Algebraic & Geometric Topology 20 (2020) 487–501

Essential loops in taut ideal triangulations

SAUL SCHLEIMER
HENRY SEGERMAN

In this note we combinatorialise a technique of Novikov. We use this to prove that, in a three-manifold equipped with a taut ideal triangulation, any vertical or normal loop is nontrivial in the fundamental group.

57M05; 57M20

1 Introduction

The notion of a taut ideal triangulation of a three-manifold is due to Lackenby [7]. He combinatorialised the angle structures introduced independently by Casson and by Rivin [12]. They in turn linearised the geometric triangulations of Thurston [13]. Each of these structures plays an important role in modern low-dimensional topology. In particular, taut ideal triangulations have a strong connection to the subject of taut foliations, introduced by Gabai [5], and to that of taut branched surfaces, due to Oertel [10]. In addition to the results of Lackenby, taut ideal triangulations play a central role in the theory of layered triangulations. One spectacular contribution has been as a prerequisite for Agol's theory of veering triangulations [1].

Novikov [9, Theorem 6.1] gives one of the early applications of foliations to the study of the fundamental group of a manifold. He starts with a loop δ in good position with respect to a foliation $\mathcal{F}$. He further supposes that $H\colon D \to M$ is a null-homotopy of δ, also in good position. Pulling back, he obtains a singular foliation $H^{-1}(\mathcal{F})$ on the disk D. The Poincaré–Hopf theorem gives combinatorial control of the singularities, which translates to topological control over the homotopy. Morally, the positivity of the Euler characteristic of the disk constrains the position of δ. We refer to Candel and Conlon [4, Chapter 9] for a history of the subject and for detailed proofs.

We introduce a combinatorial version of the Novikov technique; instead of pulling back a foliation we pull back a taut ideal triangulation. This gives a train track with stops in

Published: 23 February 2020 DOI: 10.2140/agt.2020.20.487

the disk D. We so obtain a very simple proof of a variant of one of Novikov's results. That is, suppose that M is a three-manifold, equipped with a taut ideal triangulation $\mathcal{T}$. Let $\mathcal{B} = \mathcal{T}^{(2)}$ be the resulting branched surface in M.

Theorem 3.2 *Any loop δ in M which is vertical with respect to $\mathcal{B}$ is nontrivial in $\pi_1(M)$.*

There is also an indirect proof of this using Novikov's original technique — see Calegari [3, Theorem 4.35(3)] — once we observe that $\mathcal{B}$ carries an essential lamination which extends to a taut foliation of M (see Gabai and Oertel [6, Example 5.1] as well as Lackenby [7, page 373]).

Using our techniques we also obtain a new result, as follows:

Theorem 5.1 *Any loop γ in M which is normal with respect to $\mathcal{B}$ is nontrivial in $\pi_1(M)$.*

The proof of Theorem 5.1 is more delicate than that of Theorem 3.2; new behaviour near the boundary of D must be dealt with.

From Theorems 3.2 and 5.1 we deduce that vertical, and also normal, loops are of infinite order in the fundamental group. Note that this is a bit weaker than the conclusion in the comparable situation of a train track τ in a surface — there, loops dual to, or carried by, τ are not only nontrivial but also nonperipheral.

We have a simple corollary of Theorem 5.1. Let $\widetilde{M}$ be the universal cover of M and let $\widetilde{\mathcal{B}}$ be the resulting branched surface.

1.1 Corollary *Suppose that F is a connected surface* (perhaps with boundary) *carried by $\widetilde{\mathcal{B}}$ and realised as a* (perhaps finite) *union of faces of $\widetilde{\mathcal{B}}$. Then F is a disk.* □

Previous work Gabai and Oertel prove that laminations carried by essential branched surfaces are π_1–injective [6, Lemma 2.7]. Our Theorem 5.1 is both more and less general than their work. We do not require a lamination. They do not require the manifold to be cusped.

Calegari [2, Remark 5.6] gives a very different combinatorial version of Theorem 3.2, in the closed case. He introduces the notion of a *local orientation*; this is, in a sense, dual to having a transverse taut branched surface $\mathcal{B} \subset M$ where all components of $M - \mathcal{B}$ are taut balls.

Acknowledgements We thank Marc Lackenby for helpful conversations. Segerman was supported in part by National Science Foundation grant DMS-1708239.

2 Background

Throughout the paper we will use M to denote a compact connected manifold with nonempty boundary. All boundary components will be tori or Klein bottles. Suppose that $\mathcal{T}$ is a three-dimensional triangulation; that is, a collection of *model tetrahedra* and a collection of face pairings. We will also call the faces of a model tetrahedron *model faces*, and similarly for its edges and vertices.

Let $|\mathcal{T}|$ be the quotient space: that is, we take the disjoint union of the model tetrahedra of $\mathcal{T}$ and identify model faces using the face pairings. Let $\mathcal{T}^{(k)}$ be the k–skeleton of $|\mathcal{T}|$. Let $n(\mathcal{T}^{(0)})$ be an open regular neighbourhood of the vertices of $\mathcal{T}$. We call $\mathcal{T}$ a *ideal triangulation* of M if $|\mathcal{T}| - n(\mathcal{T}^{(0)})$ is homeomorphic to M.

A *taut angle structure* on $\mathcal{T}$ is an assignment of dihedral angles, zero or π, to each model edge in $\mathcal{T}$. The assignment is required to obey two conditions. The *edge equalities* state that, for each edge $e \in \mathcal{T}^{(1)}$, the sum of the dihedral angles of its models is 2π. The *triangle equalities* state that, for any model vertex, the sum of the dihedral angles of the three adjacent model edges is π. We say that the tetrahedra of $\mathcal{T}^{(3)}$ are *taut*. See Figure 1, left.

We deduce that every taut tetrahedron has four edges with dihedral angle zero. We call the union of these four edges the *equator* of the taut tetrahedron.

Suppose now that e is an edge of $\mathcal{T}^{(1)}$. There are exactly two model edges for e with angle π; all others are zero. Obeying these dihedral angles, we isotope the two-skeleton $\mathcal{T}^{(2)}$ to obtain a smooth *branched surface* $\mathcal{B}$. See Figure 1, right. Some references would call $\mathcal{B}$ a nongeneric branched surface without vertices. See for example [3, Section 6.3].

2.1 Definition Suppose that δ is a smooth embedded loop in M, transverse to $\mathcal{B}$. Suppose that for every tetrahedron t we have that every arc d of $\delta \cap t$ *links* the equator of t. (That is, the endpoints of d are separated in ∂t by the equator of t.) Then we say that the loop δ is *vertical* with respect to $\mathcal{B}$.

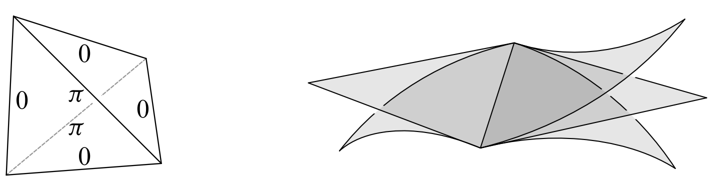

Figure 1: Left: a taut tetrahedron. Right: all faces meeting a single edge in $\mathcal{B}$.

2.2 Definition Suppose that γ is a smooth loop immersed in $\mathcal{B}$ and transverse to the edges of $\mathcal{B}$. Suppose that for every model face f of $\mathcal{B}$ and for every component J of $\gamma^{-1}(f)$, the arc J is *normal* in f. (That is, the endpoints of J lie in distinct edges of f.) Then we say that the loop γ is *normal* with respect to $\mathcal{B}$.

3 Combinatorics of null-homotopies

Suppose that δ is a loop in M which is transverse to the branched surface $\mathcal{B}$. Let $D = D^2$ be the unit disk with the usual orientation. Suppose that $H\colon D \to M$ is a null-homotopy of δ. We homotope H relative to ∂D to make H transverse to $\mathcal{B}$.

We define $\tau = H^{-1}(\mathcal{B})$. Thus τ is a *train track* in D. The *switches* of τ are exactly the points of $H^{-1}(\mathcal{B}^{(1)})$. The *stops* of τ are exactly the points of $(H|_{\partial D})^{-1}(\mathcal{B})$. The standard reference for train tracks is [11]; we also rely on [8]. We note that our track τ does not satisfy the so-called "geometry condition" [11, page 5; 8, page 52].

We call a connected component R of $D - \tau$ a *region*. Let $\mathrm{cusps}(R)$ and $\mathrm{corners}(R)$ count the number of (necessarily outwards) cusps and corners on the boundary of R. As a bit of terminology, we divide ∂R into *sides*: these are the components of ∂R minus all outward cusps and corners. Note that a side s of R may be a union of several branches of τ.

We define the *index* of R to be

$$\mathrm{ind}(R) = \chi(R) - \tfrac{1}{2}\,\mathrm{cusps}(R) - \tfrac{1}{4}\,\mathrm{corners}(R).$$

In Table 1 we give pictures of, and names to, all possible disk regions with nonnegative index. Note that index is additive under taking the union of regions [8, page 57]. Thus the sum of the indices of the regions of $D - \tau$ is exactly $\chi(D)$, that is, one. We deduce from this that there is at least one region R with positive index.

Let $r(H)$ be the number of regions of $D - \tau$. Over all null-homotopies of δ, transverse to $\mathcal{B}$, we choose H to minimise $r(H)$. We call such an H *minimal*.

3.1 Lemma *Suppose that δ is a loop in M transverse to $\mathcal{B}$. Suppose that $H\colon D \to M$ is a minimal null-homotopy of δ. Let $\tau = H^{-1}(\mathcal{B})$. Then we have the following*:

(1) *All regions of $D - \tau$ are disks.*

(2) *If s is a side of a region R, then the interior of s meets at most one switch.*

corners	index		
	1	$\frac{1}{2}$	0
0	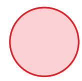	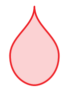	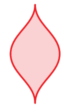
2		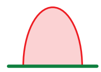	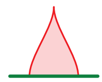
4			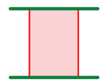

Table 1: Disk regions with nonnegative index, organised by the number of corners. These are named as follows: *nullgon*, *cusped monogon*, *cusped bigon*, *boundary bigon*, *boundary trigon* and *rectangle*.

(3) *No region R of $D-\tau$ is a nullgon.*

(4) *No region R of $D-\tau$ is a cusped monogon.*

Thus, all positive index regions of D are boundary bigons.

Proof (1) If there were a region with topology then we could compress it into the containing tetrahedron and reduce $r(H)$.

(2) Suppose that the interior of s meets at least two switches. All such switches in the interior of s are preimages under H of a single edge. Hence there is a branch $b \subset \tau$ such that $H(b)$ is a nonnormal arc. We homotope H in a neighbourhood of b to make $H(b)$ simple. This done, $H(b)$ cuts a bigon B off of the face containing $H(b)$. We then homotope H across B. This does not increase $r(H)$. If $r(H)$ does not decrease, then this move disconnects τ, and creates a region with topology, contradicting (1).

(3) Suppose that R is a nullgon. If $H(\partial R)$ is disjoint from $\mathcal{B}^{(1)}$ then the region adjacent to R is not a disk, contradicting (1). It follows that ∂R consists of an even number of branches of τ (alternating between the two faces of a tetrahedron t on either side of a π–edge of t). But this contradicts (2).

(4) Suppose that t is the taut tetrahedron containing $H(R)$. Let s be the boundary of R. We deduce that the loop s crosses the equator of t exactly once, a contradiction.

Since there are no nullgons or monogons, the only possible positive-index regions are boundary bigons. □

Equipped with this we can now prove the following:

3.2 Theorem *Any loop δ in M which is vertical with respect to $\mathcal{B}$ is nontrivial in $\pi_1(M)$.*

Proof Suppose that $H\colon D \to M$ is a minimal null-homotopy of the vertical loop δ. Applying Lemma 3.1, there must be a region R of $D - \tau$ which is a boundary bigon. Let t be the tetrahedron containing $H(R)$. Let $d = \partial R \cap \partial D$ and let $s = \partial R - d^\circ$. From the definition of vertical, we have that $H(d)$ links the equator of t. Therefore $H(s)$ crosses the equator of t an odd number of times, and thus at least once. This contradicts the fact that ∂R has no cusps. □

4 Transverse taut

In order to prove Theorem 5.1, we will use the following strengthening of the notion of a taut structure. A *transverse taut structure* on $\mathcal{T}$ is a taut structure together with a coorientation on $\mathcal{B}$ with the following property. If model faces f and f' of a model tetrahedron t share a common model edge e, then

- the edge e is part of the equator of t if and only if exactly one of the coorientations on f and f' points into t.

See Figure 2, left. It follows that the coorientations on faces incident to an edge change direction precisely twice as we go around an edge. See Figure 2, right.

Suppose that $\mathcal{T}$ is an ideal triangulation of a manifold M equipped with a taut structure. We now construct a triangulation $\widetilde{\mathcal{T}}$ of a double cover $\widetilde{M}$ of M. By construction, the lift of the taut structure on $\mathcal{T}$ to $\widetilde{\mathcal{T}}$ will support a transverse taut structure.

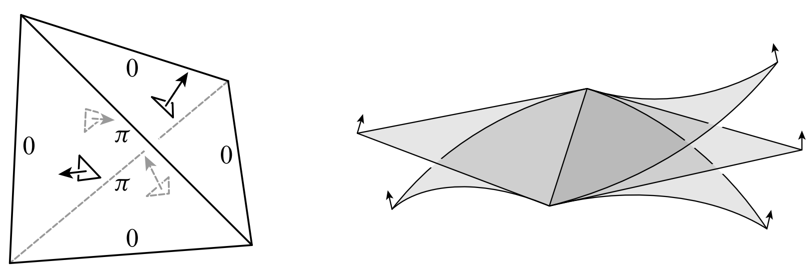

Figure 2: Left: coorientations and angles in a transverse taut tetrahedron. Right: coorientations around an edge.

For each taut tetrahedron t of $\mathcal{T}$, we arbitrarily label the two model edges with dihedral angle π as e' and e''. In $\widetilde{\mathcal{T}}$, we have two taut tetrahedra t' and t'' corresponding to t. We assign a coorientation to the model faces of t' and t'' in such a way that the coorientation points into the tetrahedron on the two model faces of t^* incident to e^*. Now suppose that t_i and t_j are tetrahedra of $\mathcal{T}$, glued to each other along model faces f_i and f_j. In $\widetilde{\mathcal{T}}$ we have tetrahedra t_i', t_i'', t_j' and t_j'', with model faces f_i', f_i'', f_j' and f_j'', respectively.

We glue t_i' to either t_j' or t_j'' as the coorientation on f_i' agrees with f_j' or f_j''. We similarly glue t_i'' to the remaining copy of t_j. Having made all such gluings, the resulting triangulation $\widetilde{\mathcal{T}}$ has a transverse taut structure by construction. It has one component if and only if the taut structure on $\mathcal{T}$ does not support a transverse taut structure.

5 Proof of the main result

5.1 Theorem *Any loop γ in M which is normal with respect to $\mathcal{B}$ is nontrivial in $\pi_1(M)$.*

Proof Suppose, for a contradiction, that the normal loop γ is null-homotopic. Thus γ lifts to a normal loop in any cover. Thus, without loss of generality, we may assume that the taut structure on $\mathcal{T}$ supports a transverse taut structure. This gives us a local notion of *upwards*. In particular, every model tetrahedron has two lower faces and two upper faces, separated by its equator.

Lemma 3.1 does not apply directly to a normal loop γ. So, let A be a model annulus with horizontal boundary circles $\partial_0 A \sqcup \partial_1 A$. Let G be a small smooth homotopy $G\colon A \to M$, moving γ slightly upwards. That is, $G(\partial_0 A) = \gamma$ and we define $\delta = G(\partial_1 A)$. We ensure that G is transverse to $\mathcal{B}$ away from $\partial_0 A$; also, we arrange that for each vertical interval J in A the tangents to $G(J)$ point upwards. We will apply Lemma 3.1 to δ.

We call δ a *raised curve*. We call the components of $\delta - \mathcal{B}$ *raised arcs*. There are six *types* of raised arc. These are shown in Figure 3. There is a cellulation of A with one-skeleton $\partial A \cup G^{-1}(\mathcal{B})$. Suppose that C is a two-cell. Let $c = C \cap \partial_0 A$ and $d = C \cap \partial_1 A$. Thus $G(c) \subset \gamma$ and $G(d) \subset \delta$. We say that $G(c)$ is the *lowering* of the raised arc $G(d)$. We record this by the *lowering map*, L, where $L(G(d)) = G(c)$. Note that $G(c)$ may be either a single vertex, a single normal arc or two normal arcs. Again, see Figure 3.

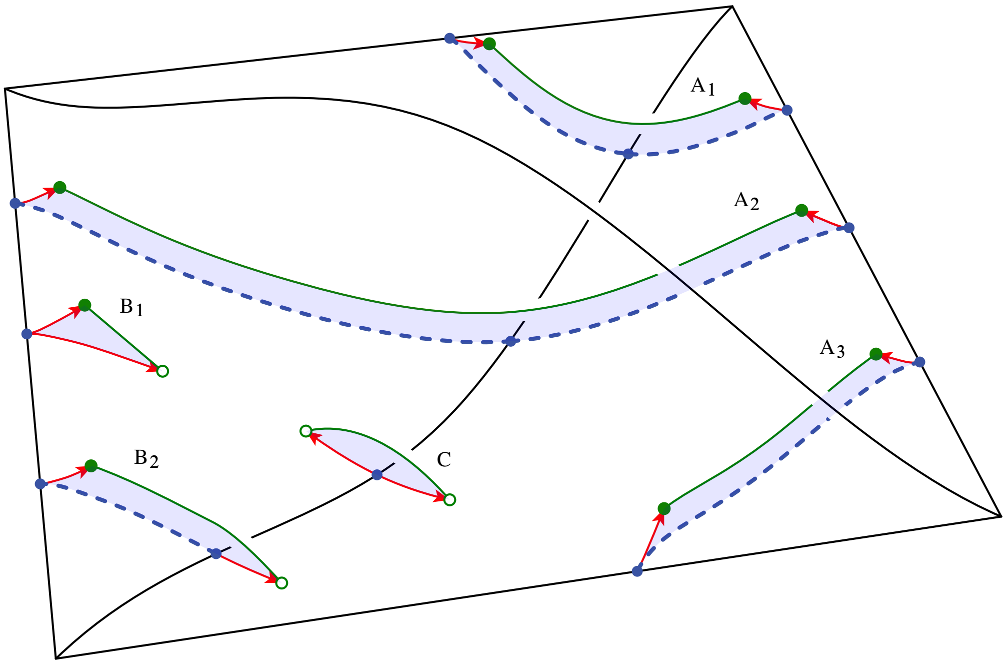

Figure 3: A taut tetrahedron containing the six possible types of raised arcs of δ. These are drawn in (solid) green. The normal arcs (or points) of γ, namely the lowerings of the raised arcs, are drawn in (dashed) blue. Images of the two-cells of the annulus A are shaded in light blue. Filled green dots indicate endpoints of raised arcs on the top two faces of the tetrahedron; open green dots indicate endpoints on the bottom two faces.

Suppose that $H\colon D \to M$ is a minimal null-homotopy of δ. Recall that $\tau = H^{-1}(\mathcal{B})$. Applying Lemma 3.1 implies that $D - \tau$ has at least two boundary bigons. Applying another small homotopy, we can retain minimality and also make H transverse to γ.

Pulling back the transverse taut structure on $\mathcal{B}$ by H gives a transverse orientation on the branches of τ which is consistent across switches. Thus, for any region R of $D - \tau$ and for any side s of R, the transverse orientation on s points either into or out of R. This gives us a classification of boundary bigons. Suppose that R is a boundary bigon and $s = \partial R - \partial D$ is its side in τ. If the transverse orientation on s points out of R then we call R a *min-bigon*. If it points into R, we call R a *max-bigon*.

5.2 Min-bigons Suppose that R is a min-bigon. We move γ up, across $H(R)$, to obtain γ'. We appeal to Lemma 3.1(2) to ensure that γ' is normal. Let δ' be the corresponding raised loop and let H' be the new null-homotopy. See Figure 4.

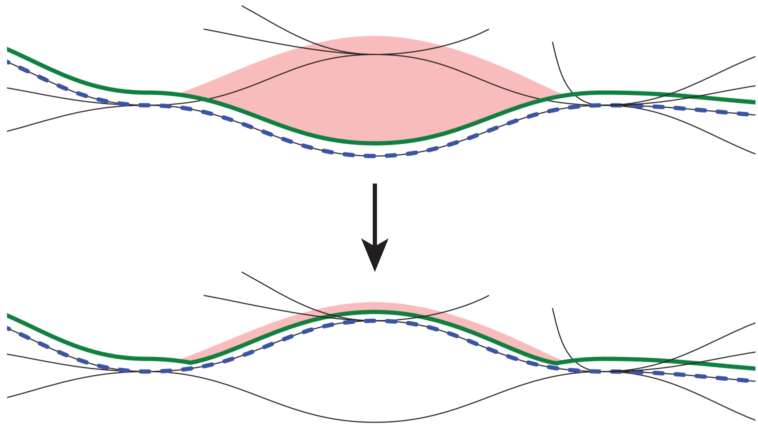

Figure 4: Pushing over a min-bigon of type A_3.

The loop γ' may be shorter than, the same length as, or longer than γ (see types $\mathrm{A}_1, \mathrm{A}_2$ and A_3 in Figure 3). However, H' has exactly one fewer region. That is, $r(H') = r(H) - 1$. We repeat this process until there are no more min-bigons.

5.3 Max-bigons Suppose that R_0 is a max-bigon. Unlike the situation of a min-bigon, a max-bigon does not give us a simple move to reduce complexity. The asymmetry stems from the fact that we raised γ rather than lowered it. Instead, our plan is to uniquely associate to R_0 two small subregions of $D - \tau$, each with index $-\frac{1}{4}$. This will imply that the index of D is at most zero. This contradiction finally proves Theorem 5.1.

We begin as follows. Let s be the side of R_0 in τ. Let $d_0 = \partial R_0 - s \subset \partial D$. We give d_0 the (tangential) orientation it receives from D. In Figure 5, this orientation will point left. Note that $H(d_0) \subset \delta$ is a raised arc. Let $c_0 = L(H(d_0))$ be its lowering.

5.4 Claim

- *The raised arc $H(d_0)$ has type* C.
- *The side s meets exactly one switch c_0' of τ.*
- *The vertices c_0 and $H(c_0')$ cobound a subedge $\epsilon_0 \subset \mathcal{B}^{(1)}$.*

Proof Let t_0 be the tetrahedron containing $H(R_0)$. By the definition of a max-bigon, the transverse orientation on s points into R_0. Thus each corner of $H(R_0)$ is contained in a lower face of t_0. Consulting Figure 3 we deduce that $H(d_0)$ is of type C. Thus each corner of $H(R_0)$ is contained in its own lower face of t_0. We deduce that s meets at least one switch of τ. By Lemma 3.1(2) the side s meets exactly one switch, which we call c_0'.

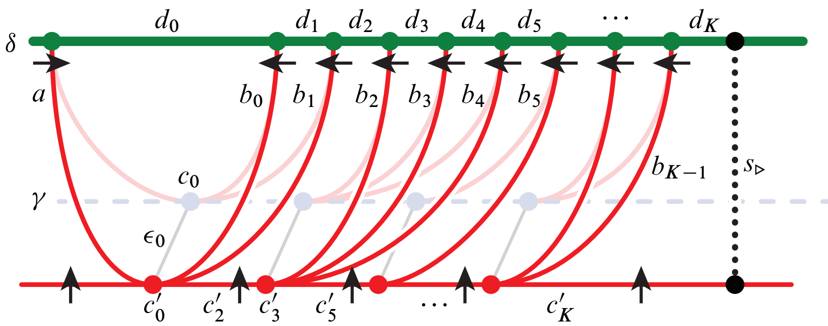

Figure 5: A possible picture of part of the annulus A (in back), the bigons B_k (in front) and the homotopies F_k (bottom). To lighten the notation in this figure, we have omitted applying H to labels of subsets of D. Transverse orientations on the branches b_i are shown with arrows. Note that $\epsilon_0 = \epsilon_1$, $c'_0 = c'_1$, $c'_3 = c'_4$, and so on.

Since H is transverse to γ, the vertices c_0 and $H(c'_0)$ are distinct. They are contained in the same edge of $\mathcal{B}^{(1)}$, namely the bottom edge e_0 of t_0. In e_0 they cobound a subedge, which we call ϵ_0. □

Let a and b_0 be the components of $s - c'_0$, where b_0 meets the right endpoint of d_0. See the leftmost region on Figure 5.

Now consider a sequence of regions $R_0, R_1, \ldots, R_n$ that meet ∂D in the sides $d_0, d_1, \ldots d_n$ as we move along ∂D to the right. Define $c_k = L(H(d_k)) \subset \gamma$, the lowering of the raised arc $H(d_k)$. Define $\gamma_k = \bigcup_{i=0}^{k} c_i \subset \gamma$.

Let b_i be the branch of τ that meets ∂D at the right corner of R_i. We now choose $n = N$ such that b_{i-1} and b_i have the same transverse orientation for $1 \le i < N$, while they have opposite transverse orientations for $i = N$. Thus, R_N is at a *local minimum* of γ, and we are going downhill to it from the *local maximum* at R_0. This downhill condition implies that for $i \in [1, N)$, the raised arc $H(d_i)$ is of type either B_1 or B_2. Again, see Figure 3.

Recall that all positive index regions are now max-bigons. Thus none of the R_i can have positive index for $i > 0$. Let K be the smallest number for which R_K has negative index, or, if there is none, then set $K = N$.

5.5 Claim *The region R_K is not a boundary trigon.*

Proof If $K < N$ then, by definition, R_K has negative index, and so is not a boundary trigon. If $K = N$ then R_N cannot be a boundary trigon since the transverse orientations on the two sides of a boundary trigon must agree, yet R_N is at a local minimum of γ. □

For all $k \in [0, K)$ we define the union $B_k = \bigcup_{i=0}^{k} R_i$. Define $\gamma'_k = \partial B_k - (a^\circ \cup b_k^\circ \cup \partial D)$. (Unlike in Section 5.2, here γ'_k is a push-off of only a section of γ.)

5.6 Definition Suppose that $g, h\colon [0,1] \to \mathcal{B}$ are paths. Suppose $F\colon [0,1] \times [0,1] \to \mathcal{B}$ is a homotopy from g to h. Thus $g(x) = F(x,0)$ and $h(x) = F(x,1)$. We say that F is *transverse* if whenever $F(x_0, t_0)$ is contained in a (1– or 2–) cell C of $\mathcal{B}$, we have that the *trace* $F(x_0, [0,1])$ lies in C.

5.7 Claim *For all $k \in [1, K)$:*

(1) *The region R_k is a boundary trigon.*

(2) *The union B_k has exactly two corners and no cusps.*

(3) *There is a transverse homotopy F_k taking γ_k to $H(\gamma'_k)$.*

Proof We will prove this by induction. Claim 5.4 implies the base case (for $k = 1$) in a manner essentially identical to the general inductive step, so we omit its proof.

Suppose that the hypotheses hold at step k. Recall that $H(d_k)$ has type B_1 or B_2, so it has precisely one lower endpoint. Let f_k be the face that contains the lower endpoint. Let p be the endpoint of γ_k, and let e_k be the edge of f_k containing p. Let β be the normal arc of γ immediately after p. Let f_β be the face containing β. Viewed in a small neighbourhood of e_k, the faces f_β and f_k are on the same side (say the right side) of e_k, and f_β is below f_k.

Let p' be the endpoint of γ'_k meeting b_k. By hypothesis (3), the transverse homotopy F_k takes p to $H(p')$, with trace lying in e_k. Since H is transverse to e_k at $H(p')$, we deduce that $H(D)$ meets both f_k and f_β at $H(p')$. Thus p' is a switch of τ with a cusp immediately below b_k, to the right of p', pointing at γ'_k (which extends to the left of p'). This cusp lies in R_{k+1}, since b_k is part of the boundary of R_{k+1}. See Figure 5.

If R_{k+1} has negative index then $k+1 = K$ and we have nothing to prove. So suppose that R_{k+1} has index zero. Consulting Table 1 we deduce that R_{k+1} is a boundary trigon. This proves hypothesis (1). Note that hypothesis (2) follows because B_k meets R_{k+1} along b_k.

Let $s_{k+1} = \partial R_{k+1} - (d_{k+1} \cup b_k^\circ)$ be the remaining side of the boundary trigon R_{k+1}. By Lemma 3.1(2) there is at most one switch in the interior of s_{k+1}. Let $c'_{k+1} = s_{k+1} - b_{k+1}^\circ - \partial D$. Note that $\gamma'_{k+1} = \gamma'_k \cup c'_{k+1}$.

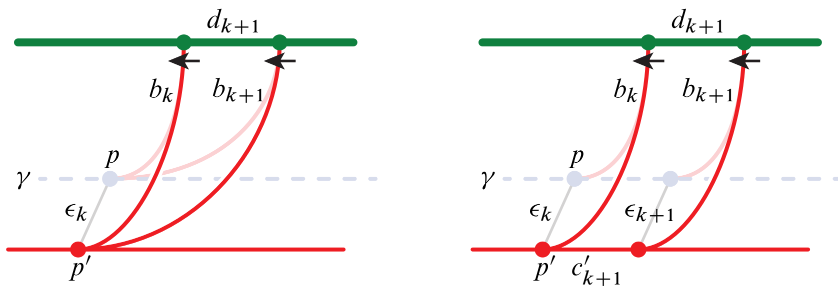

Figure 6: Extending the transverse homotopy F_k when $H(d_{k+1})$ has type B_1 (left) and type B_2 (right). As in Figure 5, we have omitted applying H to labels of subsets of D.

The path $H(s_{k+1})$ has endpoints $H(p')$ and the lower endpoint of $H(d_{k+1})$. The point $H(p')$ lies on the edge e_k. Recall that f_{k+1} is the face containing the lower endpoint of $H(d_{k+1})$. There are two cases, depending on the type of $H(d_{k+1})$:

- Suppose that $H(d_{k+1})$ has type B_1. Then $\gamma_{k+1} = \gamma_k$. In this case, e_k is a boundary edge of f_{k+1}. Since there is at most one switch in the interior of s_{k+1}, there are in fact no such switches. So $s_{k+1} = b_{k+1}$ and c'_{k+1} is a single switch, equal to p'. We deduce that $\gamma'_{k+1} = \gamma'_k$. Since $\gamma_{k+1} = \gamma_k$ and $\gamma'_{k+1} = \gamma'_k$, we set $F_{k+1} = F_k$. See Figure 6, left.
- Suppose that $H(d_{k+1})$ has type B_2. Then $\gamma_{k+1} = \gamma_k \cup c_{k+1}$. Let t_{k+1} be the tetrahedron containing $H(R_{k+1})$. In this case, the path $H(s_{k+1})$ must cross the bottom edge of t_{k+1} in order to get into f_{k+1}. Since there is at most one switch in the interior of s_{k+1}, there is exactly one. Let f be the other lower face of t_{k+1}. Thus c_{k+1} is a normal arc in f. Note that $H(c'_{k+1})$ is a properly immersed arc in f, with endpoints on the same edges as those of c_{k+1}. Thus there is a transverse homotopy E taking c_{k+1} to $H(c'_{k+1})$. Reparametrising E, we set $F_{k+1} = F_k \cup E$. See Figure 6, right.

This proves hypothesis (3). □

Let $B_\triangleright = B_{K-1}$. This is the *right-bigon* for R_0. We rerun the argument of Claim 5.7 to the left to obtain the *left-bigon* for R_0, denoted by $B_\triangleleft$.

The induction in the proof of Claim 5.7 extends to show that R_K contains a cusp pointing at γ'_{K-1}. The cusp lies between b_{K-1} and another branch on the boundary of R_K, which we call c'_K. See the far right of Figure 5. Let $Q_\triangleright$ be a small closed regular neighbourhood of b_{K-1} in R_K. The boundary of the subregion $Q_\triangleright$ has four sides; we call it a *right-quadrilateral*. The four sides are $d_K \cap N$, b_{K-1}, $c_K \cap N$ and a fourth side, $s_\triangleright$ say. Note that $s_\triangleright$ is properly embedded in R_K. The quadrilateral $Q_\triangleright$ therefore has one cusp and three corners, and so it has index $-\frac{1}{4}$. An identical argument builds the *left-quadrilateral* $Q_\triangleleft$.

Let $S(R_0) = Q_\triangleleft \cup B_\triangleleft \cup B_\triangleright \cup Q_\triangleright$.

5.8 Claim *For any max-bigons R and R',*

1. *$S(R)$ is embedded in D,*
2. *if $R \neq R'$ then $S(R)$ and $S(R')$ are disjoint, and*
3. *$S(R)$ is a rectangle.*

Proof Let $B_\triangleleft$ and $B_\triangleright$ be the right- and left-bigons for R; define $B'_\triangleleft$ and $B'_\triangleright$ similarly for R'. Note that boundary trigons in $B_\triangleleft$ have transverse orientations on their branches that disagree with the tangential orientation on ∂D. On the other hand, boundary trigons in $B'_\triangleright$ have transverse orientations that agree with the tangential orientation.

This proves that $B_\triangleleft$ and $B_\triangleright$ share only one region — the max-bigon itself — and so $B = B_\triangleleft \cup B_\triangleright$ is again a boundary bigon. The same argument shows that B and $B' = B'_\triangleleft \cup B'_\triangleright$ have no regions in common if $R \neq R'$.

We claim that ∂B and $\partial B'$ are disjoint. To see this, note that ∂B consists of an arc in ∂D and an arc in τ. The transverse orientation on the arc in τ points into B, and similarly for B'.

Let $Q_\triangleleft$, $Q_\triangleright$, $Q'_\triangleleft$ and $Q'_\triangleright$ be the quadrilaterals for R and R'. Since these are obtained by taking subsets of small regular neighbourhoods of branches in ∂B and $\partial B'$, these are all pairwise disjoint (if $R \neq R'$). This proves parts (1) and (2).

Adding the subregions $Q_\triangleleft$ and $Q_\triangleright$ replaces the two corners of B with four corners, and thus $S(R)$ is a rectangle, and we obtain (3). □

Let $D' = D - \bigcup S(R)$, where the union ranges over all max-bigons R.

5.9 Claim *The induced cellulation of D' has no regions of positive index.*

Proof Suppose that R' is a region of D' having positive index. If R' were a nullgon or monogon, then it would be a nullgon or monogon in D, contradicting Lemma 3.1. The region R' is not a boundary bigon since we removed all of them. Thus R' was created by cutting quadrilaterals out of some region R of $D-\tau$. Note that R' meets τ, meets ∂D and meets $\partial Q_{\triangleright}$ (say) along some side $s_{\triangleright}$. So R' has at least three corners. Since its index is positive, R' has exactly three corners. Thus $R = R' \cup Q_{\triangleright}$ is a boundary trigon, contradicting Claim 5.5. □

Note that D' has both outward and *inward* corners $\big($a combinatorial version of the exterior angle being $\frac{3\pi}{2}\big)$. Again following [8, page 57], we generalise our definition of index; each inward corner adds $+\frac{1}{4}$ to the overall index. Thus D' has nonpositive index. Since rectangles have index zero, from the additivity of index we deduce that D has nonpositive index, a contradiction. This concludes the proof of Theorem 5.1. □

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, J H Rubinstein, editors), Contemp. Math. 560, Amer. Math. Soc., Providence, RI (2011) 1–17 MR

[2] **D Calegari**, *Foliations transverse to triangulations of* 3*–manifolds*, Comm. Anal. Geom. 8 (2000) 133–158 MR

[3] **D Calegari**, *Foliations and the geometry of* 3*–manifolds*, Oxford Univ. Press (2007) MR

[4] **A Candel**, **L Conlon**, *Foliations, II*, Graduate Studies in Mathematics 60, Amer. Math. Soc., Providence, RI (2003) MR

[5] **D Gabai**, *Foliations and the topology of* 3*–manifolds*, J. Differential Geom. 18 (1983) 445–503 MR

[6] **D Gabai**, **U Oertel**, *Essential laminations in* 3*–manifolds*, Ann. of Math. 130 (1989) 41–73 MR

[7] **M Lackenby**, *Taut ideal triangulations of* 3*–manifolds*, Geom. Topol. 4 (2000) 369–395 MR

[8] **L Mosher**, *Train track expansions of measured foliations*, preprint (2003) `https://homepages.warwick.ac.uk/~masgar/Maths/arationality_03_12_28.pdf`

[9] **S P Novikov**, *The topology of foliations*, Trudy Moskov. Mat. Obšč. 14 (1965) 248–278 MR In Russian; translated in Trans. Mosc. Math. Soc. 14 (1965) 268–304

[10] **U Oertel**, *Homology branched surfaces: Thurston's norm on* $H_2(M^3)$, from "Low-dimensional topology and Kleinian groups" (D B A Epstein, editor), London Math. Soc. Lecture Note Ser. 112, Cambridge Univ. Press (1986) 253–272 MR

[11] **R C Penner**, **J L Harer**, *Combinatorics of train tracks*, Annals of Mathematics Studies 125, Princeton Univ. Press (1992) MR

[12] **I Rivin**, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. of Math. 139 (1994) 553–580 MR

[13] **W P Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1979) http://msri.org/publications/books/gt3m

Mathematics Institute, University of Warwick
Coventry, United Kingdom

Department of Mathematics, Oklahoma State University
Stillwater, OK, United States

s.schleimer@warwick.ac.uk, segerman@math.okstate.edu

http://homepages.warwick.ac.uk/~masgar/,
https://math.okstate.edu/people/segerman/

Received: 10 March 2019 Revised: 31 May 2019