is contained in the west side of Q. In either case, we obtain an upper bound for the s_i . We now apply the previous argument, swapping the roles of m and ℓ . This completes the proof of (1).

To prove (2), we must find a sequence of exterior cusps $c_i \in \Delta(\Gamma(x))$ whose projections $\pi_m(c_i)$ accumulate at ∂m and whose projections $\pi^{\ell}(c_i)$ accumulate at x. Let (m_i) be an increasing sequence of open initial segments of m, whose union is m. By Lemma 4.5, there is a rectangle R_i with south-west corner at x and whose southern side is m_i . Let Q_i be the union of all rectangles Q so that

- Q contains R_i and
- the west sides of Q and R_i are identical.

Since Q_i is a rectangle, by Definition 2.11 (2) we have a tetrahedron rectangle, R'_i , containing Q_i . Thus there is a cusp c_i contained in the east side of R'_i . Note that c_i is north of *m*, by the construction of Q_i . So there is a rectangle in $\Gamma(x)$ with opposite corners at *x* and c_i ; thus c_i is an exterior cusp of $\Gamma(x)$.

By construction, the projection $\pi_m(c_i)$ is not contained in m_i . Since the m_i exhaust m, the sequence of projections accumulates on ∂m . It follows that the sequence of projections $\pi^{\ell}(c_i)$ is decreasing in ℓ . By (1), the sequence $\pi^{\ell}(c_i)$ tends to x. This proves (2) for m; the proof for ℓ is similar.

We record a few consequences of the astroid lemma.

Corollary 4.13. The cusp leaves are dense in $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$.

See also [22, Lemma 4.2(1)].

Corollary 4.14. Suppose that c and d are exterior cusps for $\Gamma(x)$. Suppose their projections to ℓ^x (or m_x) are consecutive. (That is, not separated by the image of any other exterior cusp.) Then there is an edge rectangle $R \subset \Gamma(x)$ having c and d as opposite corners.

4.15. Finiteness and connectedness

Lemma 4.16. For any rectangle R, there are only finitely many skeletal rectangles containing it.

Proof. Suppose that *R* is the given rectangle. Let *x* be the south-west corner of *R*. Let $\Gamma(x) = \Gamma(x, R)$ the resulting staircase; let *m* and ℓ be its axis rays.

Let $m_R \subset m$ be the projection of R to m, along $F^{\mathcal{X}}$. Similarly, let $\ell^R \subset \ell$ be the projection of R to ℓ , along $F_{\mathcal{X}}$. See Figure 4.17.

By Lemma 4.10(2), there are exterior cusps c and d in $\Delta(\Gamma(x))$ so that $\pi_m(c)$ lies in m_R and $\pi^{\ell}(d)$ lies in ℓ^R . By Lemma 4.10(1), there are only finitely many cusps



Figure 4.17. The staircase $\Gamma(x, R)$.

 $c' \in \Delta(\Gamma(x))$ so that $\pi_m(c')$ lies between $\pi_m(c)$ and $\pi_m(d)$. Furthermore, we may replace *x* with any other corner of *R* and perform the same analysis in the corresponding staircase. This determines a finite collection of cusps. By Lemma 3.12, an edge rectangle is determined by its cusps. A skeletal rectangle is determined by the edge rectangles it contains. This gives the desired bound.

Corollary 4.18. Any edge rectangle is contained in only finitely many face rectangles.

Proof. This follows from Lemmas 4.16 and 2.30.

Before giving the next result, we require two definitions.

Definition 4.19. Suppose that *P* and *Q* are distinct tetrahedron rectangles. We say that *P* and *Q* are *face-adjacent* if their intersection, $P \cap Q$, is a face rectangle.

In general, we say that two tetrahedron rectangles P and Q are *face-connected* if there is a finite sequence $(P = P_0, P_1, ..., P_n = Q)$ of tetrahedron rectangles, where P_i and P_{i+1} are face-adjacent for all i.

Note that every tetrahedron rectangle is face-connected to itself.

Definition 4.20. Suppose that P and Q are rectangles of \mathcal{L} . We say that P west-east spans Q if there is a leaf of the induced foliation F_Q that is contained in P. We say that P properly west-east spans Q if, additionally, P - Q has two components. We define south-north spans and properly south-north spans similarly. Finally, we say that P (properly) spans Q if either P (properly) west-east spans Q, or P (properly) south-north spans Q.

Remark 4.21. Note that the definition of west-east spans is independent of any choices of orientation made as in Remark 2.5.

Lemma 4.22. Suppose that P and Q are tetrahedron rectangles. Suppose that P westeast spans Q. Then there is a sequence of tetrahedron rectangles $(P = P_0, P_1, ..., P_n = Q)$ so that

- P_i is face-adjacent to P_{i+1} and
- P_i west-east spans P_{i+1} .

Proof. Let $R = P \cap Q$. We define tet(R) to be the set of tetrahedron rectangles that contain *R*. By Lemma 4.16, the set tet(R) is finite. We induct on the size of tet(R). In the base case R = P = Q, so tet(R) has exactly one element and there is nothing to prove.

In general, let *F* be any face rectangle of *P* that south-north spans *P* and contains *R*. By Lemma 2.31, there is exactly one tetrahedron rectangle *P'* that is, via *F*, face-adjacent to *P*. Note that *P* west-east spans *P'* which in turn west-east spans *Q*. See Figure 4.23. Here the widest rectangle is *P*, the tallest is *Q*, and the remaining tetrahedron rectangle is *P'*. Set $R' = P' \cap Q$ and note that $R \subset R'$. Thus $tet(R') \subset tet(R)$. Furthermore, *P* is an element of tet(R) but is not an element of tet(R'). The induction hypothesis now implies that *P'* is face-connected to *Q*, using only the tetrahedra in tet(R'), completing the proof.



Figure 4.23. A possible picture for Lemma 4.22.

Proposition 4.24. The set of tetrahedron rectangles of \mathcal{L} is face-connected.

Proof. Suppose that *P* and *Q* are tetrahedron rectangles. Choose an arc $\gamma \subset \mathcal{L}$ connecting a point of *P* to a point of *Q*. Recall that by Lemma 2.4, the open rectangles give a basis for the topology. Also, γ is compact. Thus γ admits a finite covering by rectangles. By Definition 2.11 (2), the arc γ is covered by a finite collection of tetrahedron rectangles.

Thus we are reduced to the case where P and Q intersect. Let $R = P \cap Q$. Let R' be the rectangle so that

- R' contains R,
- the west and east sides of R' contain the west and east sides of R, respectively, and
- *R'* is maximal with respect to the above two properties.

From Definition 2.11 (2), we deduce that R' is a tetrahedron rectangle. From the construction, we deduce that both P and Q west-east span R' (and perhaps one or both equal R'). The proposition now follows from two applications of Lemma 4.22.

We deduce the following.

Corollary 4.25. There are countably infinitely many tetrahedron rectangles. Thus the same holds for face rectangles, edge rectangles, cusps, and cusp leaves.

This implies that there are only countably many cusp leaves. We deduce the following.

Corollary 4.26. The non-cusp leaves are dense in $F^{\mathcal{X}}$ and $F_{\mathcal{X}}$.

5. Locally veering triangulations

In this section, we review several combinatorial structures on triangulations of three-manifolds. We introduce the notions of *taut isomorphisms* and *locally veering triangulations*. We then follow Guéritaud [16, Section 2] to construct a locally veering triangulation from a loom space. In Proposition 5.19, we show that this construction is functorial.

5.1. Definitions

A useful example for the first several definitions is the canonical triangulation of the figure-eight knot complement. See Figure 2.21.

Let

$$t^{3} = \{x \in \mathbb{R}^{4} \mid x_{i} \geq 0 \text{ and } \sum x_{i} = 1\}$$

be the *standard* tetrahedron. This is equipped with the subspace topology. Note that the vertices of t^3 are the standard unit vectors. Their usual ordering gives an orientation to t^3 .

A copy t of t^3 is called a *model* tetrahedron. See Figure 5.2 (a). The facets (faces, edges, and vertices) of t are called *model facets*. Note that t also inherits an orientation.



Figure 5.2. Cusped model tetrahedron (a); face pairing (b).

Suppose that t and t' are model tetrahedra (which may be equal). Suppose that f and f' are faces of t and t', respectively. Suppose that $\phi: f \to f'$ is a homeomorphism induced by restricting an affine map. We call ϕ a *face pairing*. See Figure 5.2 (b).

Essentially following [28, Section 4.2], we define an ideal triangulation

$$\mathcal{T} = (\{t_{\alpha}\}, \{\phi_{\beta}\})$$

to be a collection of model tetrahedra and a collection of face pairings. The *realisation* of \mathcal{T} , denoted by $|\mathcal{T}|$, is the topological space obtained as follows:

- Take the disjoint union of the model tetrahedra.
- Quotient by the face pairings.
- Remove the zero-skeleton of the result.

The *realisation* of a model facet of \mathcal{T} is its image in $|\mathcal{T}|$. The *models* of a realised facet in $|\mathcal{T}|$ are its preimages in \mathcal{T} . In order to ensure that $|\mathcal{T}|$ is a three-manifold, we require the following:

- If ϕ is a face pairing, then so is ϕ^{-1} .
- Every model face occurs in exactly two face pairings.
- No face is paired with itself.
- Every edge has only finitely many models.
- The models of a single edge can be consistently oriented.

A *taut* structure on a model tetrahedron t is an assignment of dihedral angle of either zero or π to the model edges of t. The dihedral angles are required to satisfy the following:

Suppose that v is a model vertex of t. Suppose that e, e', e" are the model edges of t adjacent to v. Then the sum of their dihedral angles is π.

In a taut tetrahedron, the edges with dihedral angle zero are called *equatorial* while the edges with dihedral angle π are called *diagonal*. See Figure 5.4.

Following [18, Definition 1.1] (see also [20, p. 370]), we say an ideal triangulation \mathcal{T} is *taut* if all model tetrahedra are taut, and we have the following:

• Suppose that *e* is an edge of $|\mathcal{T}|$. Then the dihedral angles of the models of *e* sum to 2π .

Definition 5.3. Suppose that \mathcal{T} and \mathcal{S} are taut ideal triangulations. If $f: \mathcal{T} \to \mathcal{S}$ is an isomorphism of triangulations and sends the taut structure on \mathcal{T} to that on \mathcal{S} , then we call f a *taut isomorphism*.

Again, taut isomorphisms compose in the usual way. Thus taut triangulations, together with taut isomorphisms, form a category denoted Taut. We use $\text{Isom}(\mathcal{T}, \mathcal{S})$ to denote the set of taut isomorphisms from \mathcal{T} to \mathcal{S} ; we use $\text{Aut}(\mathcal{T})$ to denote the group of taut automorphisms. It is an exercise to check that the taut automorphism group of the triangulation

of the figure-eight knot complement, shown in Figure 2.21, is isomorphic to the symmetries of the square.

An oriented taut tetrahedron *t* is *veering* if there is a bi-colouring (by red and blue) of the model edges as follows:

• Suppose that, at a model vertex v, we have adjacent model edges e, e', e''. Suppose that this is the anticlockwise ordering on the edges, as viewed from outside of t and using the induced orientation on ∂t . Then if e has dihedral angle π , we have that e' is blue and e'' is red.

See Figure 5.4.



Figure 5.4. A veering tetrahedron. The π angle edges may be either red or blue.

Suppose now that we have fixed an orientation of a taut ideal triangulation \mathcal{V} . Following [18, Definition 1.3] (see also [1, Definition 4.1]), we define a *veering* structure on \mathcal{V} to be a colouring (by red and blue) of the edges of $|\mathcal{V}|$ that pull back to give veering structures on all of the oriented model taut tetrahedra. For an example, see Figure 2.21.

We now turn to generalising veering triangulations to non-orientable manifolds. Suppose that M is a three-manifold. Suppose that \mathcal{T} is a taut ideal triangulation of M. Suppose that e is an edge of \mathcal{T} . Let (e_i) be the collection of models of e. We order the e_i cyclically, as we walk about e in M. For each model edge e_i , let t_i be a copy of the model tetrahedron containing e_i . For each i, let ϕ_i be the face pairing, from $f_i \subset t_i$ to $g_i \subset t_{i+1}$, so that $\phi_i(e_i) = e_{i+1}$. We define $\mathcal{T}_e = (\{t_i\}, \{\phi_i\})$ to be the *model edge neighbourhood* of e in \mathcal{T} . Note that \mathcal{T}_e inherits a taut structure from \mathcal{T} . Also, its realisation $|\mathcal{T}_e|$ is a three-ball.

Choose an orientation on $|\mathcal{T}_e|$. We say that \mathcal{T} is *veering at e* if \mathcal{T}_e admits a veering colouring.

Definition 5.5. Suppose that M is a three-manifold. Suppose that \mathcal{T} is a taut ideal triangulation of M. Then \mathcal{T} is *locally veering* if \mathcal{T} is veering at every edge.

Example 5.6. The Gieseking manifold M_G can be obtained as a punctured torus bundle; the monodromy is the matrix G given in Example 2.15. Thus M_G admits a taut ideal triangulation \mathcal{V} with a single tetrahedron. The orientation double cover of M_G is the

figure-eight knot complement; the taut triangulation \mathcal{V} lifts to give the one shown in Figure 2.21. Thus \mathcal{V} is locally veering.

Locally veering triangulations are "almost" veering, in the following sense.

Proposition 5.7. Suppose that M is a three-manifold. Suppose that \mathcal{V} is a taut ideal triangulation of M. Suppose that (M, \mathcal{V}) is locally veering. Then (M, \mathcal{V}) admits a veering structure if and only if M is orientable.

Proof. Suppose that (M, \mathcal{V}) admits a veering structure. Then each model tetrahedron is veering, thus oriented. These orientations disagree across faces (because of the colours of the edges). Thus they define an orientation of M.

For the converse, we fix an orientation of M. This induces an orientation on all tetrahedra and thus on all model tetrahedra. The taut structure picks four of the six model edges to be equatorial. The orientation of the model tetrahedron and the axioms of a veering triangulation determines the colours of the equatorial model edges. The hypothesis of local veering tells us that all equatorial models of a given edge have the same colour.

We use Veer to denote the full subcategory of Taut consisting of locally veering triangulations. We use $Veer(\mathbb{R}^3)$ to denote the further full subcategory of those triangulations whose realisation is homeomorphic to \mathbb{R}^3 .

5.8. Building the triangulation

We give a (metric-free) version of Guéritaud's construction [16, Section 2]. That is, for every loom space \mathcal{L} , we give a locally veering triangulation $\mathcal{V} = V(\mathcal{L})$.

Definition 5.9. Suppose that \mathcal{L} is a loom space. Its *induced* triangulation $V(\mathcal{L})$ has *model cells*

 $\{c(P) \mid P \text{ is a skeletal rectangle of } \mathcal{L}\}.$

We identify the vertices of c(P) with the cusps of P. These are distinct by Lemma 3.11. If P and Q are skeletal rectangles with P properly contained in Q, then the set of cusps $\Delta(P)$ is a proper subset of $\Delta(Q)$. For every such pair, we take $\phi_{P,Q}$ to be the corresponding cell identification between c(P) and the corresponding subsimplex of c(Q). The realisation of $V(\mathcal{L})$ is the resulting quotient (minus the zero-skeleton).

Definition 5.10. Suppose that $f: \mathcal{L} \to \mathcal{M}$ is a loom isomorphism. We define the *induced* map $V_f: V(\mathcal{L}) \to V(\mathcal{M})$ as follows. Suppose that $P \subset \mathcal{L}$ is a skeletal rectangle. Let c(P) be the corresponding cell of $V(\mathcal{L})$; so c(P) is either an edge, face, or tetrahedron. We take $V_f(c(P)) = c(f(P))$.

We also use the notation V(f) for V_f .

5.11. Induced triangulations

We present a few properties of induced triangulations.

Lemma 5.12. Suppose that \mathcal{L} is a loom space. Let $\mathcal{V} = V(\mathcal{L})$ be its induced triangulation. Then the realisation $|\mathcal{V}|$ is a non-compact, connected three-manifold. Furthermore, $|\mathcal{V}|$ is orientable.

Proof. By Corollary 4.25, the triangulation $V(\mathcal{L})$ has infinitely many model tetrahedra. We deduce that $|\mathcal{V}|$ is non-compact. By Proposition 4.24, we have that $|\mathcal{V}|$ is connected. By Lemma 2.31, we have that every face of $|\mathcal{V}|$ meets exactly two tetrahedra of $|\mathcal{V}|$. By Corollary 4.18, we have that every edge of $|\mathcal{V}|$ meets finitely many faces, and thus finitely many tetrahedra, of $|\mathcal{V}|$. Since tetrahedron rectangles are embedded in \mathcal{L} , no tetrahedron is glued to itself. Recall also that we removed the zero-skeleton from $|\mathcal{V}|$. We deduce that $|\mathcal{V}|$ is a non-compact, connected topological space which is a three-manifold away from the midpoints of edges.

As in Remark 2.5, we fix orientations of $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$. Ordering $F^{\mathcal{L}}$ before $F_{\mathcal{L}}$, these orientations determine an orientation of \mathcal{L} as well as the cardinal directions south, east, north, and west.

Suppose that *P* is a tetrahedron rectangle in \mathcal{L} . We order the sides of *P* according to their direction: first south, then east, north, and west. This induces an ordering on the model vertices of c(P) and thus induces an orientation on c(P). Suppose that *Q* is another tetrahedron rectangle which is face-adjacent to *P*. Suppose that $R = P \cap Q$ is the shared face rectangle. We note that the composition $\phi_{R,P} \circ \phi_{R,Q}^{-1}$ reverses orientation. See Figure 5.13. Thus the midpoints of edges are also manifold points of $|\mathcal{V}|$; also our choices of orientations above determine an orientation of $|\mathcal{V}|$.



Figure 5.13. The induced orientations on the shared face are opposite. (a) The face rectangle R is shaded; (b) the face f_R is shaded.

Before discussing the induced taut structure on $V(\mathcal{L})$, we require a lemma.

Lemma 5.14. Suppose that P is an edge rectangle of \mathcal{L} . Then there are exactly two tetrahedron rectangles which contain P and properly span P.

Proof. Let a and b be the cusps of P. Let m_a , ℓ^a , m_b , and ℓ^b be the sides of P, contained in the associated cusp leaves. Now let R_P be the union of all rectangles that contain



Figure 5.15. The small rectangular neighbourhoods of m_a and m_b are shaded.

both m_a and m_b . Note that every rectangle in this union properly south-north spans P. Appealing to Definition 2.11 (1) twice and Lemma 2.4, the union R_P is non-empty. See Figure 5.15.

Since R_P can be realised as an increasing union of rectangles, it is a rectangle. Furthermore, R_P properly south-north spans P. Finally, since R_P is maximal, it is a tetrahedron rectangle, by Definition 2.11 (2). The same construction, applied to ℓ^a and ℓ^b , produces the tetrahedron rectangle R^P .

Suppose that Q is a tetrahedron rectangle containing P, and either south-north or west-east spanning P. Thus Q is either contained in R_P or in R^P , and we are done.

We now define the *induced* dihedral angle assignment on $\mathcal{V} = V(\mathcal{L})$. Suppose that R is a tetrahedron rectangle in \mathcal{L} ; let c(R) be the corresponding model tetrahedron. By Lemma 2.30, there are six edge rectangles in R. By Definition 2.9, exactly two of these span R. We give c(R) a taut structure as follows. Suppose that P is an edge rectangle contained in R. Thus $\phi_{P,R}$ gives a model edge e of c(R). We give e dihedral angle π or zero exactly as P does or does not span R.

From Lemma 5.14, we deduce the following.

Corollary 5.16. The induced dihedral angle assignment on $V(\mathcal{L})$ is a taut structure.

We now define the *induced* colouring of the one-skeleton of $V(\mathcal{L})$. Suppose that *P* is an edge rectangle, coloured according to Definition 2.28. Then we give the edge c(P) that same colour.

Consulting Figure 5.4, we can state the following.

Corollary 5.17. Orienting the foliations $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ induces a veering structure on $V(\mathcal{L})$. Thus $V(\mathcal{L})$ is locally veering.

5.18. Functorial

We summarise this section as follows.

Proposition 5.19. *Guéritaud's construction* V: $Loom(\mathbb{R}^2) \rightarrow Veer$ *is a functor from the category of loom spaces to the category of locally veering triangulations.*

Proof. Suppose that \mathcal{L} is a loom space. By Lemma 5.12, its induced triangulation $V(\mathcal{L})$ is an ideal triangulation of a non-compact, connected, orientable three-manifold. By Corol-

lary 5.16, the induced dihedral angle makes $V(\mathcal{L})$ into a taut triangulation. By Corollary 5.17, we have that $V(\mathcal{L})$ is locally veering.

Suppose now that \mathcal{M} and \mathcal{N} are also loom spaces. Suppose that $f: \mathcal{L} \to \mathcal{M}$ and $g: \mathcal{M} \to \mathcal{N}$ are loom isomorphisms. Recall that we use the notations $V_f = V(f)$ to represent the induced map.

Suppose that $P \subset \mathcal{L}$ is a skeletal rectangle. If $f = \text{Id}_{\mathcal{L}}$, then, appealing to Definition 5.10, we have

$$V_f(c(P)) = c(f(P)) = c(P).$$

Thus $V(Id_{\mathcal{L}}) = Id_{V(\mathcal{L})}$, as desired.

For general loom isomorphisms f and g, and again appealing to Definition 5.10, we have

$$\mathsf{V}_{g\circ f}(\mathsf{c}(P)) = \mathsf{c}(g(f(P))) = \mathsf{V}_g(\mathsf{c}(f(P))) = (\mathsf{V}_g \circ \mathsf{V}_f)(\mathsf{c}(P)).$$

Thus $V(g \circ f) = V(g) \circ V(f)$.

Note that $f^{-1}: \mathcal{M} \to \mathcal{X}$ is a loom isomorphism; thus $V(f^{-1})$ is an inverse for V(f). We deduce that V(f) is a bijection. Note that the triangulations $V(\mathcal{X})$ and $V(\mathcal{M})$ are defined solely in terms of the combinatorics of \mathcal{X} and \mathcal{M} . Thus the cell structures, angle assignments, and local veering structures are also preserved by V(f).

We deduce that V: Aut(\mathcal{L}) \rightarrow Aut(V(\mathcal{L})) is a group homomorphism. In fact, as discussed in Section 1.3, it is an isomorphism.

6. Convexity

In this section, we prove Theorem 6.50: the realisation of the triangulation $V(\mathcal{L})$ is homeomorphic to \mathbb{R}^3 .

Remark 6.1. One proof of Theorem 6.50 runs along the following lines. Choose transverse measures of full support for $F^{\mathcal{X}}$ and $F_{\mathcal{X}}$. This gives \mathcal{L} an incomplete, locally euclidean, metric. Suppose that $t \in V(\mathcal{X})$ is a model tetrahedron. By sending its vertices to the associated cusps (in the completed metric), and then extending to all of t via barycentric coordinates, we obtain a linear map from t to \mathcal{L} . We glue these linear maps together to obtain a piecewise-linear map $\pi: |V(\mathcal{X})| \to \mathcal{X}$. We now claim the following:

- The map π is continuous and surjective.
- Point preimages under π are copies of \mathbb{R} .
- The map π is a fibre bundle map.

The proof of surjectivity requires the astroid lemma (Lemma 4.10(1)). Since \mathcal{L} is homeomorphic to \mathbb{R}^2 , we deduce that $|V(\mathcal{L})|$ is isomorphic to a product. Thus the claims imply Theorem 6.50.

In their work, Landry, Minsky, and Taylor [22, Proposition 5.11] carry this strategy out but in a more delicate setting. They begin with a pseudo-Anosov flow without perfect fits. After drilling, the leaf space is a loom space, equipped with an action by the fundamental group of the drilled three-manifold. Here one cannot simply choose measures since, in general, measures invariant under the action need not exist.

In the remainder of Section 6, we give a very different proof. We develop synthetic notions of geodesicity and convexity in \mathcal{L} . We use these to prove that $|V(\mathcal{L})|$ admits an exhaustion by three-balls. The synthetic approach is longer than that of Remark 6.1; however it is constructive and is more revealing of the structure of \mathcal{L} . In future work, we use this combinatorial structure to give an algorithm for drilling veering triangulations along flow loops.

6.2. Geodesics and convexity

We define a *polygonal geodesic* in \mathcal{L} . These are *polygonal paths* (as in [12, Definition 3.1]) with additional properties.

Our geodesics are very similar to "staircase" geodesics in \mathbb{R}^2 when it is equipped with the L^1 metric. However, our definition is combinatorial; we do not make use of a metric on \mathcal{L} .

Definition 6.3. A segment in \mathcal{L} is a subarc of a leaf of $F^{\mathcal{L}}$ or $F_{\mathcal{L}}$. Suppose that *a* and *b* lie in $\mathcal{L} \cup \Delta(\mathcal{L})$. A polygonal path γ from *a* to *b* is

- a finite sequence $(\rho_i)_{i=0}^{n-1}$ of oriented segments and
- a finite sequence $(v_i)_{i=0}^n$ of material points and cusps

where

- $a = v_0, v_n = b$, and
- the arc ρ_i emanates from v_i and ends at v_{i+1} .

We say that v_i is a *corner* of γ if there is a rectangle *R* with consecutive sides contained in ρ_{i-1} and ρ_i , respectively.

A corner v_i is *right-turning* if R is locally to the right of γ . We define *left-turning* corners similarly.

Suppose that γ is a polygonal path in \mathcal{L} . Then the cusps of γ (that is, the elements of $\Delta(\gamma)$) lie among the endpoints v_i of the segments ρ_i . In particular, the cusps of γ need not appear only at the corners of γ .

Definition 6.4. We say that a polygonal path γ is a *geodesic* if the intersection of $\bigcup_i \rho_i$ with any leaf is connected.

Definition 6.5. Suppose that *H* is a subset of *L*. Then *H* is *convex* if for any $p, q \in H \cup \Delta(H)$, we have that *H* contains all geodesics between *p* and *q*.

6.6. U-turns

We will need a criterion for a path being a geodesic. We begin as follows.

Definition 6.7. We say that a rectangle *R* is a *U*-turn for γ if

- three sides of *R* are contained in γ and
- $R \cap \gamma$ is empty.

We say that a U-turn for γ is *maximal* if it is not properly contained in another U-turn for γ . See Figure 6.8 for several examples of maximal U-turns.



Figure 6.8. In each case, the maximal U-turn is the shaded rectangle. Any combinatorial possibility may be obtained by combining these, and adding cusps to γ .

An Euler characteristic argument gives the following.

Lemma 6.9. For any embedded polygonal loop γ in \mathcal{L} and for any segment σ of γ , there is a U-turn of γ disjoint from σ .

We use this to prove the following.

Lemma 6.10. An embedded polygonal path δ is a geodesic if and only it contains no *U*-turns.

Proof. If *R* is a U-turn for δ , then some leaf ℓ crossing the interior of *R* intersects δ twice.

For the other direction, suppose that ℓ is a leaf that intersects δ twice. Thus there is a segment $\sigma \subset \ell$ so that

- $\sigma \cap \delta = \partial \sigma$ and
- $\sigma \cup \delta$ contains an embedded polygonal loop γ .

We apply Lemma 6.9 to γ and find a U-turn in δ .

Lemma 6.11. Suppose that a and b lie in $\mathcal{L} \cup \Delta(\mathcal{L})$. Then there is a geodesic δ from a to b.

Proof. Using Lemma 2.4 or Definition 2.11 (1), we choose rectangles *A* and *B* that have *a* and *b*, respectively, in the interior of one of their sides. Since \mathcal{L} is path connected, there is a compact arc ε connecting a point of *A* to a point of *B*. We cover ε by finitely many rectangles. The union of these rectangles with *A* and *B* contains an embedded polygonal path γ from *a* to *b*.

By Lemma 6.10, it now suffices to "straighten" γ so that it contains no U-turns. To this end, we define the *complexity* of γ to be the pair

ordered lexicographically.

We now induct on the complexity. Suppose that γ is a polygonal path from *a* and *b*. Suppose that *R* is a maximal U-turn for γ . Define

$$\gamma' = (\gamma - \partial R) \cup (\partial R - \gamma).$$

(Again, see Figure 6.8.) If γ' is empty, then a = b and the desired geodesic path has no arcs. If not, then γ' is a disjoint union of a polygonal path from *a* to *b* and some number of polygonal loops. Let γ'' be the polygonal path, and let $\{\gamma_i\}$ be the polygonal loops.

We now claim that γ'' has lower complexity than γ . If γ'' has fewer corners than γ , we are done. If γ'' has the same number of corners as γ , then $\{\gamma_i\}$ is empty. Let σ be the side of the maximal U-turn *R* not contained in γ . We deduce that the interior of σ is disjoint from γ . We further deduce that σ contains a cusp in its interior; see Figure 6.8 (d). Thus γ'' has fewer maximal U-turns than γ .

6.12. Sectors

Properly speaking, cusps of \mathcal{L} are not points of \mathcal{L} . Nonetheless, we would like to treat points and cusps on an equal footing. We do so as follows.

Definition 6.13. Suppose that x is a point or a cusp of \mathcal{L} . Suppose that Q is a subset of \mathcal{L} . We say that Q meets x if either x is a point lying in Q or x is a cusp of Q.

The following is similar to Fenley's definition of quarters [12, p. 22].

Definition 6.14. Suppose that x is a point or cusp of \mathcal{L} . Let $\Lambda(x)$ be the union of all leaves meeting x. We call the components of $\mathcal{L} - \Lambda(x)$ the *sectors based at x*. Two sectors based at x are *adjacent* if they are disjoint but their closures intersect along a subarc of some leaf.

Note that a (cusp) rectangle R, with an (ideal) corner at x, determines a unique sector based at x, denoted S(x, R). Note also that the sector S(x, R) contains the staircase $\Gamma(x, R)$.

Lemma 6.15. If x is a point of \mathcal{L} , then there are exactly four sectors based at x. If x is a cusp of \mathcal{L} , then there are countably many sectors based at x; these are linearly ordered by the adjacency relation.

Proof. The first follows from Lemma 2.4 and Remark 2.2. The second follows from Definitions 3.3 and 2.11 (1), and Remark 3.9.

Definition 6.16. Suppose that p, q, and r lie in $\mathcal{L} \cup \Delta(\mathcal{L})$. Suppose that Q and R are sectors based at p so that q and r lie in the closures of Q and R, respectively. If the sectors Q and R are distinct and not adjacent, then we say that p is *between* q and r.

Lemma 6.17. Suppose that p, q, and r lie in $\mathcal{L} \cup \Delta(\mathcal{L})$. Then p is between q and r if and only if there is a geodesic from q to r passing through p.

Proof. Suppose that p is between q and r. Lemma 6.11 gives us geodesics δ from q to p and ε from p to r. Since δ is a geodesic, and since it meets p, it must be contained in the closure of a single sector at p. The same holds for ε . Thus δ and ε are contained in the closures of distinct, non-adjacent sectors based at p. Thus $\delta \cup \varepsilon$ is the desired geodesic, giving the forward direction.

The backward direction follows from the definition of geodesics.

6.18. Hulls

Here we again follow, at least in spirit, Fenley's paper [12, Section 3]; see in particular his notion of *convex polygonal paths* [12, Definition 3.2].

Definition 6.19. Suppose that *C* is a finite subset of $\mathcal{L} \cup \Delta(\mathcal{L})$. We define the *hull* of *C* as follows:

 $H(C) = \{ p \in \mathcal{L} \mid p \text{ is between some pair of elements of } C \}.$

From Lemma 6.17, we deduce the following.

Corollary 6.20. Suppose that C is a finite subset of $\mathcal{L} \cup \Delta(\mathcal{L})$. Then H(C) is the union of all geodesics connecting elements of C.

Set $H(q, r) = H(\{q, r\})$. Note that H(q, r) is closed in \mathcal{L} . The definition of hull implies that $H(C) = \bigcup_{q,r \in C} H(q, r)$. Thus, since H(C) is a finite union of closed sets, it is also closed. Recall from Definition 3.4 that $\Delta(H(C))$ is the set of cusps of the hull H(C).

Lemma 6.21. Suppose that C is a finite subset of $\mathcal{L} \cup \Delta(\mathcal{L})$. Suppose that q and r lie in $H(C) \cup \Delta(H(C))$. Then $H(q,r) \subset H(C)$.

It follows that hulls are convex. Equally well, the boundary of a hull is the union of finitely many *convex polygonal paths* in the sense of [12, Definition 3.2]. We prove a version of this in Lemma 6.23 (3).

Proof of Lemma 6.21. Set H = H(C). We prove the contrapositive. That is, we assume the following:

- *C* is a finite subset of $\mathcal{L} \cup \Delta(\mathcal{L})$,
- q and r lie in $\mathcal{L} \cup \Delta(\mathcal{L})$, and
- some point $p \in H(q, r) \subset \mathcal{L}$ does not lie in H.

We must show that one of q or r lies outside of $H \cup \Delta(H)$.

As in Remark 2.5, we orient $F^{\mathcal{L}}$ and $F_{\mathcal{L}}$ so that we may refer to the cardinal directions in \mathcal{L} . By Definition 6.16, the point p is between q and r. Let ℓ^p and m_p be the leaves through p. There are four sectors based at p; breaking symmetry, we assume that q lies in the closure of the south-west sector while r lies in the closure of the north-east sector. We have assumed that p is not between any pair of points of C. Breaking symmetry, we may assume that C lies strictly to the east of ℓ^p . Since C is finite, and appealing to either Corollary 4.13 or Corollary 4.26, there is a leaf ℓ of $F^{\mathcal{L}}$ that separates ℓ^p from C. Appealing to Definition 6.16, we deduce that H is east of ℓ .

Suppose that q is a point of \mathcal{L} . Then q is on, or west of, ℓ^p . Thus q is strictly west of ℓ . Thus q is not between any pair of points of C, and we are done.

Suppose instead that q lies in $\Delta(\mathcal{L})$. Then no cusp rectangle at q lies in H. Thus q is not an element of $\Delta(H)$, and we are done.

Suppose that q and r lie in $\mathcal{L} \cup \Delta(\mathcal{L})$. Suppose that S(q) and S(r) are sectors, based at q and r, respectively, so that r lies in S(q) and q lies in S(r). (This is possible if and only if q and r do not lie on a common leaf.) Set $S(q, r) = S(q) \cap S(r)$. Note that the closure of S(q, r) contains all geodesics from q to r. Thus by Corollary 6.20, the closure contains H(q, r).

We fix, for the remainder of Section 6, an orientation on \mathcal{L} . Suppose that δ is a geodesic from q to r. We say that x, a point or cusp of H(q, r), is to the right of δ if x meets the union of the components of $S(q, r) - \delta$ to the right of δ . We say that a geodesic ε is to the right of δ if all points of $\varepsilon - \delta$ are to the right of δ . We define to the left of similarly. See Figure 6.22.



Figure 6.22. Two geodesics δ and ε from q to r. There is a right-turning corner (on δ) at x.

Lemma 6.23. Suppose that q and r lie in $\mathcal{L} \cup \Delta(\mathcal{L})$. Then we have the following:

- (1) $\Delta(H(q, r))$ is finite.
- (2) For every cusp $c \in \Delta(H(q, r))$, there is a geodesic from q to r running through c.
- (3) H(q, r) contains a unique rightmost geodesic and also a unique leftmost geodesic. The rightmost geodesic has no right-turning corners; the leftmost geodesic has no left-turning corners.
- (4) If q and r are cusps, then the boundary of H(q, r) is the disjoint union of the rightmost and leftmost geodesics.
- (5) If q and r are cusps, then H(q, r) is a disjoint finite union of finite-sided closed disks (minus finitely many boundary points). These disks may have cusps in common.

Proof. Lemma 6.11 gives us a geodesic γ from q to r. Suppose that x is a right-turning corner of γ . Let R be the rectangle provided by Definition 6.3. Let ℓ^x and m_x be the axis rays of the staircase $\Gamma(x, R)$. Note that both ℓ^x and m_x have initial segments contained in γ .

Define m_x^{γ} to be the *projection* of γ , to m_x , along $F^{\mathcal{L}}$. That is, suppose for $y \in \mathcal{L}$, we have that ℓ^y is the leaf of $F^{\mathcal{L}}$ through y. Then we define $m_x^{\gamma} \subset m_x$ to be the polygonal path obtained by taking the closure of

$$\{m_x \cap \ell^y \mid y \in \gamma\}.$$

We define $\ell_{\gamma}^x \subset \ell^x$ similarly. See Figure 6.24.



Figure 6.24. The two projections of γ to the axis rays of a right-turning corner x.

Note that since γ is a finite polygonal path, both m_x^{γ} and ℓ_{γ}^x are bounded in m_x and ℓ^x , respectively. (However, by Lemma 4.6, the projection m_x^{γ} may meet the axis cusp c_m for m_x , if it exists. A similar statement holds for ℓ_{γ}^x .) Recall that $\Delta(\Gamma(x))$ is the set of exterior cusps for $\Gamma(x)$. We define $\Delta^{\gamma}(\Gamma(x))$ to be those cusps of $\Delta(\Gamma(x))$ whose projections to m_x and ℓ^x lie in m_x^{γ} and ℓ_{γ}^x , respectively. By the astroid lemma (Lemma 4.10(1)), we have that $\Delta^{\gamma}(\Gamma(x))$ is finite.

Suppose that c is any cusp of $\Delta(H(q, r))$, not on γ . Suppose that c is to the right of γ , that is, c lies in one of the components of $S(q, r) - \gamma$ to the right of γ . Since c is a cusp of $\Delta(H(q, r))$, it lies between q and r. Thus there are distinct, non-adjacent sectors U and V based at c whose closures contain q and r, respectively. By Lemma 6.15, there is a sector S based at c separating U from V. Let m_c and ℓ^c be the cusp leaves giving the sides of S.

By construction, neither q nor r lie in S. (They may lie in the closure of S; that is, in m_c or in ℓ^c .) Thus γ intersects both m_c and ℓ^c . Let S_{γ} be the component of $S - \gamma$ that meets c. We define $\gamma' = \gamma \cap \partial S_{\gamma}$. Note that ∂S_{γ} is a closed polygonal loop.

An Euler characteristic argument, applied to ∂S_{γ} , gives us a right-turning corner x of γ' . We deduce that ℓ^c crosses m_x^{γ} ; likewise, m_c crosses ℓ_{γ}^x . Thus c lies in $\Delta^{\gamma}(\Gamma(x))$. It follows that every cusp of $\Delta(H(q, r))$ either lies on γ or lies in $\bigcup_x \Delta^{\gamma}(\Gamma(x))$. Here x ranges over the corners of γ . Thus $\Delta(H(q, r))$ is a finite union of finite sets. This proves (1).

Suppose now that *c* is a cusp of $\Delta(H(q, r))$. Pick any geodesic γ from *q* to *r*. If *c* lies in γ , we are done. If not, then we may assume that *c* is to the right of γ . Then, as above, we find *S* and *S*_{γ}. Using these, we define

$$\varepsilon = (\gamma - \partial S_{\gamma}) \cup (\partial S_{\gamma} - \gamma).$$

This is a geodesic through c, and so gives (2).

Starting with any geodesic from q to r, (2) finds a geodesic to its right. Repeating this, we find a sequence of geodesics; by (1) this sequence ends with a geodesic ζ which passes through all of the rightmost cusps. If ζ has a right-turning corner, say at x, then the cusps immediately before and after x span an edge rectangle. Flipping over such edge rectangles, we obtain a new geodesic ρ . Since ρ has no right-turning corners, ρ is rightmost. This proves (3).

Suppose that ρ and λ are, respectively, the rightmost and leftmost geodesics in H(q, r). Suppose for a contradiction that ρ and λ intersect at a point p of \mathcal{L} . If ρ and λ cross at p then either ρ is not rightmost or λ is not leftmost. Either is a contradiction. We deduce that near p both ρ and λ lie in a single leaf, say m_p . Thus there is a cusp at the beginning of the segment $\rho \cap m_p$. Likewise there is a cusp at the end of the segment $\lambda \cap m_p$. These contradict Lemma 3.10. This proves (4).

Appealing to the Jordan curve theorem and (4), we obtain (5).

Corollary 6.25. Staircases are convex.

Proof. Suppose that $\Gamma(x)$ is the given staircase. By Lemma 4.10, the set of exterior cusps is countable. Let H_k be the hull of x together with the first k exterior cusps. By Corollary 4.14 and Lemma 6.23 (3), the boundary of H_k is contained in the closure of $\Gamma(x)$. Thus the staircase is a growing union of convex sets.

Lemmas 6.23(5) and 6.21 prove the following.

Corollary 6.26. Suppose that C is a finite subset of $\Delta(\mathcal{L})$. Then H(C) is a disjoint finite union of finite-sided disks (which may have cusps in common).

6.27. Skeletal rectangles, redux

With convexity in hand, we are equipped to prove various existence and uniqueness results.

Lemma 6.28. Suppose that three distinct edge rectangles meet three distinct cusps. Then they are all contained in a single face rectangle.

Proof. Let a, b, and c be the given cusps. Let A, B, and C be the given edge rectangles; by Lemma 3.11, we may assume that A does not meet a, and so on.

Let *H* be the hull of *a*, *b*, and *c*. By definition, the union of the closures of *A*, *B*, and *C* gives *H*. By Lemma 6.33 (3), the hull *H* is a disjoint finite union of finite sided disks, adjacent only at cusps of $\Delta(H)$. Suppose that there are two or more disks in the disjoint union. Suppose that *c*, say, meets two of these disks. Then all geodesics from *a* to *b* run through *c*. In particular, all geodesics from *a* to *b* in the edge rectangle *C* meet *c*. Thus *C* meets *c*, a contradiction. Thus *H* is a single disk.

Breaking symmetry, we may assume that a, b, and c appear in anticlockwise order around ∂H . Let ρ_a be the rightmost geodesic from b to c given by Lemma 6.23 (3). Define ρ_b and ρ_c similarly. Since a is to the left of ρ_a , we deduce that ρ_a lies in ∂H . Similarly, ρ_b and ρ_c lie in ∂H .

Since b and c are the cusps of A, we deduce that ρ_a also lies in ∂A . So ρ_a , and similarly ρ_b and ρ_c , each consist of only two segments. We deduce that the only cusps meeting ∂H are a, b, and c. Also the only segments in ∂H are those of ρ_a , ρ_b , and ρ_c .

Enumerating simple polygonal loops with at most six segments, which occur as the boundary of a convex set, we find two possibilities. One has four outward corners (one at a cusp) while the other has five (two at cusps). These are shown in Figure 6.29. The latter does not contain an edge rectangle between two of its cusps. The former is the desired face rectangle.



Figure 6.29. (a) One of the four possible face rectangles. Note that all four corners are outward. (b) One of the four possible convex hulls with six segments in its boundary and five outward corners.

Suppose that *P*, *Q*, and *R* are edge, face, and tetrahedron rectangles in \mathcal{L} , respectively. We say that c(P) is the *bottom edge of* c(Q) if *P* and *Q* west-east span each other. We similarly define what it means for c(P) to be the *bottom edge of* c(R) and what it means for c(Q) to be a *bottom face of* c(R) (there are two). We define the *top* edges and faces similarly. See Figure 2.8 (d); the shaded edge rectangle corresponds to the bottom edge of the corresponding faces and tetrahedron. Also, if *P* is an edge rectangle, then Lemma 5.14 provides two tetrahedron rectangles: R_P and R^P . These are the tetrahedron rectangles, where c(P) is the bottom and top edge of $c(R_P)$ and $c(R^P)$, respectively.

Lemma 6.30. Suppose that P and Q are edge rectangles of \mathcal{L} . Suppose that P properly west-east spans Q. Set $H = H(\Delta(P) \cup \Delta(Q))$. Then R_P and R^Q lie in H.

Proof. Suppose that *a* and *b* are the cusps of *Q*. Breaking symmetry, suppose that *a* is south of and *b* is north of *P*. Let ℓ be the cusp leaf emanating from *a* and crossing *P*. Thus either *a* is the southern cusp of R_P or ℓ crosses the southern side of R_P . The same holds for *b*, to the north of *P*. Thus R_P lies in *H*. A similar argument proves that R^Q lies in *H*.

6.31. Three-balls

For the remainder of Section 6, we fix *T* a non-empty, finite, and face-connected collection of tetrahedra in $V(\mathcal{L})$. Note that the union of the corresponding tetrahedron rectangles need not be convex. In a slight abuse of notation, we define

$$\Delta(T) = \bigcup_{t \in T} \Delta(\mathsf{c}^{-1}(t)).$$

For the remainder of Section 6, we set $H = H(\Delta(T))$: that is, H is the hull of the set of all cusps of the rectangles associated to tetrahedra $t \in T$.

Definition 6.32. The *content* of *H*, denoted by V(H), is the set of cells *k* of $V(\mathcal{L})$ so that the skeletal rectangle $c^{-1}(k)$ lies in *H*.

We next build up the combinatorial structure of V(H) in a sequence of lemmas and corollaries. The last of these is Proposition 6.46, proving that V(H) is a three-ball. From this we deduce Theorem 6.50.

Lemma 6.33. We have the following:

- (1) $T \subset V(H)$.
- (2) V(H) is finite.
- (3) The hull H is a finite-sided closed disk in \mathcal{L} (minus finitely many boundary points).

Proof. A tetrahedron rectangle is the (interior of the) convex hull of its cusps. This gives the first conclusion (1).

By Corollary 6.26, we know that *H* is a finite disjoint union of finite-sided disks. Thus $\Delta(H)$ is finite. Thus V(H) is finite, giving (2).

Suppose that there is more than one disk component in the union given by Corollary 6.26. We form a bipartite graph as follows: we take one set of nodes for the disk components, another set of nodes for the cusps meeting two or more disks, and edges for a disk-cusp pair where the latter is a cusp of the former. By Lemma 6.21, this graph is a tree. Any leaf of the tree gives a disk that contains at least one tetrahedron rectangle $c^{-1}(t)$ for some $t \in T$. Since there are at least two disks, the tree has at least two leaves. We deduce that T is not face-connected, a contradiction. This gives (3).

Remark 6.34. Recall that the loom space \mathcal{L} is homeomorphic to \mathbb{R}^2 . We equip the latter with its usual (anti-clockwise) orientation, and transport this to \mathcal{L} . Recall that $H = H(\Delta(T))$. The orientation of \mathcal{L} restricts to give an orientation of H.

Lemma 6.33 (3) says that *H* is homeomorphic to a closed disk minus finitely many boundary points. The boundary of *H* is a collection of open intervals, each meeting exactly two cusps. The orientation of *H* induces a circular ordering of these intervals and thus on $\Delta(H)$.

Lemma 6.35. Suppose that a and b are distinct cusps of $\Delta(H)$; suppose that a and b span an edge e of V(H). Then exactly one of the following occurs:

- a and b are adjacent in the circular order on $\Delta(H)$ or
- there is a tetrahedron $t \in T$ so that $c^{-1}(t)$ properly spans $c^{-1}(e)$.

Proof. Set $P = c^{-1}(e)$. We orient the geodesics of P from a to b. Let λ and ρ be the leftmost and rightmost of these; note that each is connected.

From Lemma 6.33 (3), the fact that $\partial P = \lambda \cup \rho$ has exactly two components, and the fact that *P* is open, we deduce that H - P has exactly two components. We label these *K* and *K'*, with *K* to the left of *P* and *K'* to the right. Every cusp of *H* is also a cusp of *K* or of *K'*; only *a* and *b* are cusps of both. Note that $K \cup P$ is convex because it is obtained by cutting *H* (a convex set by Lemma 6.21) along ρ , a rightmost geodesic. Similarly, $K' \cup P$ is convex.

Suppose that *a* and *b* are not adjacent in the circular order on $\Delta(H)$. Thus, each of $\Delta(K)$ and $\Delta(K')$ contains at least one cusp of $\Delta(H) - \{a, b\}$. Since $K' \cup P$ is convex, the set $\Delta(K)$ contains at least one cusp from $\Delta(T) - \{a, b\}$, say *c*. Similarly, $\Delta(K')$ contains a cusp *c'* belonging to $\Delta(T) - \{a, b\}$.

Suppose that *t* and *t'* are tetrahedra of *T* so that *c* and *c'* are cusps of $R = c^{-1}(t)$ and $R' = c^{-1}(t')$, respectively. Since *T* is face-connected, there is a sequence of tetrahedra $(t = t_0, t_1, t_2, ..., t_n = t')$ so that t_i and t_{i+1} share a face. Let $R_k = c^{-1}(t_k)$. Note that $R' = R_n$ is not contained in $K \cup P$, since it meets *c'*. Let *k* be the first index so that R_k intersects *K'*. By induction, for *i* between 0 and *k*, we have that R_i intersects *K*. Thus R_k intersects *K* and *K'*; so R_k properly spans *P*; thus t_k is the desired tetrahedron of *T*.

We now give a partial converse to Lemma 6.35.

Lemma 6.36. Suppose that a and b are adjacent cusps in the circular order on $\Delta(H)$. Then the interior of H(a, b) is an edge rectangle in \mathcal{L} .

We call the resulting edge, e = c(H(a, b)), a *coastal edge* of V(H).

Proof of Lemma 6.36. The set H(a, b) lies in the hull H by Lemma 6.21.

Breaking symmetry, suppose that b is anticlockwise of a in ∂H . Applying Lemma 6.23 (3), let ρ be the rightmost geodesic from a to b. Thus ρ lies in ∂H . Recall also that ρ has no right-turning corners. By Lemma 3.10, the geodesic ρ has at least two segments. We deduce that ρ has exactly two segments; these meet at one left-turning corner, say at $x \in \mathcal{L} \cup \Delta(\mathcal{L})$. However, x is not a cusp by Lemma 3.10. Let ℓ^x and m_x be the leaves through x. Breaking symmetry, we suppose that ρ first runs along m_x and then along ℓ^x . See Figure 6.37. Let $\Gamma(x)$ be the staircase at x that meets both a and b.



Figure 6.37. Part of the staircase $\Gamma(x)$ meeting both *a* and *b*.

Let λ be the leftmost geodesic from *a* to *b*. By Corollary 6.25, the staircase $\Gamma(x)$ contains λ . If λ has exactly two segments then the lemma is proved. For a contradiction, suppose that λ has more than two segments. By Lemma 6.23 (3), there is at least one cusp along λ . Let *c* be the first such. Thus H(*a*, *c*) is the closure of an edge rectangle.

If a and c are adjacent in the circular order on $\Delta(H)$, then the leftmost geodesic from a to c is again in ∂H . We deduce that no tetrahedron rectangle in H meets a; see Figure 6.37. In particular, a does not lie in $\Delta(T)$. Also a does not meet the hull of any pair of cusps in $\Delta(T)$. Thus a does not meet H, a contradiction.

Therefore, *a* and *c* are not adjacent. Applying Lemma 6.35, we find a tetrahedron *t* so that $c^{-1}(t)$ properly spans the interior of H(a, c). Thus $c^{-1}(t)$ crosses one of ℓ^x or m_x . Either is a contradiction. Again, see Figure 6.37.

Definition 6.38. We say that *e*, an edge of V(H), is a *lower edge for H* if it has the following property. For any edge e' of V(H) the edge rectangle $c^{-1}(e')$ does not properly west-east span $c^{-1}(e)$.

We say that a face f of V(H) is a *lower face for* H if all three edges of f are lower edges for H.

The definition implies the following.

Lemma 6.39. Lower edges do not link each other with respect to the given circular order on $\Delta(H)$.

From Lemma 6.36, we deduce the following.

Corollary 6.40. Coastal edges are lower edges for H.

We now show that lower edges give rise to tetrahedra in V(H).

Corollary 6.41. Suppose that e is a non-coastal, lower edge for H. Then the tetrahedron t, having e as its bottom edge, lies in V(H).

Proof. Suppose that *e* is the given lower edge. Let $P = c^{-1}(e)$. Define U(e) to be those tetrahedra *t'* so that $R' = c^{-1}(t')$ properly south-north spans *P*. Since *e* is non-coastal, applying Lemma 6.35, the set U(e) is non-empty.

Suppose that t' is any tetrahedron in U(e). Again, let R' be the corresponding tetrahedron rectangle. Since e is lower, R' does not properly west-east span P. Thus R' instead

properly south-north spans *P*. Let e' be the upper edge of t'. Let $P' = c^{-1}(e')$. We now apply Lemma 6.30 to the edge rectangles *P* and *P'*. The resulting tetrahedron $t = c(R_P)$ gives the result.

Lemma 6.42. The lower edges and faces for H form a triangulated disk in V(H).

Proof. Let L(H) be the subcomplex of $V(\mathcal{X})$ consisting of the lower edges and faces for H. By Corollary 6.40, all coastal edges of V(H) belong to L(H). Since T is non-empty, the set $\Delta(H)$ has at least four elements. From this and Lemma 6.33 (3), we deduce that L(H) contains a simple edge cycle.

Suppose that ℓ is any simple edge cycle in L(H). By Lemmas 3.11 and 3.12, the edge cycle ℓ has length at least three. If ℓ has length exactly three, then by Lemma 6.28 and by the definition of lower faces, there is a face f contained in L(H) spanning ℓ . If ℓ has length greater than three then we must show that ℓ has a chord in L(H). This and induction then proves the lemma.

Let H_{ℓ} be the hull of the cusps of ℓ . Let e' be any edge of ℓ . Let $P' = c^{-1}(e')$ be the corresponding edge rectangle. Thus P' lies in H_{ℓ} and separates \mathcal{L} . All edges of $\ell \subset L$ are lower. Thus their edge rectangles do not properly span each other. Therefore, all cusps of $\Delta(H_{\ell}) - \Delta(P')$ meet a common component of $\mathcal{L} - P'$. Orient e' so that this component is to the left of P'. Let $\rho_{e'}$ be the rightmost geodesic in P'. Since H_{ℓ} is convex, $\rho_{e'}$ lies in the boundary of H_{ℓ} .

Claim 6.43. H_{ℓ} is a disk and $\partial H_{\ell} = \bigsqcup_{e} \rho_{e}$, where the union ranges over the edges of ℓ .

Proof. Note that ℓ is simple: that is, it enters and exits each of its cusps exactly once. Thus $\bigsqcup_e \rho_e$ meets all cusps of ℓ ; for each cusp, exactly one arc enters the neighbourhood of that cusp and exactly one arc leaves. Since $\bigsqcup_e \rho_e$ lies in ∂H_ℓ , the arcs of $\bigsqcup_e \rho_e$ do not intersect in \mathcal{X} . Thus $\bigsqcup_e \rho_e$ is a polygonal loop in \mathcal{X} . By the Jordan curve theorem, $\bigsqcup_e \rho_e$ bounds a disk D in \mathcal{X} . By Corollary 6.26, the hull H_ℓ is a finite union of finite-sided disks, meeting at cusps. Thus D is exactly one of these disks. The claim follows.

Recall that the simple edge loop ℓ has length at least four. We now find a chord of ℓ in L(H).

Applying Lemma 6.9 to ∂H_{ℓ} , we obtain an (open) rectangle $R \subset H_{\ell}$ with three consecutive sides s, s', and s'' contained in ∂H_{ℓ} and otherwise disjoint from ∂H_{ℓ} . By Claim 6.43, there is a cusp b of ℓ meeting s'. Let a and c be the (distinct) cusps of ∂H_{ℓ} immediately clockwise and anticlockwise of b, respectively. There are now two cases as a and c do or do not span an edge of V(\mathcal{L}).

First suppose that there is an edge rectangle $P = c^{-1}(e)$ with cusps at *a* and *c*. See Figure 6.44 for the three possible configurations (up to horizontal and vertical reflections).

Since ℓ has length greater than three, we deduce that *e* is not an edge of ℓ . If *e* is lower for *H*, then we have found a chord for ℓ and we are done. Suppose instead that *e* is not lower for *H*. Thus there is some edge *e'* of V(*H*) so that $P' = c^{-1}(e')$ properly west-east spans *P*. Repeating the argument if needed, we may also assume that *e'* lies



Figure 6.44. The three possible configurations for an edge rectangle with cusps at *a* and *c*.



Figure 6.45. The three possible configurations when there is no edge rectangle with cusps at a and c.

in L(H). Then P' cannot west-east span any edge rectangle of $c^{-1}(\ell)$ as these are all in L(H). The only possible configuration is shown in Figure 6.44 (a), where P' has a cusp at b. By Lemma 6.39, the edge e' does not link any edge of ℓ . Thus e' is a chord for ℓ , a contradiction.

Next suppose that there does not exist an edge rectangle with cusps at a and c. See Figure 6.45 for the three possible configurations (up to horizontal and vertical reflections).

In all cases, the rightmost geodesic from *a* to *c* consists of exactly two segments. The leftmost geodesic from *a* to *c* necessarily meets at least one cusp, say *d*. The cusp *d* is both an exterior cusp for the staircase based at *b* and is a cusp of $\Delta(H)$. Thus there is an edge rectangle $P' = c^{-1}(e')$ with cusps at *b* and *d*. In all three cases, any edge rectangle P'' that strictly west-east spans P' necessarily strictly west-east spans the edge rectangle H(a, b). However, c(H(a, b)) is an edge of ℓ , thus lower for *H*. Thus P'' is not an edge rectangle of *H*. We deduce that the edge e' is lower for *H*. Thus e' is a chord for ℓ , as desired.

Proposition 6.46. Suppose that T is a finite, face-connected, non-empty collection of tetrahedra in $V(\mathcal{L})$. Let $H = H(\Delta(T))$. Then the realisation of V(H) is a closed three-ball (minus finitely many boundary points).

Proof. Lemma 6.42 gives us a triangulated disk L(H), whose edges and faces are lower for H and whose boundary consists of the coastal edges for H. Set $L_0 = L(H)$. We now use induction to obtain a sequence of triangulated disks $L_k \subset V(H)$, with $\partial L_k = \partial L_0$.

We say that a non-boundary edge e of L_k is *flippable* if both of its adjacent faces in L_k , say f and f', have e as their bottom edge. (That is, all three of $c^{-1}(e)$, $c^{-1}(f)$, and $c^{-1}(f')$ west-east span each other. See Figure 2.8 (d).)

If L_k has a flippable edge, then let e_k be one such. Consulting Figure 2.8 (d), we see that f and f' are the bottom faces of a tetrahedron t_k in V(H). Let g and g' be the top faces of t_k . We define the result of *flipping* L_k across e_k to be the triangulated disk

$$L_{k+1} = (L_k - (f \cup f')) \cup (g \cup g').$$

On the other hand, if L_k has no flippable edge, then the induction is complete.

Claim 6.47. Suppose that e' is an edge of V(H) that is not equal to e_j , for any j < k. Then either e' lies in L_k or there is an edge e of L_k which properly west-east spans e'.

Proof. For k = 0, this follows from the definition of L_0 . Suppose by induction that the claim holds at stage k. Suppose that e' is an edge of V(H) that is not e_j for any j < k + 1. Suppose that no edge e of L_{k+1} properly west-east spans e'.

Suppose that e' lies in L_k . Since $e' \neq e_k$, we deduce that e' lies in L_{k+1} . Thus, in this case we are done.

Suppose instead that e' does not lie in L_k . By induction, there is some edge e of L_k so that $c^{-1}(e)$ properly west-east spans $c^{-1}(e')$. If $e \neq e_k$ then e lies in L_{k+1} , contrary to assumption. Thus $e = e_k$ is the only edge of L_k whose rectangle properly west-east spans $c^{-1}(e')$; we deduce that the equatorial edges of t_k do not give properly spanning rectangles. Thus e' is the top edge of t_k . Thus e' lies in L_{k+1} , and we are done.

Claim 6.48. For every tetrahedron t' of V(H), there is an n so that $t' = t_n$.

Proof. Let D(t') be the collection of tetrahedra *s* in V(H) which have $c^{-1}(s)$ west-east spanning $c^{-1}(t')$. The set D(t') is finite and partially ordered by the spanning relation. Note that *t'* lies in D(t').

Suppose that $c^{-1}(e)$ is an edge rectangle in H that west-east spans $c^{-1}(t')$. Let e' be the top edge of t'. Thus $c^{-1}(e)$ properly west-east spans $c^{-1}(e')$. By Lemma 6.30, there is a tetrahedron t so that e is the bottom edge of t and $c^{-1}(t)$ lies in the convex hull of the cusps of $c^{-1}(e)$ and $c^{-1}(e')$. By Lemma 6.21, the tetrahedron t lies in D(t').

Recall that t_k is the tetrahedron lying between the triangulated disks L_k and L_{k+1} . Define $D_k(t') = D(t') - \{t_0, t_1, \dots, t_{k-1}\}$. If $D_k(t')$ is empty, then we are done. Otherwise, suppose that $s' \in D_k(t')$ has a bottom edge which is not flippable in L_k . Thus some edge e' of s', other than the top edge of s', is not contained in L_k . By Claim 6.47, there is an edge e of L_k that properly west-east spans e'. By the previous paragraph, s' was not a minimum of $D_k(t')$. That is, all minima of the partial order on $D_k(t')$ have bottom edges which are flippable in L_k .

Any flippable edge in L_k , if not removed, remains flippable in L_{k+1} . Thus there is a k' > k so that $D_{k'}(t')$ has fewer elements than $D_k(t')$.

Claim 6.49. For every non-coastal edge e of L_0 , there is an n so that $e = e_n$.

Proof. This follows from Corollary 6.41 and Claim 6.48.

Claim 6.48 and Lemma 6.33 (2) imply that the realisation of V(H) is a finite collection of finitely triangulated three-balls, perhaps intersecting at separating edges. Claim 6.49 implies that there are no separating edges. Thus the realisation V(H) is a closed three-ball (minus finitely many boundary points). This ends the proof of Proposition 6.46.

Theorem 6.50. Suppose that \mathcal{L} is a loom space. Then the realisation of its induced triangulation $V(\mathcal{L})$ is homeomorphic to \mathbb{R}^3 .

Proof. Choose an ordering $(t_i)_{i \in \mathbb{N}}$ for the tetrahedra of V(\mathcal{L}). Applying Proposition 4.24, we arrange matters so that any initial subsequence of (t_i) is face-connected. Let H_n be the convex hull of the cusps of the first *n* tetrahedra. By Proposition 6.46, the realisation of $V(H_n)$ is a closed three-ball (minus finitely many boundary points). Taking interiors, we find that $|V(\mathcal{L})|$ is an increasing union of open three-balls. The theorem now follows from a result of Brown [5].

Acknowledgements. We thank the referee for their careful reading, which greatly improved the work. We thank Sabetta Matsumoto for sourcing the fabric shown in Figure 1.1.

Funding. The second author was supported in part by National Science Foundation grants DMS-1708239 and DMS-2203993.

References

- [1] I. Agol, Ideal triangulations of pseudo-Anosov mapping tori. In Topology and geometry in dimension three, pp. 1–17, Contemp. Math. 560, American Mathematical Society, Providence, RI, 2011 Zbl 1335.57026 MR 2866919
- [2] I. Agol, Veering triangulations and pseudo-Anosov flows. 2015, http://www.ias.edu/math/ events/princeton-university-mathematics-department-colloquium-4, visited on 12 February 2024
- [3] C. Bonatti and A. Wilkinson, Transitive partially hyperbolic diffeomorphisms on 3-manifolds. Topology 44 (2005), no. 3, 475-508 Zbl 1074.57015 MR 2122214
- [4] R. Bowen and P. Walters, Expansive one-parameter flows. J. Differential Equations 12 (1972), 180-193 Zbl 0242.54041 MR 341451
- [5] M. Brown, The monotone union of open *n*-cells is an open *n*-cell. Proc. Amer. Math. Soc. 12 (1961), 812-814 Zbl 0103.39305 MR 126835
- [6] M. Brunella, Surfaces of section for expansive flows on three-manifolds. J. Math. Soc. Japan 47 (1995), no. 3, 491–501 Zbl 0828.58024 MR 1331326
- [7] D. Calegari, Foliations and the geometry of 3-manifolds. Oxford Math. Monogr., Oxford University Press, Oxford, 2007 Zbl 1118.57002 MR 2327361

- [8] V. Delecroix and C. Ulcigrai, Diagonal changes for surfaces in hyperelliptic components: a geometric natural extension of Ferenczi–Zamboni moves. *Geom. Dedicata* 176 (2015), 117– 174 Zbl 1326.37023 MR 3347575
- [9] S. Fenley, The structure of branching in Anosov flows of 3-manifolds. *Comment. Math. Helv.* 73 (1998), no. 2, 259–297 Zbl 0999.37008 MR 1611703
- [10] S. Fenley, Surfaces transverse to pseudo-Anosov flows and virtual fibers in 3-manifolds. *Topology* 38 (1999), no. 4, 823–859 Zbl 0926.57009 MR 1679801
- S. Fenley, Pseudo-Anosov flows and incompressible tori. Geom. Dedicata 99 (2003), 61–102 Zbl 1047.37017 MR 1998929
- [12] S. Fenley, Ideal boundaries of pseudo-Anosov flows and uniform convergence groups with connections and applications to large scale geometry. *Geom. Topol.* 16 (2012), no. 1, 1–110 Zbl 1279.37026 MR 2872578
- [13] S. Fenley and L. Mosher, Quasigeodesic flows in hyperbolic 3-manifolds. *Topology* 40 (2001), no. 3, 503–537 Zbl 0990.53040 MR 1838993
- [14] S. Frankel, S. Schleimer, and H. Segerman, From veering triangulations to link spaces and back again. 2022, arXiv:1911.00006
- [15] A. Giannopoulos, S. Schleimer, and H. Segerman, A census of veering structures. 2023, https:// math.okstate.edu/people/segerman/veering.html, visited on 12 February 2024
- [16] F. Guéritaud, Veering triangulations and Cannon–Thurston maps. J. Topol. 9 (2016), no. 3, 957–983 Zbl 1354.57025 MR 3551845
- [17] G. Hector and U. Hirsch, Introduction to the geometry of foliations. Part A. Foliations on compact surfaces, fundamentals for arbitrary codimension, and holonomy. 2nd edn., Aspects Math. E1, Friedrich Vieweg & Sohn, Braunschweig, 1986 Zbl 0628.57001 MR 0881799
- [18] C. D. Hodgson, J. H. Rubinstein, H. Segerman, and S. Tillmann, Veering triangulations admit strict angle structures. *Geom. Topol.* 15 (2011), no. 4, 2073–2089 Zbl 1246.57034 MR 2860987
- [19] M. Keane, Interval exchange transformations. *Math. Z.* 141 (1975), 25–31 Zbl 0278.28010 MR 357739
- M. Lackenby, Taut ideal triangulations of 3-manifolds. Geom. Topol. 4 (2000), 369–395
 Zbl 0958.57019 MR 1790190
- [21] M. Landry, Stable loops and almost transverse surfaces. Groups Geom. Dyn. 17 (2023), no. 1, 35–75 Zbl 07672076 MR 4563328
- [22] M. P. Landry, Y. N. Minsky, and S. J. Taylor, Flows, growth rates, and the veering polynomial. Ergodic Theory Dynam. Systems 43 (2023), no. 9, 3026–3107 MR 4624489
- [23] J. W. Milnor, *Topology from the differentiable viewpoint*. University Press of Virginia, Charlottesville, VA, 1965 Zbl 0136.20402 MR 226651
- [24] L. Mosher, Laminations and flows transverse to finite depth foliations. 1996, https://web. archive.org/web/20190829013413/http://andromeda.rutgers.edu/~mosher/, visited on 12 February 2024
- [25] A. Parlak, S. Schleimer, and H. Segerman, Veering, code for studying taut and veering ideal triangulations. 2021, http://github.com/henryseg/veering, visited on 12 February 2024
- [26] H. Segerman, Veering Dehn surgery. 2015, Talk at Moab Topology Conference, http://math. okstate.edu/people/segerman/talks/veering_Dehn_surgery.pdf, visited on 12 February 2024
- [27] S. Smale, Differentiable dynamical systems. Bull. Amer. Math. Soc. 73 (1967), no. 6, 747–817
 Zbl 0202.55202 MR 228014
- [28] W. P. Thurston, The geometry and topology of three-manifolds. 1978, Lecture notes, https:// library.slmath.org/books/gt3m/, visited on 12 February 2024

- [29] W. P. Thurston, Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. Bull. Amer. Math. Soc. (N.S.) 6 (1982), no. 3, 357–381 Zbl 0496.57005 MR 648524
- [30] W. P. Thurston, On the geometry and dynamics of diffeomorphisms of surfaces. Bull. Amer. Math. Soc. (N.S.) 19 (1988), no. 2, 417–431 Zbl 0674.57008 MR 956596
- [31] C. C. Tsang, Veering branched surfaces, surgeries, and geodesic flows. 2024, To appear in *New York J. Math.*, arXiv:2203.02874
- [32] R. L. Wilder, *Topology of manifolds*. Amer. Math. Soc. Colloq. Publ. 32, American Mathematical Society, New York, 1949 Zbl 0039.39602 MR 29491
- [33] A. Wright, Translation surfaces and their orbit closures: An introduction for a broad audience. EMS Surv. Math. Sci. 2 (2015), no. 1, 63–108 Zbl 1372.37090 MR 3354955
- [34] A. Zorich, Flat surfaces. In Frontiers in number theory, physics, and geometry. I. On random matrices, zeta functions, and dynamical systems, pp. 437–583, Springer, Berlin, 2006 Zbl 1129.32012 MR 2261104

Received 5 September 2021.

Saul Schleimer

Mathematics Institute, University of Warwick, Zeeman Building, Coventry CV4 7AL, UK; s.schleimer@warwick.ac.uk

Henry Segerman

Department of Mathematics, Oklahoma State University, 401 MSCS Building, Stillwater, OK 74078, USA; henry.segerman@okstate.edu