

Explicit angle structures for veering triangulations

DAVID FUTER
FRANÇOIS GUÉRITAUD

Agol recently introduced the notion of a *veering* triangulation, and showed that such triangulations naturally arise as layered triangulations of fibered hyperbolic 3–manifolds. We prove, by a constructive argument, that every veering triangulation admits positive angle structures, recovering a result of Hodgson, Rubinstein, Segerman, and Tillmann. Our construction leads to explicit lower bounds on the smallest angle in this positive angle structure, and to information about angled holonomy of the boundary tori.

57M50, 57R05

1 Introduction

Let \bar{M} be a compact oriented 3–manifold whose boundary $\partial\bar{M}$ is a nonempty union of tori. Suppose \bar{M} is decomposed into finitely many truncated tetrahedra Δ_i , glued in pairs along all their hexagonal faces. The *boundary triangles* or *cuspidal triangles* obtained by truncating tips of the tetrahedra remain unglued, and give a (possibly non-simplicial) triangulation of $\partial\bar{M}$. We then say that $\tau = \{\Delta_1, \dots, \Delta_n\}$ determines an *ideal triangulation* of the noncompact manifold $M = \bar{M} \setminus \partial\bar{M}$ (by removing the cuspidal triangles). Note that $\partial\bar{M}$ has two types of edges: cuspidal edges (contained in $\partial\bar{M}$) and interior edges (contained in M).

Definition 1.1 We will assign a real variable θ_j , called an *angle*, to every pair of opposite edges of every tetrahedron. Thus there are $3n$ variables for the n tetrahedra. The angles associated to Δ_i are $\theta_{3i-2}, \theta_{3i-1}, \theta_{3i}$. We impose the following system of equations:

- (1) For each tetrahedron Δ_i , the angle sum is $\theta_{3i-2} + \theta_{3i-1} + \theta_{3i} = \pi$.
- (2) For each (interior) edge of M surrounded by angles $\theta_{j_1}, \dots, \theta_{j_s}$, one has $\sum \theta_{j_i} = 2\pi$.

An angle vector $(\theta) = (\theta_1, \dots, \theta_{3n})$ satisfying equations (1) and (2) is called:

- a *generalized angle structure* on τ if $\theta_j \in \mathbb{R}$ for all j ,
 - a *taut angle structure* on τ if $\theta_j \in \{0, \pi\}$ for all j ,
 - a *positive angle structure* or simply an *angle structure* on τ if $\theta_j > 0$ for all j .
- Note that by (1), this implies $\theta_j \in (0, \pi)$ for all j .

These related definitions are ordered according to (roughly) increasing strength. Casson, Luo, and Tillmann showed that every ideal triangulation of a manifold with torus boundary admits a generalized angle structure [13, Theorem 1]. If M is irreducible and acylindrical, Lackenby showed that *some* ideal triangulation of M admits a taut angle structure [10]. On the other hand, positive angle structures are rarer and more powerful: Casson and Lackenby showed that if an ideal triangulation of M admits angle structures, then M must carry a complete hyperbolic metric. See Lackenby [11, Corollary 4.6] and our earlier paper [4, Theorem 1.2].

The theme of this paper is to deform a taut angle structure into a (positive) angle structure, with explicit information about the angles.

A taut angle structure endows every boundary triangle in the tiling of $\partial\bar{M}$ with a triple of angles $(0, 0, \pi)$. Therefore, every vertex v in the tiling of $\partial\bar{M}$ has a link consisting of two π -angles, also called *wide* angles, separated by two (possibly empty) sequences of 0-angles, also called *thin* angles. These two sequences of thin angles are called the two *fans* adjacent to v . *Wide* and *thin* will always refer to the same angles at v , even when we will start assigning them other values than $0, \pi$.

A nice source of taut angle structures comes from *layered* triangulations, which are constructed in the following way. Let S be a surface with punctures and ρ an ideal triangulation of S . Let $\phi: S \rightarrow S$ be a pseudo-Anosov diffeomorphism, and suppose M is the mapping torus $M = (S \times [0, 1]) / \sim_\phi$, where $(x, 1) \sim_\phi (\phi(x), 0)$. Find a path from the triangulation ρ to its pushforward $\phi_*(\rho)$, via a sequence of *diagonal exchanges*. Each such diagonal exchange can be seen as a flattened tetrahedron, with angles of π on the exchanged diagonals and 0 on the periphery. The union of these tetrahedra gives a taut ideal triangulation of M .

The taut angle structure on a layered triangulation admits an additional global property: a coherent choice of transverse orientation σ on the 2-skeleton, such that for each tetrahedron, the two faces sharing one π -angle have σ pointing inward, and the two faces sharing the other π -angle have σ pointing outward. We call this property *transverse-taut*.¹ It is not hard to see that if a taut angle structure is not transverse-taut, it must have a transverse-taut double cover; see Lemma 5.4. To summarize the

¹Following Lackenby [10], most authors call *taut* what we here call *transverse-taut*, as this notion is inspired by that of a taut foliation. By contrast, a taut angle structure can be called *angle-taut* for short.

implications,

$$\text{layered} \implies \text{transverse-taut} \implies \text{taut angle structure.}$$

Agol recently introduced the following notion [1].

Definition 1.2 A taut angle structure is *veering* if for every vertex v of $\partial\bar{M}$, either all the triangles of the fans of v have their π -angle immediately *before* v in the counterclockwise cyclic order, or all the triangles of the two fans of v have their π -angle immediately *after* v in the counterclockwise cyclic order. We say v is *left-veering* in the first case, *right-veering* in the second. A veering taut angle structure will be called a *veering structure* for short. We note that this definition uses the orientation on M in an essential way.

See Figure 1 (ignoring the colors and labels for the moment) for the two types of vertices in $\partial\bar{M}$ (left-veering is left). Notice that every fan is nonempty: for, if v had an empty fan, then the two wide triangles incident to v would have another common vertex w that would be neither left- nor right-veering. In other words, the rightmost fan of Figure 1 has the minimal number of triangles, namely 1. We will call a fan *short* if it contains just one triangle, and *long* otherwise.

Notice also that if $e = vv'$ is an interior edge of M (not a cusp edge!), the vertices v and v' veer in the same direction: right- or left-veeringness is an intrinsic property of the edge e .

Agol proved that veering structures are pleasantly common: in particular, there is a *canonical* layered, veering triangulation on any pseudo-Anosov mapping torus, provided that the singularities of the invariant foliations are punctures. This canonical veering

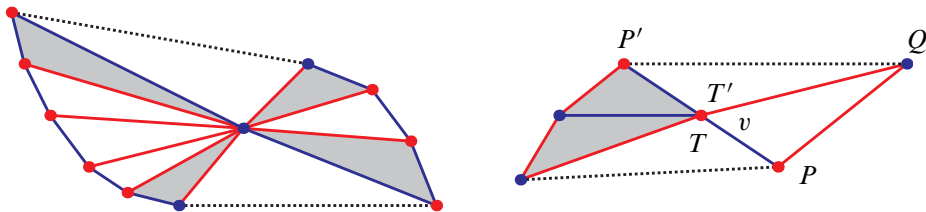


Figure 1: Two vertex links in $\partial\bar{M}$ showing fans of 5, 3, 2, and one triangles (the last one a so-called *short* fan). Obtuse angles are π , acute angles are 0. The polygon on the left surrounds a left-veering vertex, and the polygon on the right a right-veering vertex. Vertices and edges receive colors (red/blue) from the veeringness condition. Hinge triangles in the fans are shaded.

triangulation is in fact produced algorithmically by a standard (weighted) train track splitting procedure.

Shortly after Agol introduced the notion of *veering*, Hodgson, Rubinstein, Segerman, and Tillmann [7] proved

Theorem 1.3 (Hodgson, Rubinstein, Segerman and Tillman [7, Theorem 1.5]) *Let τ be a veering ideal triangulation of M . Then the veering (taut angle) structure on τ can be deformed to a positive angle structure.*

The proof in [7] is non-explicit in that it uses duality in linear programming. Namely, the linear problem of finding positive (θ_j) has a so-called *dual* linear problem, which by work of Rivin [15], Kang–Rubinstein [9], Luo–Tillmann [13], and others can be reduced to checking the absence of certain types of normal surfaces in the triangulation of M . Our aim in this paper is to give a constructive proof of Theorem 1.3, which provides effective information about the angles. In fact, we show the following.

Theorem 1.4 *Let τ be a veering ideal triangulation of M . Then τ admits a positive angle structure (θ) , whose angles satisfy*

$$\theta_i \geq \frac{\pi}{12 d_{\max}^2} \geq \frac{\pi}{12 (e_{\max} - 3)^2},$$

where d_{\max} is the maximum number of triangles in a fan of any vertex $v \in \partial \bar{M}$, and e_{\max} the largest degree of an edge of τ .

The quadratic dependence on d_{\max} or e_{\max} is sharp. On the other hand, the constant $\pi/12$ is very likely not sharp. See Remark 5.2 for more details.

Knowing explicit angles for the triangulation has several benefits. First, it is satisfying to know how to find actual examples, even though linear optimization algorithms can find them efficiently on a computer as soon as no obstruction exists.

Second, the lower bound of Theorem 1.4 is useful for algorithms in computational topology (see, for example, Cooper and Tillman [2] and Lackenby [12]). This is because known lower bounds on the angles dramatically reduce the search time required to enumerate normal surfaces of a particular genus.

Third, our construction endows every cusp with a canonical *slope* attached to the veering structure. This invariant seems not to have been pointed out before.

Fourth, our construction produces positive angle structures with zero angular holonomy. That is, the turning angle about every embedded essential curve in $\partial \bar{M}$ is zero; see

Section 6 for definitions. At present it seems to be unknown whether the existence of positive angle structures implies the existence of positive angle structures with vanishing holonomy.

Fifth, we can play further with this idea and exhibit large values of the holonomy (see Section 6). These values are not always the largest possible, but they still give interesting lower bounds on holonomy that might be useful for Dehn surgery arguments. In the transverse–taut case, we can also produce by deformation some taut angle structures distinct from the initial one (and usually non-veering).

Sixth, an ambitious goal is to find hyperbolic shapes on the tetrahedra Δ_i that glue up coherently to give the hyperbolic metric on M . By the Casson–Rivin program (see Rivin [14] and the authors’ paper [5]), this amounts to finding a critical point of the *volume functional* $\mathcal{V}: \mathcal{A}(\tau) \rightarrow \mathbb{R}$, defined on the space $\mathcal{A}(\tau)$ of positive angle structures. In practice, proving the existence of a critical point of \mathcal{V} requires a careful parametrization of $\mathcal{A}(\tau)$ and its boundary (see the second author’s paper [6]). The explicit deformations described in Sections 4 and 5 are a step toward this detailed parametrization.

Recall that the layered, veering triangulations of mapping tori constructed by Agol [1] are canonically determined by the pseudo–Anosov monodromy. The existence of a critical point of \mathcal{V} for these triangulations would show a deep interaction between combinatorics (of train tracks, say) and hyperbolic geometry. We remark, as a cautionary note, that Issa has found an example of a 13–tetrahedron triangulation coming from Agol’s construction, for which \mathcal{V} does not have a critical point [8]. In other words, this triangulation is not positively oriented in the hyperbolic metric on the ambient manifold.

Finally, the combinatorial understanding provided here might help in addressing such questions as: *Is every veering angle structure virtually layered (that is, a finite quotient of a layered structure)? Is there a universal bound on the volume of the union of the tetrahedra belonging to a fan?*

1.1 Organization

This paper is organized as follows. In Section 2, we reformulate Definition 1.2 in terms of a coloring of the edges, and explore a number of consequences for the triangulation of $\partial\bar{M}$. In Section 3, we recall the notion of *leading–trailing deformations* of a generalized angle structure, naturally associated to closed curves on $\partial\bar{M}$. In Section 4, we use leading–trailing deformations to unflatten a veering taut angle structure on τ into a positive angle structure, proving Theorem 1.3. In Section 5, we analyze how far

the angles can be unflattened, and prove Theorem 1.4. Finally, in Section 6, we explore the possible holonomies of the angle structures on a veering triangulation.

1.2 Acknowledgements

The majority of the arguments presented in this paper were discovered during the first author’s visit to Université de Lille in April 2009. At the time, Agol’s definition of veering triangulations had not yet been formulated, but in retrospect we were actually studying a special case: namely, layered triangulations of “Penner–Fathi” mapping tori defined by alternating Dehn twists (see Fathi [3]). We are grateful to Université de Lille 1 for its support and hospitality during this visit.

The appearance of the papers of Agol [1] and Hodgson, Rubinstein, Segerman and Tillman [7] prompted us to reconsider our construction, and realize that it actually works for all veering structures. We thank Ian Agol, Craig Hodgson, Hyam Rubinstein, Henry Segerman, and Stephan Tillmann for their stimulating ideas as well as remarks on the draft.

The first author is supported in part by NSF Grant No. DMS–1007221, and the second author is supported in part by the ANR program ETTT (ANR-09-BLAN-0116-01).

2 Observations from the cusp

Agol’s definition of veering structures can be reformulated in the following way. Note that Definition 1.2 implies that the edges of M can be partitioned into two families. We color the right-veering edges *red* and the left-veering edges *blue*. Hodgson, Rubinstein, Segerman, and Tillmann showed the following characterization of *veering*:

Lemma 2.1 (Hodgson, Rubinstein, Segerman and Tillman [7, Proposition 1.4]) *A taut angle structure is veering if and only if every tetrahedron Δ_i can be sent by an orientation-preserving diffeomorphism to the one depicted in Figure 2: namely a thickening of the unit square, with*

- the π -angles on the diagonals,
- the rising diagonal in front,
- the vertical (thin) edges blue, and
- the horizontal (thin) edges red.

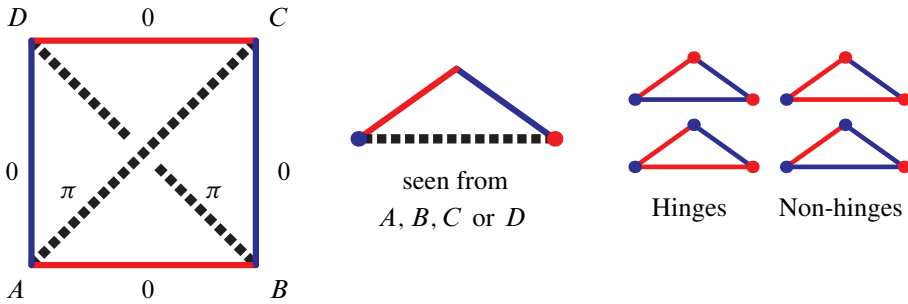


Figure 2: A colored tetrahedron (left) and its cusp view (middle). Given a triangle in the cusp view, the colors of its base and top vertex determine its type (hinge or non-hinge). Throughout this paper, the lighter color (printed in greyscale) is red, and the darker color is blue.

Note that in this characterization of veeringness, the colors of the two wide diagonals (black in Figure 2) may be anything: that is, they are determined by the adjacent tetrahedra. This induces a partition of the tetrahedra of M into *hinges*, whose diagonals (that is, wide angles) bear different colors, and *non-hinges*, whose diagonals are the same color. We are importing this terminology from our earlier paper [6].

Next, we can color all the edges and vertices of the tiling of $\partial\bar{M}$ in red and blue: the color of a vertex v of $\partial\bar{M}$ is the color of the edge in M incident to v , while the color of an edge $e \subset \partial M$ is the color of the *base* of the ideal triangle in M whose opposite tip was truncated to yield e . A consequence of the red/blue characterization of veeringness is that for each triangle T of $\partial\bar{M}$, if we draw the π -angle on top, then the left edge and right vertex are red, while the right edge and left vertex are blue. The bottom edge and top vertex could be any color (determined by some adjacent triangle T'); these two colors disagree if and only if T is a truncation of a hinge tetrahedron (called a hinge triangle). Note that hingeness is an inherent property of the tetrahedron, inherited by all four of its boundary triangles. See Figure 2, right.

We can now revisit Figure 1 and apply colors according to the rule above. Notice that all vertices and edges in the figure receive a determined color, except for the bases of the two triangles that have wide angles at v . We can next make a series of observations:

Observation 2.2 The vertex v is connected to precisely two vertices P, P' with the same color as v (belonging to the two wide triangles incident to v). In the cyclic order for the neighbors of v , these two vertices are not consecutive, because each fan is nonempty. Therefore, if v is for example red, then v has blue neighbors on both sides (in both fans).

Observation 2.3 As a consequence, if we draw all edges in $\partial\bar{M}$ that connect two vertices of the same color (incidentally, such an edge is always of the other color: check Figure 2, right), then this defines a system of disjoint curves $\gamma_1, \gamma_2, \dots$ on $\partial\bar{M}$, passing through all vertices. We claim that no curve γ_i can bound a disk in $\partial\bar{M}$. For, suppose without loss of generality that γ_1 has red vertices and bounds an *innermost* disk. By Observation 2.2, any vertex v on this curve must have blue neighbors on both sides of γ_1 . Hence there is a blue vertex inside γ_1 that belongs to some γ_i , contradicting the assumption that γ_1 was innermost.

Therefore the complement of the union of the curves γ_i is a union of annuli \mathcal{L}_i , because each component of $\partial\bar{M}$ has Euler characteristic 0. The number of parallel curves γ_i inside each torus component T of $\partial\bar{M}$ is even, because their colors alternate. Note that the *slope* of the γ_i in $\mathbb{P}H_1(T, \mathbb{Z}) \simeq \mathbb{Q} \cup \{\infty\}$ is an invariant of the veering structure. See Observation 2.9 for the meaning of this slope in the context of layered triangulations of mapping tori.

Observation 2.4 Inside each annulus \mathcal{L}_i , every edge connects one boundary component of \mathcal{L}_i to the other (indeed the edge has ends of distinct colors because it does not belong to the γ_i). So \mathcal{L}_i has the structure of a *ladder*, with two *ladderpoles* γ_i, γ_{i+1} (red and blue) connected by many *rungs*. Two consecutive rungs always have a common endpoint, and each triangle of \mathcal{L}_i is bounded by two rungs and one ladderpole segment (see Figure 3 in anticipation).

The combinatorics of the rungs inside each annulus \mathcal{L}_i could be expressed as a cyclic sequence of “rights” and “lefts,” echoing the situation with punctured torus bundles [6]. However, in our more general setting the sequences for distinct annuli are generally unrelated.

Observation 2.5 For v a red vertex, the two π -angles at v are *clockwise* just after the ladderpole through v , for the cyclic order on the link of v . For v a blue vertex, the π -angles are *counterclockwise* just after the ladderpole. The rest of the space on each side of the ladderpole is occupied by a full fan of v . This is visible in Figure 1.

Observation 2.6 We can examine Figure 1 to determine which triangles of the fans of v are hinges. (Ignore the two triangles outside the fans of v , that is, the ones having π -angles at v : they belong to other fans. For example the triangle T' belongs to a fan of P' and a fan of Q .) A triangle is hinge if and only if its vertex at the wide angle has a color different from the opposite edge, so examination yields the following:

- In a long fan (of 2 or more triangles), only the first and last triangle are hinges.

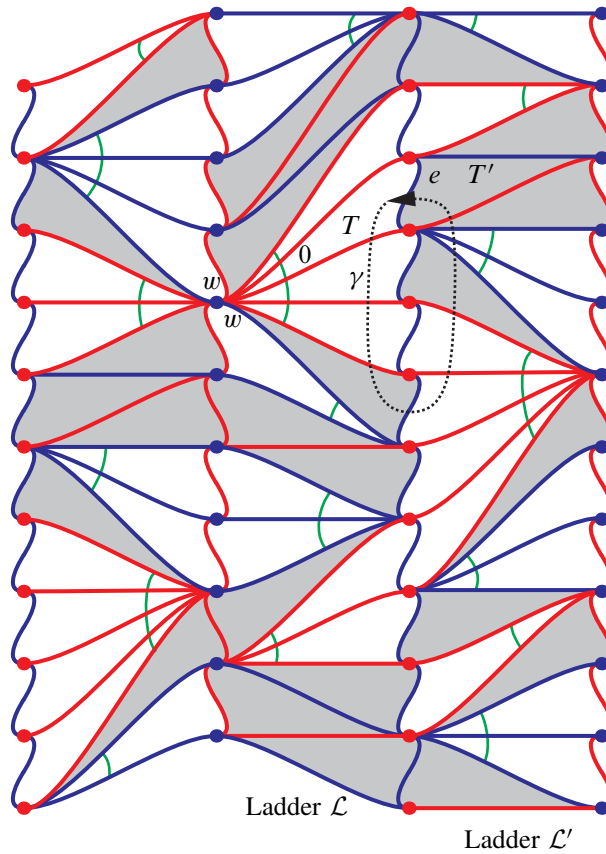


Figure 3: View of $\partial\bar{M}$ with all colors. Ladderpoles are vertical, hinge triangles are shaded, and wide vs. thin angles are distinguished by a smoothing convention at each vertex. Ladderpoles may have different lengths, as the cutouts at the top and bottom suggest. Flat angles (after Lemma 4.1 of the “rescuing algorithm”) are marked in green. In Lemma 4.2, the curve γ rescues the flat angle marked 0, via the deformation D^γ .

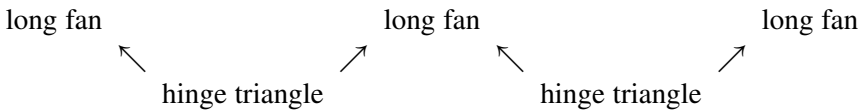
- In a short fan, the (single) triangle is not a hinge.

Hinge triangles are shaded in Figure 1.

Observation 2.7 A triangle of a *short* fan of a vertex v (such as vPQ in Figure 1) belongs to a *long* fan of the vertex at its other thin corner (here Q), since the angles \widehat{vQP} and $\widehat{vQP'}$ are both thin. In fact, the fan at Q is not just long but has length ≥ 3 , since the triangle vPQ , being non-hinge, cannot be the first nor the last triangle of its fan at Q .

This implies in particular that *there are* hinge triangles, and in fact that there are some inside every annulus \mathcal{L}_i of every boundary torus of \bar{M} .

Therefore, every hinge triangle belongs to precisely two long fans (while every non-hinge triangle belongs to precisely one long fan). Since every long fan also contains precisely two hinge triangles, we can jump from one long fan to the next according to the following scheme:



where arrows denote inclusion. By construction, this sequence of fans follows one of the ladders \mathcal{L}_i .

All the above observations are summarized in Figure 3, a view of the triangulation of $\partial\bar{M}$, also called the “cusp view”. Ladderpoles are vertical, hinge triangles are shaded, and we use a train track-like smoothing convention to emphasize which angles near any given vertex are the wide angles (look, for example, for the two w -labels for “wide”).

We close this section with two observations that hold true under additional hypotheses.

Observation 2.8 Suppose that the veering triangulation τ is transverse–taut. Recall this means there is a transverse orientation on all the faces (for example, towards the reader in the first panel of Figure 2) that is consistent over all tetrahedra. This orientation induces an “upward” orientation along the ladders in each cusp triangulation. Inside each cusp, ladders of two types alternate: in one type, the wide angle of each triangle is always *above* the base (as in ladder \mathcal{L} in Figure 3); in the other type, the wide angle is always *below* the base (as in ladder \mathcal{L}'). We call the first type of ladder *ascending*, and the second type *descending*.

If the transverse orientation of Figure 2 points toward the reader, the truncated vertices at A and C belong to ascending ladders, while the truncated vertices at B and D belong to descending ladders. In particular, each tetrahedron has two vertices in each type of ladder; this fact will be crucial in Sections 5 and 6.

Observation 2.9 Suppose that the veering triangulation τ comes from Agol’s construction of a layered triangulation of a mapping torus with monodromy ϕ , with punctures at the singularities of the ϕ -invariant foliations $\mathcal{F}, \mathcal{F}'$. Then the germs of the singular leaves of $\mathcal{F}, \mathcal{F}'$ incident to a puncture of the fiber define a slope in each cusp of M . One can prove that these slopes are the same as the slopes of the ladderpoles γ_i . At each cusp of M , the number of prongs of \mathcal{F} (or \mathcal{F}') is equal to the

number of pairs of ladders, multiplied by the intersection number of the ladder slope with the fiber slope.

3 Leading–trailing deformations

Let $GAS(\tau)$ be the set of generalized angle structures on a triangulation τ , where the angles θ_j can take any value in \mathbb{R} . In this section, we exhibit a spanning set for the tangent space $T_p GAS(\tau)$. All the facts in this section are proved in our earlier paper [5, Section 4], and we point to that paper for more detail.

Definition 3.1 A *normal curve* on a component of $\partial\bar{M}$ is an embedded, oriented closed curve γ transverse to the 1–skeleton of the triangulation of $\partial\bar{M}$, such that γ enters and exits each triangle through different edges.

For each edge e crossed by a normal curve γ , there is a real variable θ_{j_e} associated to the angle opposite e in the triangle *entered* by γ , and a real variable $\theta_{j'_e}$ associated to the angle opposite e in the triangle *left* by γ . Let ε_j be the j th basis vector of \mathbb{R}^{3n} . Then the *leading–trailing deformation* associated to γ is the vector

$$D^\gamma := \sum_e (\varepsilon_{j_e} - \varepsilon_{j'_e}).$$

Here each edge e can appear many times in the sum (as often as it is crossed by γ) and the indices j_e, j'_e may be swapped according to the direction in which γ crosses e .

Lemma 3.2 (Futer and Guéritaud [5, Lemma 4.5]) *The vector D^γ is tangent to $GAS(\tau)$. That is, if $(\theta) = (\theta_j)_{1 \leq j \leq 3n}$ is a generalized angle structure, then $(\theta) + tD^\gamma$ is also one, for every real t .*

In other words, given an angle structure (θ) , one may increase (resp. decrease) by t all the angles opposite edges crossed by γ , according to the direction of crossing (with multiplicity). An equivalent and sometimes useful way of seeing the deformation D^γ is as follows: whenever γ traverses a triangle abc by entering through ab and leaving through bc , increase the angle at c and decrease the angle at a . See Figure 4.

Remark 3.3 Suppose that γ crosses a boundary triangle of a (truncated) tetrahedron Δ . Then the vector D^γ deforms the dihedral angles of four edges in Δ , which are adjacent to all four boundary triangles of Δ . As a result, the deformation D^γ can affect the shapes of boundary triangles that do not intersect γ , including truncation triangles that belong to completely different components of $\partial\bar{M}$.

Remark 3.4 Let ϵ be an edge in the interior of M that connects two vertices $v, v' \in \partial\bar{M}$. Then the clockwise loops γ_v and $\gamma_{v'}$ that encircle v and v' , respectively, induce *identical* deformations $D^{\gamma_v} = D^{\gamma_{v'}} \in T_p GAS(\tau)$. We call this common deformation D^ϵ . This is in contrast to other curves γ , each of which lives on a particular boundary component of $\partial\bar{M}$.

The deformation D^{γ_v} is illustrated in Figure 4, right. A comparison of this picture with Figure 1 should convince the reader that D^{γ_v} (resp. $-D^{\gamma_v}$) can be used to unflatten the triangles in the two fans of a blue (resp. red) vertex v , although it might place a negative angle at the thin vertices of the two wide triangles incident to v . The next section develops this idea.

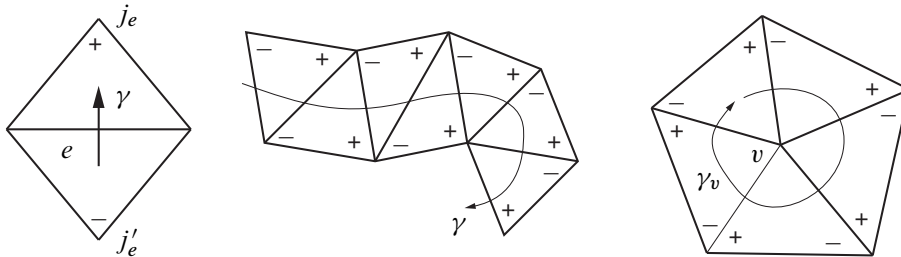


Figure 4: Various (segments of) curves γ and the associated perturbations D^γ , expressed via signs $+$ and $-$.

4 Rescuing the zeros

In this section, we describe an algorithm to deform any veering structure to a positive angle structure, by applying various deformations D^γ . This algorithm will *rescue* (that is, unflatten) the zero angles of the tetrahedra, one at a time. The hinge tetrahedra are rescued in Lemma 4.1, and the non-hinge tetrahedra in Lemma 4.2. Together, these lemmas imply Theorem 1.3.

For each edge ϵ in the interior of M , let D^ϵ be the clockwise deformation about ϵ , as described in Remark 3.4 and shown in Figure 4, right. Define

$$(4-1) \quad D := \sum_{\epsilon \text{ blue}} D^\epsilon - \sum_{\epsilon \text{ red}} D^\epsilon.$$

Lemma 4.1 *Let (θ) be a veering taut angle structure on M . Choose any $t \in (0, \pi/4)$, and deform (θ) to the generalized angle structure $(\theta') = (\theta) + tD$, for the deformation D in (4-1). Then (θ') has the following properties:*

- (1) Every hinge tetrahedron (or triangle) has positive angles.
- (2) Every non-hinge tetrahedron has non-negative angles.
- (3) All 0-angles are opposite ladderpole segments in Figure 3.

In fact, after applying the deformation tD , every non-hinge triangle will have precisely one vanishing angle, namely the one opposite the ladderpole². The vanishing angles are marked green in Figure 3.

Proof Fix a tetrahedron Δ , hinge or non-hinge. It has two blue edges with 0 angles (call these e_1, e_2), two red edges with 0 angles (call these e_3, e_4), and two diagonals with π angles (call these e_5, e_6). We will treat the deformations D^{e_j} one at a time, as though e_1, \dots, e_6 are distinct edges in M . For, if some of these edges (for concreteness, e_1 and e_2) are identified to the same edge $e \subset M$, the deformation of D^e affects Δ in two different ways, one that corresponds to e_1 and the other that corresponds to e_2 . Thus the cumulative effect of the deformations D^{e_j} will be the same, regardless of whether some edges of Δ are identified.

First, consider the four thin edges of Δ . The clockwise deformation tD^{e_1} decreases the π -angles by t and increases the red 0-angles by t . The effect of tD^{e_2} is exactly the same. In a similar fashion, each of the counterclockwise deformations $-tD^{e_3}$ and $-tD^{e_4}$ decreases the π -angles by t and increases the blue 0-angles by t . We conclude that after performing the deformations along the four thin edges of Δ_i , the dihedral angles of the tetrahedron become

$$(4-2) \quad (\pi - 4t, 2t, 2t).$$

Next, consider the diagonals e_5, e_6 of Δ . Observe that deformations D^{e_5} and D^{e_6} along opposite edges of Δ have precisely the same effect on Δ . If Δ is a hinge tetrahedron, then its two diagonals have opposite colors, and the effect of $tD^{e_5} - tD^{e_6}$ cancels out completely. Thus Δ ends up with angles $(\pi - 4t, 2t, 2t)$, which are always positive for $t \in (0, \pi/4)$. This proves (1).

If Δ is not a hinge and both diagonals are red, then each of the deformations tD^{e_5} and tD^{e_6} increases the thin red angles by t , and decreases the thin blue angles by t . Thus, after performing the deformations along both e_5 and e_6 , the angles of Δ become

$$(4-3) \quad (\pi - 4t, 4t, 0),$$

²It is worth describing what happens in the special case of punctured torus bundles. After applying Lemma 4.1, in the notation and terminology of our paper [6] we have " $w_i \equiv 2t$ " for all i , so all "hinge conditions" are satisfied and one only needs to relax the "concavity conditions".

where the blue angles are 0. Observe in Figure 3 (or Figure 1) that in a non-hinge triangle coming from a tetrahedron with red diagonals, the ladderpole connects red vertices and is opposite the blue thin angle. Thus Δ will have non-negative angles everywhere, and 0 precisely at the angles opposite the ladderpole segments.

If Δ has blue diagonals, the effect is exactly symmetric. The deformations along e_5 and e_6 will make the angles of Δ become $(\pi - 4t, 0, 4t)$, with zeros at the thin red angles opposite the ladderpole segments. This proves (2) and (3). \square

Next, we rescue the non-hinge tetrahedra, using the following inductive procedure.

Lemma 4.2 *Suppose that the tetrahedra $\Delta_1, \dots, \Delta_n$ are endowed with a generalized angle structure $(\theta) = (\theta_1, \dots, \theta_{3n})$, such that these angle assignments satisfy properties (1), (2), and (3) of Lemma 4.1. Suppose as well that some angle θ_j is 0.*

Then there is a homotopically trivial closed curve $\gamma \subset \partial M$, such that for small $t > 0$ the structure $(\theta') = (\theta) + tD^\gamma$ again satisfies properties (1), (2), and (3) of Lemma 4.1, but has strictly fewer vanishing angles than (θ) .

Proof Call a tetrahedron *degenerate* if it has an angle labeled 0, and *non-degenerate* otherwise. Similarly, call a boundary triangle *flat* if it has an angle of 0, and *non-flat* otherwise.

Since M is connected, there must be a degenerate tetrahedron Δ that is glued along face f to a non-degenerate tetrahedron Δ' . By looking at the appropriate truncated vertex of f (three possible choices), we can assume that in $\partial \bar{M}$, we have the 0-angle in a triangle $T \subset \Delta$, opposite an edge e shared with a non-flat triangle $T' \subset \Delta'$. By property (3), e must be a ladderpole segment that forms part of the boundary between annuli \mathcal{L} and \mathcal{L}' , with $T \subset \mathcal{L}$ and $T' \subset \mathcal{L}'$. See Figure 3.

Begin the curve γ by crossing edge e from T' into T . (When crossing into T opposite the vanishing angle, we instantly unflatten the angle of 0, by definition of D^γ .) Next, travel vertically through the fan of T in annulus \mathcal{L} , until coming to a hinge triangle that shares an edge with \mathcal{L}' . (Note that a hinge triangle at the top or bottom of the fan of T must be adjacent to \mathcal{L}' ; inspection in Figure 3 shows that when traversing the fan of \mathcal{L} , the curve γ must exit each triangle of the fan across the edge connecting the two *thin* angles.) Cross back into \mathcal{L}' at the hinge, and travel in the opposite direction, until γ returns to triangle T' and closes up.

Note that when γ travels vertically and crosses the rungs of a ladder, it never leaves a triangle opposite an angle of 0, hence all angles decreased by D^γ are strictly positive. In addition, γ leaves \mathcal{L} through a hinge triangle that has positive angles by (1), and

leaves \mathcal{L}' through a triangle T' that has positive angles by hypothesis. Thus D' never decreases an angle of 0, and preserves properties (1), (2), and (3) while “rescuing” at least one tetrahedron. \square

We may repeatedly apply Lemma 4.2 until all angles are positive. This completes our constructive proof of Theorem 1.3. \square

5 A lower bound on the smallest angle

In this section, we prove Theorem 1.4. We will first prove the theorem under the additional hypothesis that the triangulation τ is transverse–taut, and then apply a covering argument to extend the result to all veering triangulations.

Proposition 5.1 *Let τ be a transverse–taut, veering ideal triangulation of M . Then τ admits a positive angle structure (θ) , whose angles satisfy*

$$(5-1) \quad \theta_i \geq \frac{\pi}{12 d_{\max}^2} \geq \frac{\pi}{12 (e_{\max} - 3)^2},$$

where d_{\max} is the maximum number of triangles in a fan of any vertex $v \in \partial \bar{M}$, and e_{\max} the largest degree of an edge of τ .

Remark 5.2 The quadratic dependence on d_{\max} is optimal, even in the special case of punctured torus bundles. In the notation of our paper [6], a non-hinge tetrahedron in the i th spot of a syllable LR^nL will have an angle of the form

$$x_i = 2w_i - (w_{i-1} + w_{i+1}),$$

where all parameters must satisfy $w_j \in (0, \pi/2)$. The requirement $x_i > 0$ implies that the sequence of parameters w_{i-1}, w_i, w_{i+1} is concave, hence is called the *concavity condition*. Now, a fan of length $(n + 1)$ imposes a concave sequence w_0, \dots, w_n that must stay at most distance $\pi/2$ above the line segment from $(0, w_0)$ to (n, w_n) . Summing by parts twice, we see that the range condition $w_j \in (0, \pi/2)$ cannot be satisfied if $x_i \geq 4\pi/n^2$ for all $i = 1, \dots, n - 1$. This indicates that the quadratic behavior is sharp but the constant $\pi/12$ is probably far from sharp.

In fact, [6, Proposition 10.1] implies that when the monodromy is $L^n R^n$ and (θ) is the *geometric* structure on the triangulation (that is, the unique angle structure under which the tetrahedra glue up to give the hyperbolic metric on M), the smallest angle will be approximately $4\pi^2/n^3$ for large n .

Proof of Proposition 5.1 The proof will use the same deformations D^ν as in Lemmas 4.1 and 4.2, with careful choices of coefficient t .

Given a transverse orientation on the faces of τ (say, towards the reader in the first panel of Figure 2), Observation 2.8 implies that each tetrahedron has exactly two tips in ascending ladders, and two tips in descending ladders. Consider the collection \mathcal{T} of all cusp triangles belonging to ascending ladders: \mathcal{T} has cardinality $2n$ if there are n tetrahedra. The set \mathcal{T} is naturally endowed with a fixed-point-free involution σ taking each triangle to the only other triangle in \mathcal{T} that belongs to the same tetrahedron.

Let T be a triangle in \mathcal{T} , belonging to some ascending ladder \mathcal{L} . Define the *height* of T , denoted $H(T)$, to be the length of the shortest path *down* the ladder \mathcal{L} that connects T to a hinge triangle: for example, hinge triangles have height 0; their nonhinge neighbors *immediately above* in the ascending ladder have height 1, and so on. There is no obvious *a priori* relationship between $H(T)$ and $H(\sigma(T))$, except that they can only vanish simultaneously (when $T, \sigma(T)$ belong to a hinge tetrahedron).

Further, as in the proof of Lemma 4.2, define γ_T to be the oriented path that enters T through the ladderpole edge (leaving another ladder $\mathcal{L}' \neq \mathcal{L}$), travels $H(T)$ rungs down the ladder \mathcal{L} until it reaches a hinge triangle, then crosses back into \mathcal{L}' and travels across the rungs of \mathcal{L}' to close up. For a hinge triangle T , by convention we define γ_T to be the trivial (that is, empty) path.

Consider the generalized angle structure (θ') given by Lemma 4.1 with $t = \frac{\pi}{6}$. By equations (4-2) and (4-3), the angles of (θ') have the following properties:

- Hinge triangles are equilateral, with angles $(\frac{\pi}{3}, \frac{\pi}{3}, \frac{\pi}{3})$;
- Nonhinge triangles have angles $(\frac{\pi}{3}, \frac{2\pi}{3}, 0)$, with $\frac{\pi}{3}$ at the wide angle and 0 belonging to a long fan.

Define d to be the number of *non-hinge* triangles in the longest fan in an ascending ladder. Equivalently, $d = \max_{T \in \mathcal{T}} H(T)$ is the largest number of adjacent green arcs in an ascending ladder in Figure 3. Note that by the definition of d_{\max} , one has $d + 2 = d_{\max}$. (The inequality $d + 2 \leq d_{\max}$ is immediate, since d_{\max} includes hinge triangles and counts both ascending and descending ladders. Equality holds because to each fan in a descending ladder at an endpoint of an edge e of M corresponds a fan of the same length in an ascending ladder at the other endpoint of e .) Also, by construction, every cusp triangle (whether in the ascending collection \mathcal{T} or not) is crossed by at most $2d$ curves γ_T : this is because the curves γ_T are nested by families of at most d along the ladderpoles, and each ladder has 2 poles.

If $d = 0$, then all tetrahedra are hinges, $d_{\max} = 2$ and $e_{\max} = 6$. Inequality (5-1) easily holds for the equilateral hinge triangles. Thus we may suppose that $d \geq 1$.

Claim 5.3 Assume that $d \geq 1$, and let $\kappa = \frac{\pi}{24}$. Then the angle structure

$$(5-2) \quad (\theta) := (\theta') + \sum_{T \in \mathcal{T}} \frac{\kappa}{d_{\max}^2} \cdot H(\sigma(T)) D^{\gamma_T}$$

is positive, with smallest angle equal to at least $2\kappa/d_{\max}^2$. (We may see the sum above as being over all triangles $T \in \mathcal{T}$, even though hinge triangles contribute 0.)

This claim clearly finishes the proof that all angles are at least $2\kappa/d_{\max}^2$. The remaining inequality in the Proposition (relating angles to the maximum edge degree e_{\max}) follows from the observation that $e_{\max} \geq d_{\max} + 3$: for, d_{\max} only counts the thin angles in one fan at v , but there must also be two wide angles, and at least one angle in the other fan. Thus all that remains is to prove the claim.

First, consider a hinge tetrahedron Δ which contributes cusp triangles $T_1, T_2 \in \mathcal{T}$ and $T_3, T_4 \notin \mathcal{T}$. Each of T_1, \dots, T_4 is crossed by at most $2d$ curves γ_T , and each curve carries weight at most

$$\frac{\kappa \cdot \max H}{d_{\max}^2} = \frac{\kappa d}{d_{\max}^2}.$$

Each angle of Δ starts out with a value of $\pi/3 = 8\kappa$ and is affected by at most $4 \cdot 2d \cdot \kappa d/d_{\max}^2$. Thus all the (θ) -angles of Δ are at least

$$\theta_j \geq 8\kappa - \frac{8\kappa d^2}{d_{\max}^2} = 8\kappa \left(\frac{d_{\max}^2 - d^2}{d_{\max}^2} \right) = 8\kappa \left(\frac{(d+2)^2 - d^2}{d_{\max}^2} \right) = 8\kappa \left(\frac{4d+4}{d_{\max}^2} \right) \geq \frac{64\kappa}{d_{\max}^2}.$$

Next, consider a non-hinge triangle Δ . The same calculation as above applies to show that the angles of Δ that are nonzero for the angle structure (θ') are still at least $64\kappa/d_{\max}^2$ for the angle structure (θ) . It only remains to deal with the flat angle of Δ .

We still assume that Δ has cusp triangles $T_1, T_2 \in \mathcal{T}$ and $T_3, T_4 \notin \mathcal{T}$. By definition, σ exchanges T_1 and T_2 . Let $T'_3 \in \mathcal{T}$ be the neighbor of T_3 sharing a ladderpole segment with T_3 , and $T'_4 \in \mathcal{T}$ be the neighbor of T_4 sharing a ladderpole segment with T_4 . The following deformations D^{γ_T} all affect the 0-angle of Δ :

- $D^{\gamma_{T_1}}$ (positively because γ_{T_1} enters T_1 through the ladderpole);
- $D^{\gamma_{T_2}}$ (positively because γ_{T_2} enters T_2 through the ladderpole);
- $D^{\gamma_{T'_3}}$ (negatively because $\gamma_{T'_3}$ leaves T_3 through the ladderpole);
- $D^{\gamma_{T'_4}}$ (negatively because $\gamma_{T'_4}$ leaves T_4 through the ladderpole).

Any additional deformations D^{γ_T} may only affect the 0-angle of Δ positively: this occurs when γ_T happens to enter a triangle of Δ when it crosses back into the

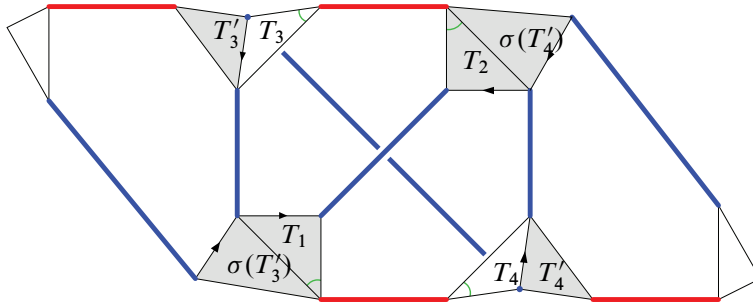


Figure 5: A non-hinge tetrahedron Δ and its cusp triangles T_1, \dots, T_4 , together with two adjacent tetrahedra. Inside each cusp, the ladderpole edges (always connecting vertices of the same color) are shown with their upward orientation induced by the transverse–taut structure. Triangles of the ascending collection \mathcal{T} are shaded, and the 0–angles of Δ for (θ') are marked in green as in Figure 3. For a non-hinge tetrahedron with four red edges, the figure would be reflected across a line of slope $+1$, with red and blue interchanged.

descending ladder \mathcal{L}' (in other words, when T'_3 or T'_4 is a hinge and happens to be the nearest hinge below T in the ascending ladder \mathcal{L}).

The key observation is that $\sigma(T'_3)$ and $\sigma(T'_4)$ are the neighbors of T_1 and T_2 down the ascending ladders: see Figure 5. This implies

$$(5-3) \quad H(\sigma(T'_3)) = H(T_1) - 1 \quad \text{and} \quad H(\sigma(T'_4)) = H(T_2) - 1.$$

As a result, by equation (5-2), the angle $\theta'_j = 0$ of Δ will become

$$\begin{aligned} \theta_j &\geq \frac{\kappa}{d_{\max}^2} \cdot (H(\sigma(T_1))D^{\gamma_{T_1}} + H(\sigma(T_2))D^{\gamma_{T_2}} + H(\sigma(T'_3))D^{\gamma_{T'_3}} + H(\sigma(T'_4))D^{\gamma_{T'_4}})_j \\ &\geq \frac{\kappa}{d_{\max}^2} \cdot (H(\sigma(T_1)) + H(\sigma(T_2)) - H(\sigma(T'_3)) - H(\sigma(T'_4))) \\ &= \frac{\kappa}{d_{\max}^2} \cdot (H(T_2) + H(T_1) - [H(T_1) - 1] - [H(T_2) - 1]) \\ &= \frac{2\kappa}{d_{\max}^2}, \end{aligned}$$

completing the proof. □

To prove Theorem 1.4 in general, we need to lift the angle structure (θ) to a transverse–taut double cover. The following lemma establishes the existence of such a cover.

Lemma 5.4 *Let τ be an ideal triangulation of M with a taut angle structure. If τ is not transverse–taut, then there is a double cover $N \rightarrow M$, such that the lift of τ to N is transverse–taut.*

Proof The (unoriented) transverse direction defines a line bundle $B \rightarrow M$, which is a natural subbundle of TM . If B is orientable (that is, it admits a nonzero section), then τ is transverse–taut. If not, then B lifts to a line bundle \tilde{B} over the universal cover \tilde{M} of M , and \tilde{B} is orientable because \tilde{M} is simply connected. The group of deck transformations of \tilde{M} which preserve the orientation of \tilde{B} has index 2, and the corresponding double cover N of M satisfies the conditions. \square

Proof of Theorem 1.4 Let τ be a veering, taut triangulation of M . If the taut angle structure on τ is actually transverse–taut, then the proof is complete by Proposition 5.1. Otherwise, Lemma 5.4 guarantees a double cover $N \rightarrow M$, such that the lift $\tilde{\tau}$ of τ is transverse–taut. Note that tetrahedra, angles, and edge degrees all lift to finite covers. Thus $d_{\max}(M) = d_{\max}(N)$ and $e_{\max}(M) = e_{\max}(N)$.

By Proposition 5.1, the transverse–taut triangulation $\tilde{\tau}$ of N admits a positive angle structure (θ) , with all angles satisfying inequality (5-1). Let σ be the involution of N that acts as a deck transformation of the (regular) double cover $N \rightarrow M$. Then, because the polytope of positive angle structures is convex,

$$(\theta') = \frac{(\theta) + \sigma(\theta)}{2}$$

is also an angle structure on N , also satisfies (5-1), and is σ –equivariant. Projecting the angles of (θ') down to M completes the proof. \square

6 Holonomies

6.1 Definitions

In this section, we explore the possible holonomies of the positive angle structures on a veering triangulation. Informally, the holonomy of a curve γ in $\partial\bar{M}$ is its total turning angle according to the angles θ_j of the angle structure. We will show that the holonomy of γ only depends on the homology class $[\gamma]$, and is a linear functional on homology classes. In other words, the holonomy naturally lives in the first cohomology of $\partial\bar{M}$.

Definition 6.1 Let V be a torus component of $\partial\bar{M}$, endowed with its possibly non-simplicial triangulation. Let $\gamma \subset V$ be an oriented, normal closed curve, as

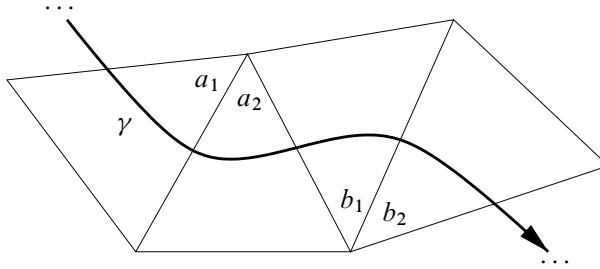


Figure 6: Here $\mathcal{A}_\gamma = a_1 + a_2 + (\dots)$ and $\mathcal{B}_\gamma = b_1 + b_2 + (\dots)$.

in Definition 3.1. Then a component of γ in a triangle of V cuts off exactly one corner the triangle, which is either to the left or right of γ .

Let (θ) be a generalized angle structure on the triangulation τ . Then every corner of each triangle on V receives a (real-valued) “angle” assignment. Let \mathcal{A}_γ be the sum of angles to the left of γ , and \mathcal{B}_γ the sum of angles to the right of γ . Note that each θ_j may appear several times in \mathcal{A}_γ and \mathcal{B}_γ . Then we define the *turning angle of γ* to be

$$(6-1) \quad t_\theta(\gamma) = \mathcal{A}_\gamma - \mathcal{B}_\gamma .$$

See Figure 6. Recall that every primitive homology class in $H_1(V, \mathbb{Z})$ is represented by a simple closed curve, which can be taken to be normal after a small isotopy. If $[\gamma]$ is a nonzero, primitive homology class and $n \in \mathbb{Z}$, define the *angular holonomy of $n[\gamma] \in H_1(V, \mathbb{Z})$* to be

$$(6-2) \quad h_\theta(n[\gamma]) = n(\mathcal{A}_\gamma - \mathcal{B}_\gamma),$$

for an arbitrary normal curve γ representing $[\gamma]$.

Equation (6-2) begs the question of whether holonomy is well-defined. In fact, we have

Proposition 6.2 *Let V be a torus component of $\partial \bar{M}$, and (θ) a generalized angle structure on τ . Let $\gamma \subset V$ be an oriented, normal, homologically non-trivial closed curve. Then the turning angle of γ only depends on the homology class of γ , hence the holonomy $h_\theta([\gamma])$ is well-defined. Furthermore, for $a, b \in \mathbb{Z}$ and homology classes $\omega, \eta \in H_1(V, \mathbb{Z})$,*

$$h_\theta(a\omega + b\eta) = a \cdot h_\theta(\omega) + b \cdot h_\theta(\eta).$$

In other words, $h_\theta: H_1(\partial \bar{M}) \rightarrow \mathbb{R}$ is a linear functional, hence $h_\theta \in H^1(\partial \bar{M}, \mathbb{R})$.

To appreciate the non-triviality of this claim, observe that the trivial homology class is represented by a curve γ encircling a single vertex of V (see Figure 4). All the angles

cut off by γ lie to one side of the curve, and by Definition 1.1 these angles sum to 2π . Thus the turning angle of this trivial curve is $\pm 2\pi$, depending on the orientation of γ , rather than 0.

We note that angular holonomy is also a restriction of a well-known definition of complex-valued holonomy, which keeps track of complex-valued shape parameters instead of just real-valued angles. (See Thurston’s notes [16, Chapter 4]. See also [5, Section 2], where holonomy is defined only for normal curves, without any claim of linearity.) The turning angle of equation (6-1) is precisely the imaginary part of the (log) holonomy from [16]. The issue of finding the right multiple of 2π is identical to the issue of finding the correct branch of a complex-valued logarithm.

Proof of Proposition 6.2 Let $\tilde{V} \cong \mathbb{R}^2$ be the universal cover of the torus V . The orientation of V , the triangulation of V , and the generalized angle structure (θ) all lift to \tilde{V} . We begin by assigning a coherent direction to every oriented edge of \tilde{V} .

Claim 6.3 *Let X be the collection of all oriented edges of \tilde{V} . Then there exists a map*

$$\psi_\theta: X \rightarrow \mathbb{R}/2\pi\mathbb{Z},$$

assigning each $x \in X$ a direction $\psi_\theta(x)$, with the following property. Whenever $x, y \in X$ have the same tail vertex p , with y immediately following x for the counterclockwise cyclic order at p , and x' denotes the same edge as x with opposite orientation, then

$$(6-3) \quad \psi_\theta(x') = \psi_\theta(x) + \pi \quad \text{and} \quad \psi_\theta(y) = \psi_\theta(x) + \theta_j,$$

where θ_j is the angle between (the projections to V of) x and y .

Consider the graph Γ with vertex set X and edges of the form xx' and xy (with x, x', y as in the claim). To define the function ψ_θ , it suffices to set $\psi_\theta(x_0) = 0$ for one arbitrary $x_0 \in X$. Then, equation (6-3) gives a way to extend the definition to every other $x \in X$. Because Γ is connected, each $x \in X$ will receive (at least one) direction vector. The main content of Claim 6.3 is that the definition is consistent, that is, no contradictions arise in extending the definition over Γ .

Up to isomorphism, Γ is the infinite trivalent planar graph obtained from the 1-skeleton of \tilde{V} by replacing every v -valent vertex with a v -gon. The complement of Γ in the plane has two types of components:

- (1) v -gons coming from vertices p of \tilde{V} . The definition is consistent along the boundary of such a v -gon because the v angles around p add up to $2\pi \equiv 0 [2\pi]$ for the angle structure (θ) .

- (2) Hexagons coming from triangles of \tilde{V} with angles $\theta_i, \theta_j, \theta_k$. The definition is consistent along the boundary of such a hexagon because $\pi + \theta_i + \pi + \theta_j + \pi + \theta_k = 4\pi \equiv 0 [2\pi]$ for the angle structure (θ) .

Since \mathbb{R}^2 is simply connected, every closed loop in Γ bounds some combination of v -gons and hexagons. Thus ψ_θ is well-defined, proving Claim 6.3. Moreover, given the generalized angle structure θ , the map ψ_θ is clearly unique up to an additive constant.

Claim 6.4 *Let $T'\tilde{V}$ denote the space of oriented tangent directions at points of \tilde{V} , that is, tangent vectors of $T\tilde{V}$ up to multiplication by positive scalars. There exists a continuous map*

$$\Psi_\theta: T'\tilde{V} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$$

that extends ψ_θ , in the sense that for any open edge e of \tilde{V} , if a nonzero tangent vector $[u] \in T'e$ defines on e the orientation of $x \in X$, then $\Psi_\theta([u]) = \psi_\theta(x)$.

We first define Ψ_θ on one triangle τ of \tilde{V} : up to a diffeomorphism, identify τ with some fixed equilateral triangle of the plane. For a vector $[u] \in T'\tau$ (pointing into τ if its root is on $\partial\tau$), let $\alpha([u]) \in [0, 2\pi)$ be its argument: we may assume that $[u]$ is parallel to one of the three sides of τ if and only if $\alpha([u]) \in \frac{\pi}{3}\mathbb{Z}$. For such $[u]$, declare $\Psi_\theta([u]) = \psi_\theta(x)$ where x is the (oriented) edge parallel to $[u]$. For other $[u]$, use interpolation, setting $\Psi_\theta([u]) = f_\theta \circ \alpha([u])$ for some continuous function f_θ such that the net variation of f_θ between two consecutive multiples of $\frac{\pi}{3}$ is θ_i or θ_j or θ_k (the (θ) -angles of τ). Note that the space of such interpolating maps f_θ is contractible. This definition is continuous over $T'\tau$ (that is, over all arguments) because $\theta_i + \theta_j + \theta_k + \theta_i + \theta_j + \theta_k = 2\pi \equiv 0 [2\pi]$.

Similar definitions for other triangles τ of \tilde{V} fit together consistently. Moreover, as the direction $[u_t]$ rotates by a full turn (counterclockwise) around a chosen basepoint p of \tilde{V} , the net variation of $\Psi_\theta([u_t])$ is precisely 2π . This follows from the definition for p in the interior of a triangle τ ; from $\theta_i + \theta_j + \theta_k = \pi$ when p is in the interior of an edge; and from the fact that the (θ) -angles add up to 2π around any vertex of V when p is a vertex of \tilde{V} . Claim 6.4 is proved.

The spaces $T'V$ (oriented tangent directions to V) and $T'\tilde{V}$ have the same universal cover $\widehat{T'V}$, and $\mathbb{R}/2\pi\mathbb{Z}$ has universal cover \mathbb{R} . Therefore Ψ_θ lifts to $\tilde{\Psi}_\theta: \widehat{T'V} \rightarrow \mathbb{R}$, and by computing net variation of $\tilde{\Psi}_\theta$ along a path, this induces a map

$$\bar{h}_\theta: \pi_1(T'V) \rightarrow \mathbb{R},$$

which is clearly a representation. Its values are linear combinations of the θ_i (and of 2π) with integer coefficients; therefore it does not depend on the choices made so far (for example, of the function f_θ). Moreover, any non-vanishing vector field on V (for instance, a constant vector field in a flat metric) defines a direction field, which provides every loop γ in V with a canonical lift $\bar{\gamma}$ to $T'V$. The value $\bar{h}_\theta(\bar{\gamma})$ does not depend on the choice of direction field, because the angular difference between two such fields is just a scalar function on V . Therefore \bar{h}_θ descends to a representation

$$\hat{h}_\theta: \pi_1(V) = H_1(V, \mathbb{Z}) \rightarrow \mathbb{R}.$$

It remains to relate the representation \hat{h}_θ defined via $T'V$ to the angular holonomy h_θ defined via turning angles. The connection is as follows. Let γ be a smooth, embedded, normal curve in V . Then the torus V can be foliated by “pushoff” curves isotopic to γ . The tangent directions to these pushoff curves define a nonzero direction field on V . By Claims 6.3 and 6.4, the tangent direction to γ changes by exactly $\mathcal{A}_\gamma - \mathcal{B}_\gamma$ as we walk around γ . Therefore,

$$\hat{h}_\theta([\gamma]) = \bar{h}_\theta(\bar{\gamma}) = \mathcal{A}_\gamma - \mathcal{B}_\gamma = t_\theta(\gamma).$$

But we have already shown that \hat{h}_θ is a representation, that is, it depends only on the (primitive) homology class $[\gamma]$. Therefore, the turning angle $t_\theta(\gamma)$ of the curve γ depends only on the homology class, which implies the holonomy $h_\theta([\gamma])$ is well-defined. This definition of holonomy for primitive classes extends linearly to all of $H_1(V, \mathbb{Z})$, via equation (6-2).

Finally, the linearity of h_θ follows immediately because $h_\theta = \hat{h}_\theta$ is a representation and \mathbb{R} is commutative. □

6.2 How to deform the holonomy

Lemma 6.5 *The angular holonomy of a veering taut angle structure is always $h_{\text{taut}} = 0$.*

Proof This can be seen from Figure 3. Indeed, a curve γ parallel to the ladderpoles cuts only thin angles off the triangles it crosses, so $\mathcal{A}_\gamma = \mathcal{B}_\gamma = 0$ and $h_{\text{taut}}(\gamma) = 0$. For more general curves, the main observation is that the triangulation of $\partial\bar{M}$ can be turned into an *oriented* train track, for example, in Figure 3 by orienting all rungs from left to right, all blue ladderpoles downward, and all red ladderpoles upward. This way, at each vertex, the two π -angles separate the incoming edges (branches) from the outgoing ones. Next, consider an oriented curve $\bar{\gamma}$ carried by the train track, and assume for simplicity that $\bar{\gamma}$ consists only of rungs (and visits all ladders in cyclic order from left to right, possibly several times). Since at every vertex of $\bar{\gamma}$ the angles

on either side of $\bar{\gamma}$ sum to π , it is easy to see that $\bar{\gamma}$ can be perturbed to a curve γ transverse to the train track, with trivial holonomy $h_{\text{taut}}(\gamma) = 0$. Since h_{taut} is linear on $H_1(\partial\bar{M}, \mathbb{R})$, we conclude that $h_{\text{taut}} = 0$. \square

For non-veering (even taut) angle structures, one may in general have non-trivial holonomy. The key to creating nonzero holonomy is the following relationship between holonomy and deformations.

Lemma 6.6 (Futer and Guéritaud [5, Lemma 4.4]) *Let (θ) be a generalized angle structure, and let γ and δ be oriented closed curves on ∂M . Then, for all $t \in \mathbb{R}$, the deformation tD^γ has the following effect on the holonomy of δ :*

$$h_{\theta+tD^\gamma}(\delta) = h_\theta(\delta) + 2t \cdot \iota(\gamma, \delta)$$

where ι denotes algebraic intersection number. (For our purposes, the sign convention in $\iota(\gamma, \delta)$ will be irrelevant.)

One immediate consequence of Lemma 6.6 is that if γ is homologically trivial, or γ belongs to a different boundary torus than δ , then D^γ does not affect the holonomy of δ at all. Thus, because $h_{\text{taut}} = 0$ and all the curves γ used in Sections 4 and 5 are trivial, it follows that all the angle structures constructed so far have $h_\theta = 0$. To obtain non-trivial holonomy, one must deform along homologically non-trivial curves.

6.3 Holonomy of the rung direction

For example, consider a cusp with $2k$ ladders ($k \geq 1$). Let δ be a closed curve that has intersection number 1 with the slope of the ladders \mathcal{L}_i , and let $\gamma_1, \dots, \gamma_k$ be consistently oriented curves which travel up *every other* ladder (that is, non-adjacent ladders), so that whenever some γ_s traverses a triangle, it enters through one of the $0 - \pi$ edges and exits through the $0 - 0$ edge. Then, the deformation

$$D := \sum_{s=1}^k \lambda_s D^{\gamma_s}$$

does not exit the space $\overline{\mathcal{A}(\tau)}$ of nonnegative angle structures on τ , for small nonnegative λ_s (it decreases the π 's and increases the 0 's). In fact we can take $\lambda_s = \pi/4$ for all s , because each tetrahedron will suffer at most 4 deformations (the γ_s can cross each of its 4 cusp triangles at most once). Moreover, if we take the γ_s in the *other* set of k ladders and reverse their orientation, the same construction works. Moreover still, we

can choose to do this on all cusps simultaneously. By Lemma 6.6, the deformation D perturbs the holonomy of δ by

$$2 \cdot \iota(\delta, \sum_{s=1}^k \lambda_s \gamma_s) = \pm \frac{k\pi}{2},$$

the sign depending on the choice of ladder set in the cusp containing δ . We can summarize the construction in the following result:

Proposition 6.7 *Let T_1, \dots, T_c be the cusps of M , carrying $2k_1, \dots, 2k_c$ ladders (annuli) respectively. Choose $\varepsilon_1, \dots, \varepsilon_c \in \{-1, 1\}$. Let $\delta_1, \dots, \delta_c$ be homology classes in $H_1(\partial\bar{M}, \mathbb{Z}) \simeq \mathbb{Z}^{2c}$ that are \mathbb{Z}^2 -complements of the slopes of the annuli in T_1, \dots, T_c respectively. Then, there exists a nonnegative angle structure (θ) such that*

$$(6-4) \quad h_\theta(\delta_i) = \frac{\varepsilon_i k_i \pi}{2}$$

for every $i \in \{1, \dots, c\}$.

Recall that the space $\mathcal{A}(\tau)$ of positive angle structures on τ is a convex polytope, whose closure $\overline{\mathcal{A}(\tau)}$ contains every non-negative angle structure. Thus there are (strictly) positive angle structures with holonomy arbitrarily close to $\varepsilon_i k_i \pi / 2$, for $\varepsilon_i \in \{-1, 1\}$. In fact, by convexity of $\overline{\mathcal{A}(\tau)}$, one can actually take $\varepsilon_i \in [-1, 1]$ in equation (6-4), getting a full Cartesian product Π of possible holonomies $(h_\theta(\delta_1), \dots, h_\theta(\delta_c))$.

A little more can be said if the triangulation of M is transverse-taut. Recall from Observation 2.8 that ascending and descending ladders then alternate in each cusp of M , and that each truncated tetrahedron then has precisely two boundary triangles in ascending ladders, and two boundary triangles in descending ladders.

Therefore, if we take all curves γ_i along ladders of the *same* type (a notion consistent across all cusps), we can make sure each tetrahedron suffers at most two (not four) deformations, and therefore choose $\lambda_s = \frac{\pi}{2}$ instead of $\frac{\pi}{4}$ in the definition of the deformation D above. However, we can no longer choose the ladder set independently inside each cusp. This is summarized in the following Proposition.

Proposition 6.8 *Suppose M has a veering, transverse-taut triangulation with cusp tori T_1, \dots, T_c carrying $2k_1, \dots, 2k_c$ ladders respectively. Take $\varepsilon \in \{-1, 1\}$, and let $\delta_1, \dots, \delta_c$ be homology classes in $H_1(\partial\bar{M}, \mathbb{Z}) \simeq \mathbb{Z}^{2c}$ that have intersection number ε with the upwards-oriented ladderpoles in T_1, \dots, T_c respectively. Then, for any subset $J \subset \{1, \dots, c\}$, there exists a nonnegative angle structure (θ) such that*

$$h_\theta(\delta_i) = \begin{cases} \varepsilon k_i \pi & \text{if } i \in J; \\ 0 & \text{otherwise.} \end{cases}$$

In summary, Proposition 6.7 states that the c -tuple $(h_\theta(\delta_1), \dots, h_\theta(\delta_c))$ can take any value inside some parallelepiped Π centered around the origin of \mathbb{R}^c , and Proposition 6.8 states that in the transverse–taut case, one pair of opposite octants of Π can be further homothetized by a factor of 2. Therefore, in the transverse–taut case, Proposition 6.8 implies Proposition 6.7 by averaging out. See Figure 7.

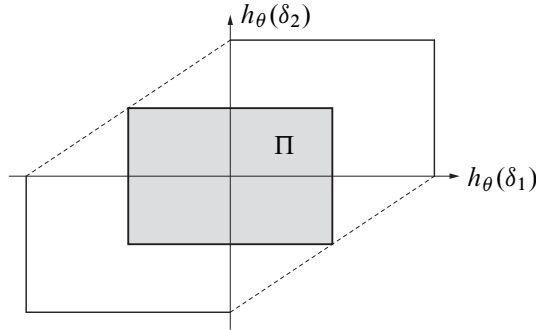


Figure 7: Illustration of Propositions 6.7 and 6.8.

Example 6.9 For once-punctured torus bundles, the bounds of Proposition 6.8 can be seen to be optimal. However, we have found layered triangulations with two cusps, where all points provided by Proposition 6.8 (for $J \neq 0$) belong to the boundary of the space of achievable holonomies, but where each individual $h_\theta(\delta_1)$ or $h_\theta(\delta_2)$ can take even larger values.

Consider the punctured rectangle

$$[0, 2] \times [0, 1] \setminus \{0, 1, 2\} \times \{0, 1\}$$

and identify opposite edges to get a twice-punctured torus T . The linear map

$$\varphi_0 := \begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}^3 = \begin{bmatrix} 5 & 8 \\ 8 & 13 \end{bmatrix}$$

preserves $2\mathbb{Z} \oplus \mathbb{Z}$ and therefore induces a mapping class of T , still denoted φ_0 and preserving each puncture. The mapping torus $T \times [0, 1] / \sim_{\varphi_0}$ is a 6-fold cover of the once-punctured torus bundle with monodromy $\begin{bmatrix} 1 & 1 \\ 1 & 2 \end{bmatrix}$ (the figure–8 knot complement) and as such receives a natural layered triangulation with 12 tetrahedra. We alter this construction by defining

$$\varphi = \varphi_0 \circ \lambda ,$$

where λ is a Dehn twist along the curve $\ell = \{\frac{3}{2}\} \times [0, 1]$ of T . Since ℓ traverses just one tetrahedron, it turns out we can account for λ by just one extra diagonal exchange,

so $M := T \times [0, 1]/\varphi$ receives a natural layered (in fact, veering) triangulation with 13 tetrahedra. The number of pairs of annuli is $k_i = 2$ for both cusps. Numerical computation shows that if δ_1, δ_2 are consistently oriented loops around the two punctures, then the space of pairs $(h_\theta(\delta_1), h_\theta(\delta_2))$, as (θ) runs over all nonnegative angle structures, is the convex hull Q of the six points

$$\pm\left(\frac{27\pi}{13}, \frac{\pi}{13}\right), \quad \pm(2\pi, 2\pi), \quad \pm\left(\frac{\pi}{13}, \frac{27\pi}{13}\right).$$

In particular, the points $\pm(2\pi, 2\pi)$, $\pm(0, 2\pi)$ and $\pm(2\pi, 0)$ that are provided by Proposition 6.8 all belong to the boundary of Q , but of course $\frac{27\pi}{13} > 2\pi$.

6.4 Exotic taut angle structures

If one takes $J = \{1, \dots, c\}$ in Proposition 6.8, then every tetrahedron Δ of M suffers *precisely* two deformations, and these deformations are in the same direction. This is true because whatever the colors of the diagonals in Figure 2, the triangles cut off at A and C (resp. B and D) receive precisely the same colors on all their edges and vertices. As a result, after deformation with $\lambda_s \equiv \frac{\pi}{2}$, each tetrahedron Δ is flat again. Since we can choose $\varepsilon = 1$ or $\varepsilon = -1$ in Proposition 6.8, it follows that

Proposition 6.10 *Every veering, transverse–taut triangulation comes with at least two “exotic” taut angle structures, distinct from the given one (and usually not veering).*

We have also checked that some, but not all, punctured-torus bundles can admit even more taut angle structures: this happens precisely when the monodromy, as a cyclic word in two letters $R = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ and $L = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$, can be decomposed into a product of terms of the form $(RL^*R)(LR^*L)$, where the stars denote arbitrary nonnegative exponents.

6.5 Holonomy of the ladderpole direction

We will find nonnegative angle structures (θ) that realize large values of $h_\theta([\ell])$, where $[\ell] \in H_1(\partial\bar{M}, \mathbb{R})$ is represented by a curve along the ladderpole direction. By the deformation formula of Lemma 6.6, this will involve applying deformations D^γ for a curve γ that intersects the ladderpoles essentially.

Proposition 6.11 *Let M be a manifold with c cusps, endowed with a veering triangulation τ . Let ℓ_1, \dots, ℓ_c be simple closed curves in $\partial\bar{M}$ along the ladderpole directions (with any orientations). Then there exists a nonnegative angle structure (θ) on M such that $h_\theta([\ell_i]) = \frac{\pi}{4}$ for all $1 \leq i \leq c$.*

Proposition 6.12 *Suppose in addition that the triangulation of M is transverse–taut and that the orientations of the curves ℓ_i all agree (or all disagree) with the transverse–taut structure. Then for any subset J of $\{1, \dots, c\}$ there exists a nonnegative angle structure (θ) such that $h_\theta([\ell_i]) = \frac{\pi}{2}$ if $i \in J$ and 0 otherwise.*

Note that just as with Propositions 6.7 and 6.8, there exists a *positive* angle structure $(\theta) \in \mathcal{A}(\tau)$ with holonomy arbitrarily close to the values specified above.

Again, Figure 7 illustrates the relationship between Propositions 6.11 and 6.12. For punctured torus bundles, it is easy to see that Proposition 6.12 gives the optimal bound when the monodromy has the form $R^a L^b$ for positive integers a, b . However, for $R^{a_1} L^{b_1} \dots R^{a_s} L^{b_s}$ the optimal bound becomes $s \frac{\pi}{2}$.

Proof of Proposition 6.11 We start by applying Lemma 4.1 with the extremal value $t = \frac{\pi}{4}$. The angles of the resulting nonnegative angle structure (θ) are now as follows:

- Hinge triangles have angle 0 at the wide angle, and $\frac{\pi}{2}$ elsewhere.
- Nonhinge triangles have angle π at the thin angle adjacent to the ladderpole segment, and 0 elsewhere.

On any cusp being considered, the chosen orientation of ℓ , referred to as “upward”, induces a partition of the $2k$ ladders into two classes: namely, call *ascending* those ladders whose triangles have their wide angle *above* the base for the orientation of ℓ , and *descending* the other ladders. (In the absence of a transverse–taut structure, this does not have to be consistent over all c cusps: a tetrahedron may have all its boundary triangles in descending ladders.)

We want to apply the deformation tD^γ to the angle structure (θ) , for a carefully chosen curve γ . Focus on one cusp. Start γ by crossing from an ascending ladder \mathcal{L} to a descending one \mathcal{L}' , at a hinge of \mathcal{L} . Then, let γ travel *down* \mathcal{L}' until it enters a hinge, then cross over to the next ascending ladder \mathcal{L}'' . Let γ travel *up* \mathcal{L}'' until the first hinge, and so on. See Figure 8. Notice that ladders are travelled in the opposite direction compared to Lemma 4.2.

If we orient all ladderpoles consistently with ℓ , then this curve γ always intersects ladderpoles from the same side, and will eventually close up (possibly not at its starting point, in which case we just drop the initial dead arc of γ). It is easy to check that the deformation D^γ only decreases positive angles for (θ) ; moreover we can build one such curve γ_i per cusp, crossing the ladderpoles in either direction. Since γ always crosses the ladderpoles of each cusp in the same direction, $\iota(\gamma, \ell_i) = m_i \geq 1$.

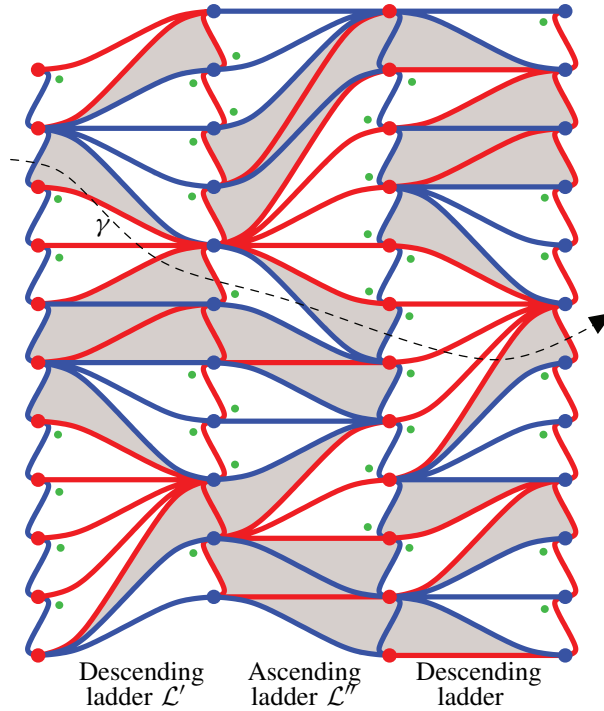


Figure 8: The same triangulation as in Figure 3, with a horizontal deformation curve γ . Green dots indicate angles of π in nonhinge triangles for (θ) ; hinge triangles have angles $0, \frac{\pi}{2}, \frac{\pi}{2}$ with the 0 in the wide corner.

We claim that

$$(\theta') = (\theta) + \sum_{i=1}^c \frac{\pi}{8} D^{\gamma_i}$$

is a nonnegative angle structure. To see this, just notice that each of the $4n$ triangles of $\partial \bar{M}$ is crossed at most once by the union of the γ_i . As a result, any angle that gets decreased (and was therefore at least $\frac{\pi}{2}$ for (θ)) is decreased by at most $4 \cdot \frac{\pi}{8}$, and thus stays nonnegative. By Lemma 6.6, $h_{\theta'}([\ell_i]) = \frac{\pi}{4} m_i$, completing the proof. \square

Proof of Proposition 6.12 If in addition M is transverse-taut and the ℓ_i are consistently oriented, we claim that

$$(\theta) + \sum_{i \in J} \frac{\pi}{4} D^{\gamma_i}$$

is still nonnegative. Of course, $\bigcup_{i \in J} \gamma_i$ still crosses each triangle of $\partial \bar{M}$ at most once.

The claim is easy to check for a nonhinge tetrahedron Δ : the π -angle of Δ gets decreased at most four times by $\frac{\pi}{4}$ (once per cusp triangle of Δ) while the other angles only get increased.

For Δ a hinge tetrahedron, we must discuss two possible cases. Note that (without loss of generality, up to exchanging colors), the curves γ_i always cross from one ladder to the next at a hinge tetrahedron Δ with the upper diagonal red and the lower diagonal blue (this is true in Figure 8; recall the diagonals of a tetrahedron are seen as wide vertices of the corresponding cusp triangles). Call such tetrahedra *hinges of type 1*. When γ_i exits a *descending* ladder, it decreases the $\frac{\pi}{2}$ -angle *clockwise* from the wide angle in a hinge of type 1. When γ_i exits an *ascending* ladder, it decreases the $\frac{\pi}{2}$ -angle *counterclockwise* from the wide angle in a hinge of type 1. Moreover, $\bigcup_{i \in J} \gamma_i$ visits hinges of type 1 *only* as the γ_i leave a ladder for an adjacent one. Therefore, since Δ has only two tips in ascending (resp. descending) ladders, each $\frac{\pi}{2}$ -angle of Δ is decreased at most twice by $\frac{\pi}{4}$ and thus stays nonnegative.

For a hinge tetrahedron Δ' of the other type (upper diagonal blue and lower diagonal red), the argument is similar: cusp triangles of Δ' are only (possibly) visited by $\bigcup_{i \in J} \gamma_i$ as γ_i *enters* a new ladder (the rest of the time γ_i travels up or down fans); for example in Figure 8 the very first visible triangle crossed by γ is such a hinge. Again, γ_i decreases the $\frac{\pi}{2}$ -angle of Δ' clockwise from the wide angle if it encounters Δ' in a descending ladder, and decreases the other $\frac{\pi}{2}$ -angle of Δ' if it encounters Δ' in an ascending ladder. Since Δ' has exactly two cusp triangles in each type of ladder, Δ' is still nonnegative after applying $\sum_{i \in J} \frac{\pi}{4} D^{\gamma_i}$. \square

Remark 6.13 A little more work would show that the slope of the curve γ (like that of the ladderpoles) is also an invariant of the veering structure, that is, it does not depend on the starting point and starting direction chosen to construct γ . In the layered (veering) case, it would be interesting to relate this slope to the one defined by the fiber.

References

- [1] **I Agol**, *Ideal triangulations of pseudo-Anosov mapping tori*, from: “Topology and geometry in dimension three”, (W Li, L Bartolini, J Johnson, F Luo, R Myers, JH Rubinstein, editors), Contemp. Math. 560, Amer. Math. Soc., Providence, RI (2011) 1–17 MR2866919
- [2] **D Cooper, S Tillmann**, *The Thurston norm via normal surfaces*, Pacific J. Math. 239 (2009) 1–15 MR2449008
- [3] **A Fathi**, *Démonstration d’un théorème de Penner sur la composition des twists de Dehn*, Bull. Soc. Math. France 120 (1992) 467–484 MR1194272

- [4] **D Futer, F Guéritaud**, *Angled decompositions of arborescent link complements*, Proc. Lond. Math. Soc. 98 (2009) 325–364 MR2481951
- [5] **D Futer, F Guéritaud**, *From angled triangulations to hyperbolic structures*, from: “Interactions between hyperbolic geometry, quantum topology and number theory”, (A Champanerkar, O Dasbach, E Kalfagianni, I Kofman, W Neumann, N Stoltzfus, editors), Contemp. Math. 541, Amer. Math. Soc., Providence, RI (2011) 159–182 MR2796632
- [6] **F Guéritaud**, *On canonical triangulations of once-punctured torus bundles and two-bridge link complements*, Geom. Topol. 10 (2006) 1239–1284 MR2255497 With an appendix by David Futer
- [7] **CD Hodgson, JH Rubinstein, H Segerman, S Tillmann**, *Veering triangulations admit strict angle structures*, Geom. Topol. 15 (2011) 2073–2089 MR2860987
- [8] **A Issa**, *Construction of non-geometric veering triangulations of fibered hyperbolic 3-manifolds*, Master’s thesis, University of Melbourne (2012)
- [9] **E Kang, JH Rubinstein**, *Ideal triangulations of 3-manifolds II: Taut and angle structures*, Algebr. Geom. Topol. 5 (2005) 1505–1533 MR2186107
- [10] **M Lackenby**, *Taut ideal triangulations of 3-manifolds*, Geom. Topol. 4 (2000) 369–395 MR1790190
- [11] **M Lackenby**, *Word hyperbolic Dehn surgery*, Invent. Math. 140 (2000) 243–282 MR1756996
- [12] **M Lackenby**, *An algorithm to determine the Heegaard genus of simple 3-manifolds with nonempty boundary*, Algebr. Geom. Topol. 8 (2008) 911–934 MR2443101
- [13] **F Luo, S Tillmann**, *Angle structures and normal surfaces*, Trans. Amer. Math. Soc. 360 (2008) 2849–2866 MR2379778
- [14] **I Rivin**, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. of Math. 139 (1994) 553–580 MR1283870
- [15] **I Rivin**, *Combinatorial optimization in geometry*, Adv. in Appl. Math. 31 (2003) 242–271 MR1985831
- [16] **WP Thurston**, *The geometry and topology of three-manifolds*, Princeton Univ. Math. Dept. Lecture Notes (1980) Available at <http://www.msri.org/gt3m/>

Department of Mathematics, Temple University
Philadelphia, PA 19122, USA

Laboratoire Paul Painlevé, CNRS UMR 8524, Université de Lille 1
59650 Villeneuve d’Ascq, France

dfuter@temple.edu, Francois.Gueritaud@math.univ-lille1.fr

Received: 13 January 2011 Revised: 29 May 2012

