

AG
T

*Algebraic & Geometric
Topology*

Volume 24 (2024)

**Dynamics of veering triangulations:
infinitesimal components of their flow graphs and applications**

IAN AGOL
CHI CHEUK TSANG



Dynamics of veering triangulations: infinitesimal components of their flow graphs and applications

IAN AGOL
CHI CHEUK TSANG

We study the strongly connected components of the flow graph associated to a veering triangulation, and show that the infinitesimal components must be of a certain form, which have to do with subsets of the triangulation which we call “walls”. We show two applications of this knowledge: first, we fix a proof in the original paper by the first author which introduced veering triangulations; and second, give an alternate proof that veering triangulations induce pseudo-Anosov flows without perfect fits, which was initially proved by Schleimer and Segerman.

57M50; 37D20, 37E30

1 Introduction

Veering triangulations were introduced in [Agol 2011] to study mapping tori of pseudo-Anosov homeomorphisms. In that setting, these are ideal triangulations of the punctured mapping tori that encode a folding sequence of train tracks associated to the pseudo-Anosov homeomorphism. In [Agol 2011], veering triangulations were used to prove a quantitative refinement of a result obtained initially by Farb, Leininger, and Margalit [Farb et al. 2011], which states that given a bound on the normalized dilatation of the pseudo-Anosov monodromy, there are only finitely many homeomorphism types of such punctured mapping tori. However, the proof presented in [Agol 2011] contained a gap: an estimate by Ham and Song [2007] was applied to a matrix which records how weights on the branches of the train tracks distribute under the folding moves, but the estimate only applies to irreducible matrices, and there exist examples for which this fails to be the case.

Meanwhile, since [Agol 2011] appeared, the study of veering triangulations has developed in many other directions; see for example [Guéritaud 2016; Hodgson et al. 2011; 2016; Landry 2018]. In particular, Guéritaud [2016] reconstructed the veering triangulations for punctured mapping tori in terms of the pseudo-Anosov suspension flow. Among other things, this shows that the veering triangulation is canonically associated to each fibered face of the Thurston unit ball in H_2 , instead of individual fiberings in the interior of the face. Generalizing this, the first author and Guéritaud showed that if a 3-manifold has a pseudo-Anosov flow without perfect fits, then the manifold obtained by drilling out the singular orbits of the flow admits a veering triangulation (see [Landry et al. 2023]).

Recently, Schleimer and Segerman proved the converse of this: if a 3–manifold admits a veering triangulation, then appropriate Dehn fillings of it carry pseudo-Anosov flows without perfect fits. Their construction involves isotoping the stable and unstable branched surfaces in order to form a dynamic pair in the sense of Mosher [1996]. The isotopy is in turn constructed by analyzing how the branched surfaces behave with respect to a canonical decomposition of the 3–manifold into veering solid tori. Details of this construction will appear in [Schleimer and Segerman \geq 2024]. We remark that Schleimer and Segerman’s construction is in fact part of a big program showing that veering triangulations and pseudo-Anosov flows without perfect fits are in correspondence with each other. In particular, their construction is inverse to that of Agol and Guéritaud in both directions, in a suitable sense. The first three parts of their program are [Schleimer and Segerman 2020; 2021; Frankel et al. 2019]. Also see the introduction of [Frankel et al. 2019] for an outline of the whole program.

This paper was born out of an attempt to reprove Schleimer and Segerman’s construction using more direct means. There is a standard way of constructing pseudo-Anosov flows on a 3–manifold starting from an embedded graph with desirable properties. See the concept of templates, introduced by Birman and Williams [1983a; 1983b] (under the name of “knot-holders”), and the concept of dynamic pairs, introduced by Mosher [1996]. The general strategy is to thicken up the graph using flow boxes, then collapse along the complement to obtain a flow on the whole 3–manifold. Given a veering triangulation, there is a natural candidate to apply this strategy to: the flow graph. (See Section 2 for definitions of objects associated to veering triangulations.) However, the fact that the flow graph might not be strongly connected presents difficulties in both constructing the flow and analyzing it.

This turns out to be exactly the same issue underlying the gap in [Agol 2011]. In this paper, we explain a way of addressing this. We show that the problematic infinitesimal components must arise in a certain form, which have to do with subsets of the veering triangulation which we call “walls”. By throwing out these components, we obtain the reduced flow graph, which shares many of the same features as the flow graph.

By analyzing this reduced flow graph, we are able to tackle the proof in [Agol 2011]. The explicit quantitative bound we get is the following:

Theorem 4.1 *If M is the punctured mapping torus of a pseudo-Anosov homeomorphism $\phi: S_{g,n} \rightarrow S_{g,n}$, where the normalized dilatation $\lambda(\phi)^{2g-2+\frac{2}{3}n} \leq P$, then M has a veering triangulation with at most*

$$\frac{P^9 - 1}{2} \left(\frac{2 \log P^9}{\log(2P^{-9} + 1)} - 1 \right)$$

tetrahedra.

Note that the bound we obtain is worse than that stated in [Agol 2011] by an exponent of $2 + \varepsilon$, but there might be room for improvement.

To reprove Schleimer and Segerman’s construction, we apply the general strategy outlined above to the reduced flow graph: we thicken up the graph by flow boxes, and collapse along its complement, with

the help of the unstable branched surface. When the filling slopes satisfy a necessary condition, it is not difficult to see that the resulting flow is pseudo-Anosov. Furthermore, we can leverage the structure of the unstable branched surface to show that the pseudo-Anosov flow obtained has no perfect fits relative to the “filled” orbits. The precise result is as below; see Sections 2 and 5 for relevant definitions.

Theorem 5.1 *Suppose M admits a veering triangulation. Let $l = (l_i)$ denote the collection of ladderpole curves on each boundary component. Then $M(s)$ admits a transitive pseudo-Anosov flow ϕ if $|\langle s, l \rangle| \geq 2$.*

Furthermore, there are closed orbits c_i isotopic to the cores of the filling solid tori, such that each c_i is $|\langle s_i, l_i \rangle|$ -pronged, and ϕ is without perfect fits relative to $\{c_i\}$.

We remark that there are 2 notions of pseudo-Anosov flows common in the 3-manifold topology literature: topological pseudo-Anosov flows and smooth pseudo-Anosov flows. Theorem 5.1 holds for both notions.

We also remark that, in contrast to the work by Schleimer and Segerman, it is not clear whether our construction provides a correspondence between veering triangulations and pseudo-Anosov flows. We pose this as a question more carefully in Section 6.

The results we discuss in this paper run parallel to those in [Landry et al. 2023]. The walls in this paper are closely related to what are called AB regions in [Landry et al. 2023]. In the setting of [Landry et al. 2023], one starts with a pseudo-Anosov flow without perfect fits and builds a veering triangulation transverse to the flow. Then it turns out that the flow graph encodes the orbits of the flow, with AB cycles governing the extent of over- and under-counting. In this paper we go in the opposite direction, starting with a veering triangulation and constructing a pseudo-Anosov flow without perfect fits. We end up with a similar conclusion: the flow graph encodes the orbits of the flow, with infinitesimal cycles of walls governing the extent of over- and under-counting. We will explain these connections between this paper and [Landry et al. 2023] where relevant.

Here is an outline of this paper. In Section 2, we review the notion of veering triangulations and related constructions. In Section 3, we study the infinitesimal components of the flow graph and show that they must arise from walls. We also show how to define the reduced flow graph by throwing away these infinitesimal components. In Section 4, we analyze the reduced flow graph of layered veering triangulations to fix the gap in [Agol 2011] and prove Theorem 4.1. In Section 5, we reprove Schleimer and Segerman’s result (Theorem 5.1) using the reduced flow graph. Finally in Section 6, we discuss some questions coming out this paper. We remark that Sections 4 and 5 are independent of each other, and the reader can jump to either of them after reading Section 3.

Acknowledgments The first author discovered the gap in [Agol 2011] with the help of the veering triangulation census. We thank Andreas Giannopolous, Saul Schleimer, and Henry Segerman for making the data in the census available. We also thank Saul Schleimer and Henry Segerman for sharing with us an early version of [Schleimer and Segerman \geq 2024]. We thank Michael Landry for the support

and encouragement throughout this project. We thank Rafael Potrie on MathOverflow for pointing us in the direction of Mario Shannon's thesis, and Mario Shannon for guiding us through the material in Section 5.5. We thank Anna Parlak and Samuel Taylor for comments on an early version of this paper. We thank the referee for suggestions in improving the exposition and clarity of the paper.

Agol was supported by MSRI and the Simonyi Professorship at IAS during part of this project. Both authors are partially supported by the Simons Foundation, grant 376200.

Notational conventions Throughout this paper:

- \bar{M} will be an oriented compact 3-manifold with torus boundary components, while M will be the interior of such a manifold. We will sometimes conflate a torus end of M with the corresponding boundary component of \bar{M} .
- $X \setminus\!\!\setminus Y$ will denote the metric completion of $X \setminus Y$ with respect to the induced path metric from X . In addition, we will call the components of $X \setminus\!\!\setminus Y$ the complementary regions of Y in X .
- \tilde{X} will mean a universal cover of X , unless otherwise stated.

2 Background: veering triangulations

We recall the definition of a veering triangulation.

Definition 2.1 An *ideal tetrahedron* is a tetrahedron with its 4 vertices removed. The removed vertices are called the *ideal vertices*. An *ideal triangulation* of M is a decomposition of M into ideal tetrahedra glued along pairs of faces.

A *taut structure* on an ideal triangulation is a labeling of the dihedral angles by 0 or π , such that

- (1) each tetrahedron has exactly two dihedral angles labeled π , and they are opposite to each other;
- (2) the angle sum around each edge in the triangulation is 2π .

Intuitively this means that there is a degenerate geometric structure on the triangulation where every tetrahedron is flat.

A *transverse taut structure* is a taut structure along with a coorientation on each face, such that for any edge labeled 0 in a tetrahedron, exactly one of the faces adjacent to it is cooriented out of the tetrahedron.

A *transverse taut ideal triangulation* is an ideal triangulation with a transverse taut structure.

We will always take the convention that the face coorientations are pointing upwards in our figures and descriptions.

Definition 2.2 A *veering structure* on a transverse taut ideal triangulation is a coloring of the edges by red or blue, so that looking at each flat tetrahedron from above, when we go through the 4 outer edges




Figure 1: A tetrahedron in a transverse veering triangulation. There are no restrictions on the colors of the top and bottom edges.

counterclockwise, starting from an endpoint of the inner edge in front, the edges are colored red, blue, red, blue in that order. We call such a tetrahedron a *veering tetrahedron*.

A *veering triangulation* is a transverse taut ideal triangulation with a veering structure.

Figure 1 shows a veering tetrahedron in a veering triangulation.

From now on Δ will denote a veering triangulation on M .

Remark 2.3 Some authors call the notion we have just defined a transverse veering triangulation, and call the version of this notion without face coorientations veering triangulations instead. These two versions do not differ by much, since one can always take a double cover to make the faces coorientable. However, we will just be studying the version with face coorientations in this paper. Without the face coorientations, our main object of study, the flow graph (Definition 2.7), cannot be defined.

We recall some combinatorial facts and constructions for veering triangulations.

Definition 2.4 A veering tetrahedron in Δ is called a *toggle tetrahedron* if the colors on its top and bottom edges differ. It is called a *red/blue fan tetrahedron* if both its top and bottom edges are red/blue respectively.

Note that some authors call toggle and fan tetrahedra hinge and nonhinge respectively.

Proposition 2.5 [Futer and Guéritaud 2013, Observation 2.6] *Every edge e in Δ has one tetrahedron above it, one tetrahedron below it, and two stacks of tetrahedra, in between the tetrahedra above and below, on either of its sides.*

Each stack must be nonempty. Suppose e is red (resp. blue). If there is exactly one tetrahedron in one stack, then that tetrahedron is a red (resp. blue) fan tetrahedron. If there are $n > 1$ tetrahedron in one stack, then going from top to bottom in that stack, the tetrahedra are: one toggle tetrahedron, $n - 2$ blue (resp. red) fan tetrahedra, and one toggle tetrahedron.




Figure 2: The unstable branched surface meets the faces of a veering tetrahedron in the unstable train track, which undergoes a folding move as one goes from the top faces to the bottom faces. The flow graph can be embedded in the unstable branched surface.

Definition 2.6 For each tetrahedron of Δ , define a branched surface inside it by placing a quadrilateral with vertices on the top and bottom edges and the two side edges of the same color as the top edge, then adding a triangular sector for each side edge of the opposite color to the top edge, with a vertex on that side edge and attached to the quadrilateral along an arc going between the two faces adjacent to that side edge. We also require that the arcs of attachment for the two triangular sectors intersect only once on the quadrilateral. See Figure 2, top left. These branched surfaces in each tetrahedron can be arranged to match up across faces, thus glue up to a branched surface in M , which we call the *unstable branched surface* B . The intersection of the unstable branched surface with the faces of Δ is called the *unstable train track*. Notice that as one goes from the top faces to the bottom faces of each veering tetrahedron, the unstable train track undergoes a *folding move*. See Figure 2, bottom.

The branch locus of B is a union of circles, smoothly carried by the branch locus and intersecting transversely at *double points* of the branch locus. We call these circles the *components* of the branch locus. Orient the circles so that they intersect the 2-skeleton of the veering triangulation negatively, i.e. they disagree with the coorientation of the faces whenever they meet. This orientation has the special property that at double points of the branch locus, it always points from the side with more sectors to the side with less sectors.

Note that the unstable branched surface, when considered as a cell complex by taking the 0-cells to be the double points of its branch locus, the 1-skeleton to be the branch locus, and the 2-cells to be the sectors, is dual to the veering triangulation. As such, it makes sense for us to talk about, for example, the sector of B dual to a given edge of Δ .

Also, notice that each sector of B is of the form of a diamond. Each sector has two upper edges and at least two lower edges, at least one on either side; one top vertex, two side vertices, some vertices between the lower edges on either side, and one bottom vertex. We refer to [Frankel et al. 2019, Section 6.13] for details.

Definition 2.7 Let Δ be a veering triangulation. The *flow graph* Φ is a directed graph with vertex set $V(\Phi)$ equals to the set of edges of Δ , and edge set $E(\Phi)$ defined by adding 3 edges per tetrahedron, going from the top edge and the two side edges of opposite color to the top edge, into the bottom edge.

We note that the flow graph was first defined in [Landry et al. 2024]. However the reader is cautioned that there the flow graph is oriented in the opposite direction compared to here.

The flow graph Φ can be embedded in the unstable branched surface B in the following way. Place each vertex of Φ at the top vertex of the sector dual to the corresponding edge of Δ , and place the edges of Φ exiting that vertex in the interior of the sector, as shown in Figure 2, top right. See also [Landry et al. 2024, Section 4.3]. The flow graph inherits the structure of a (nongeneric) oriented train track from this embedding, by making the edges meeting a vertex tangent to a vertical line. We will consider the flow graph to be a subset of B in this way from now on.

The ideal triangulation Δ on M induces a triangulation $\partial\Delta$ on the torus boundary components of \overline{M} , by considering them as links of the ideal vertices. The branch locus of the unstable branched surface inside each tetrahedron opens up towards two opposite ideal vertices; hence each side determines an oriented interval between two edges of a face in $\partial\Delta$. Join these intervals end-to-end. Since each face in $\partial\Delta$ contains at most one such interval, the paths must close up to form oriented parallel loops. Also note that since every edge in Δ is the top edge of some tetrahedron, each boundary component of \overline{M} will receive at least one loop.

Definition 2.8 [Futer and Guéritaud 2013] The homotopy class of the union of these oriented parallel loops on each boundary component of \overline{M} is called the *ladderpole curve* on the boundary component. The slopes they determine is called the *ladderpole slope* on the boundary component.

We end this section by analyzing the complementary regions of B in M and the complementary regions of Φ in B .

Proposition 2.9 *The components of $M \setminus B$ are (once-punctured cusped polygons) $\times S^1$, where each cusp $\times S^1$ represents the ladderpole slope on the corresponding boundary component of \overline{M} , and the sum over all cusp circles in each component represents the ladderpole curve on the corresponding boundary component of \overline{M} .*

Proof Let T be such a component. T is a neighborhood of a torus end of M . Inside each tetrahedron t , $T \cap t$ is homeomorphic to the product of $T \cap t \cap B$ with $[0, \infty)$. Hence the product structures glue up to give a homeomorphism $T \cong T^2 \times [0, \infty)$. ∂T inherits the branch locus of B as cusp circles, with these representing the ladderpole classes and slopes as described in the statement by definition. An

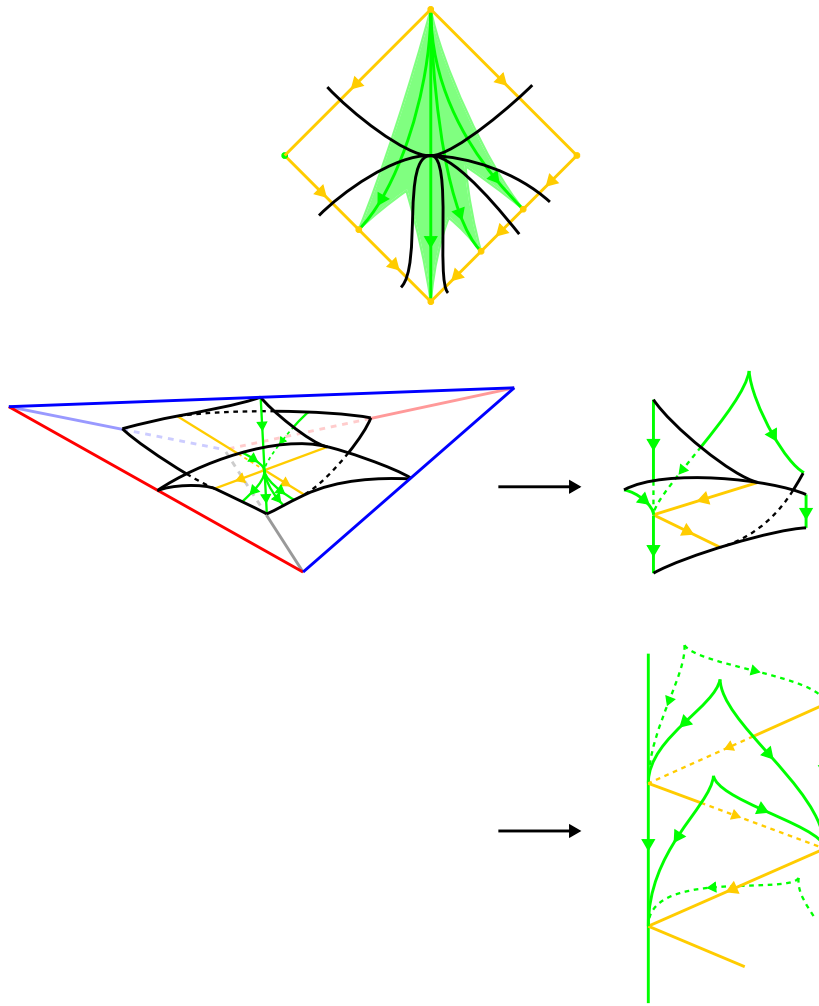


Figure 3: The complementary regions of Φ in B are annuli/Möbius bands with tongues.

identification of ∂T with T^2 sending these cusp circles to longitudes extend to an identification of T with $(\text{once-punctured cusped } n\text{-gon}) \times S^1$, where n is the number of cusp circles. \square

Lemma 2.10 *The components of $B \setminus \Phi$ are annuli or Möbius bands with tongues, ie they can be obtained by attaching triangular sectors (“tongues”) along arcs to a smooth annulus or Möbius band and smoothing all the arcs of attachment in the same direction with respect to the core of the annulus/Möbius band. (This terminology is borrowed from [Mosher 1996].)*

Moreover, the arcs of attachment zigzag along the annulus/Möbius band. More precisely, the arcs on the annulus/Möbius band lift to $y = \pm x + 2i$, for $i \in \mathbb{Z}$, in the universal cover $[0, 1] \times \mathbb{R}$. These arcs are subintervals of the branch locus of B and are oriented downwards (ie decreasing y in the above model). See Figure 3, bottom.

Proof The topology of the complementary regions is unchanged if we thicken up Φ to be a regular neighborhood of itself in B within each sector. Hence we can take a neighborhood of Φ in each sector of B as in Figure 3, top, where the black lines in the figure are where the faces of Δ meet the sector.

With this adjustment in place, it is straightforward to analyze the portion of each complementary region within a veering tetrahedron. There are two of these for each tetrahedron, corresponding to the two side edges of the same color as the top edge. We show the form of these components in Figure 3 middle. In particular they intersect the faces of the tetrahedron in train track switches.

Components of $B \setminus \Phi$ can be obtained by gluing these pieces along the train track switches on the top and bottom, and the result of the gluing must be an annulus or Möbius band with tongues with the arcs of attachments of the tongues as described. \square

3 Infinitesimal components of the flow graph

Definition 3.1 A directed graph G is said to be *strongly connected* if for every ordered pair of vertices (v, w) there a directed edge path going from v to w .

The *adjacency matrix* of G is defined to be the matrix $A \in \text{Hom}(\mathbb{R}^{V(G)}, \mathbb{R}^{V(G)})$ where each entry A_{wv} is the number of edges going from v to w . It is easy to see that G being strongly connected is equivalent to its adjacency matrix A being *irreducible*, ie for every v and w , the entry $(A^n)_{wv}$ is positive for some $n > 0$.

For reasons explained in the introduction, it would be convenient to have a strongly connected flow graph associated to a given veering triangulation. However, a brief search through the veering triangulation census [Giannopolous et al. 2019] gives examples which this property fails, for instance: eLAKbbcdhqwqj_2102 (which is layered) and fLAMcaccdeejsnaxk_20010 (which is nonlayered). See [Giannopolous et al. 2019] for what these codes mean. For completeness, we also note that there are examples which this property holds: cPcbbbdxm_10 (which is layered), and gLLAQbddeeffennmann_011200 (which is nonlayered). As we shall see later, further examples can be constructed by taking covers.

Our goal in this section is to analyze how this property can fail and explain how to prune the flow graph to make a version of this property hold.

3.1 Strongly connected components

We first set up some notation for discussing strong connectivity.

Definition 3.2 Let G be a directed graph, let v and w be vertices of G . Write $v \succsim w$ if there is a directed edge path going from v to w . Write $v \sim w$ if $v \succsim w$ and $v \lesssim w$. Note that \sim is an equivalence relation; call equivalence classes of \sim *strongly connected components*.

Construct a directed graph G/\sim with vertex set equal to the equivalence classes $[v]$ of \sim , and an edge from $[v]$ to $[w]$ if $v \succ w$ and $[v] \neq [w]$. Call a strongly connected component $[v]$ of G an *infinitesimal component* if $[v]$ has an incoming edge in G/\sim .

A subset V of $V(G)$ is called a *minimal set* if it has the property that if $v \in V$ and $v \succ w$ then $w \in V$.

It is easy to see that a finite directed graph is a disjoint union of strongly connected graphs if and only if it has no infinitesimal components. Similarly, a finite directed graph is strongly connected if and only if it has no proper minimal sets. In this sense, infinitesimal components, or minimal sets in general, are the obstruction to strong connectivity. Hence our approach to understanding the failure of strong connectivity of a given flow graph is to analyze its infinitesimal components and minimal sets.

3.2 Walls

In this section, we introduce the concept of walls of a veering triangulation. These present one way in which infinitesimal components of flow graphs can arise. What we will show in Section 3.3 is that these account for all infinitesimal components of flow graphs.

Definition 3.3 Let Δ be a veering triangulation on a 3-manifold M . Suppose there are tetrahedra $t_{i,j}$, for $1 \leq i \leq w+1$ and $j \in \mathbb{Z}/h$, where $w \geq 2$, such that for $2 \leq i \leq w$:

- (1) The bottom edge of $t_{i,j}$ is the top edge of $t_{i,j+1}$
- (2) If i is odd, the bottom edge of $t_{i,j}$ are side edges of $t_{i-1,j}$ and $t_{i+1,j}$ and no other tetrahedra in Δ .
If i is even, the bottom edge of $t_{i,j}$ are side edges of $t_{i-1,j+1}$ and $t_{i+1,j+1}$ and no other tetrahedra in Δ .

Then we call the indexed multiset $\{t_{i,j}\}$ a wall, and call w the *width* of the wall. We emphasize that in this definition, $t_{i,j}$ for different i and j may not be different tetrahedra in Δ .

Notice that for $w \geq 3$, it is possible to extract a proper subcollection of $\{t_{i,j}\}$ which will form a wall of smaller width. A wall will be called *maximal* if the collection of tetrahedra $\{t_{i,j}\}$ cannot be enlarged into a wall of larger width. Similarly, one can always replace h with a multiple of h by renaming the tetrahedra appropriately. In the sequel we will implicitly assume that h is the minimum possible value that satisfies the definition.

Suppose we have a wall $\{t_{i,j}\}$ of a veering triangulation Δ . Notice that for $2 \leq i \leq w$, by Proposition 2.5 and Definition 2.7, the bottom edge of $t_{i,j}$ has only one outgoing edge in the flow graph Φ , namely the edge going to the bottom edge of $t_{i,j+1}$. Hence there is a cycle c_i in the flow graph passing through the top/bottom edges of $t_{i,j}$, for which there are edges entering c_i but no edges exiting c_i . In other words, c_i is an infinitesimal component of the flow graph for each $2 \leq i \leq w$. We call these c_i the *infinitesimal cycles* of the wall.

Meanwhile for each j there is an edge of Φ going from the bottom edge of $t_{1,j}$ to the bottom edge of $t_{1,j+1}$, since the former is either the top edge of $t_{1,j+1}$ or a side edge of opposite color as the top edge

of $t_{1,j+1}$. Let c_1 be the cycle formed by these edges. Similarly, there is a cycle c_{w+1} passing through the bottom edges of $t_{w+1,j}$. We call c_1 and c_{w+1} the *boundary cycles* of the wall.

Here is a convenient way to visualize a wall, at least for those with width $w \geq 3$. Inside each tetrahedron, there is a unique quadrilateral carried by the unstable branched surface with boundary lying on the unstable train track. If $w \geq 3$, by applying Proposition 2.5 to the bottom edges of $t_{i,j}$ for $2 \leq i \leq w$, we can see that the tetrahedra $t_{i,j}$ are all fan tetrahedra, and they are either all blue fan tetrahedra or all red fan tetrahedra. Hence the quadrilaterals that lie inside them are adjacent to each other in B , and their union forms a surface carried by B that is the image of an annulus or Möbius band under an immersion. The quadrilaterals form a tiling of the annulus or Möbius band, with 4 quadrilaterals meeting at each vertex, resembling a tiled wall. In fact, this is the reason why we have chosen to call the collection of tetrahedra a wall. In particular, we note that in this case Definition 3.3(1) holds for $i = 1, w + 1$ as well, and so the infinitesimal and boundary cycles c_i lie within the quadrilateral tiling as vertical loops.

If $w = 2$, we can try to repeat the above argument, but the picture is not as clean. By the same argument as above, $t_{1,j}$ and $t_{3,j}$ are fan tetrahedra, but the bottom edge of $t_{1,j}$ may be different from the top edge of $t_{1,j+1}$, and similarly for $t_{3,j}$. As a result, the union of quadrilaterals inside the tetrahedra do not nicely tile up an annulus or Möbius band, but instead at some vertices some quadrilaterals from the bottom might “peel away”.

We remark that the quadrilaterals we considered above are among the ones considered in [Hodgson et al. 2011], and this idea of looking at how quadrilaterals tile up also appeared there.

In Figure 4, we present 3 ways of illustrating a wall of width 4 in order to aid the reader’s intuition. The first picture shows the quadrilateral tiling as mentioned above (Figure 4, top left). The second picture shows a layered view in terms of a folding sequence of the unstable train track on the faces of tetrahedra in the wall (Figure 4, right). The third picture shows the portion of the unstable branched surface in a small neighborhood of the wall (Figure 4, bottom left). The flow graph contains vertical lines as subgraphs in the picture. These are the infinitesimal and boundary cycles of the wall.

In Figure 5, we also demonstrate the same 3 viewpoints for a wall of width 2, since as pointed out before, the combinatorics for width 2 walls are slightly more general than that for higher width walls.

We caution the reader that c_i for different i may be the same cycle in Φ . In fact, call a wall *twisted* if $t_{i,j} = t_{w+2-i,j+h'}$ for some h' , and call a wall *untwisted* otherwise. For a twisted wall, c_i and c_{w+2-i} are the same cycle for each i ; for an untwisted wall, c_i are distinct cycles. Equivalently, a wall is twisted if and only if the surface tiled by the quadrilaterals is homeomorphic to a Möbius band.

Another point of caution is that two distinct maximal walls $\{t_{i,j}\}$ and $\{t'_{i',j'}\}$ can share a tetrahedron, say $t_{i,j} = t'_{i',j'}$. By maximality, this is only possible if $i = 1$ or $w + 1$ and $i' = 1$ or $w' + 1$, where w and w' are the widths of the walls respectively. In this case, the appropriate boundary cycles of the two walls touch within the shared tetrahedron. In fact, this behavior can happen within a single wall as well, that is,

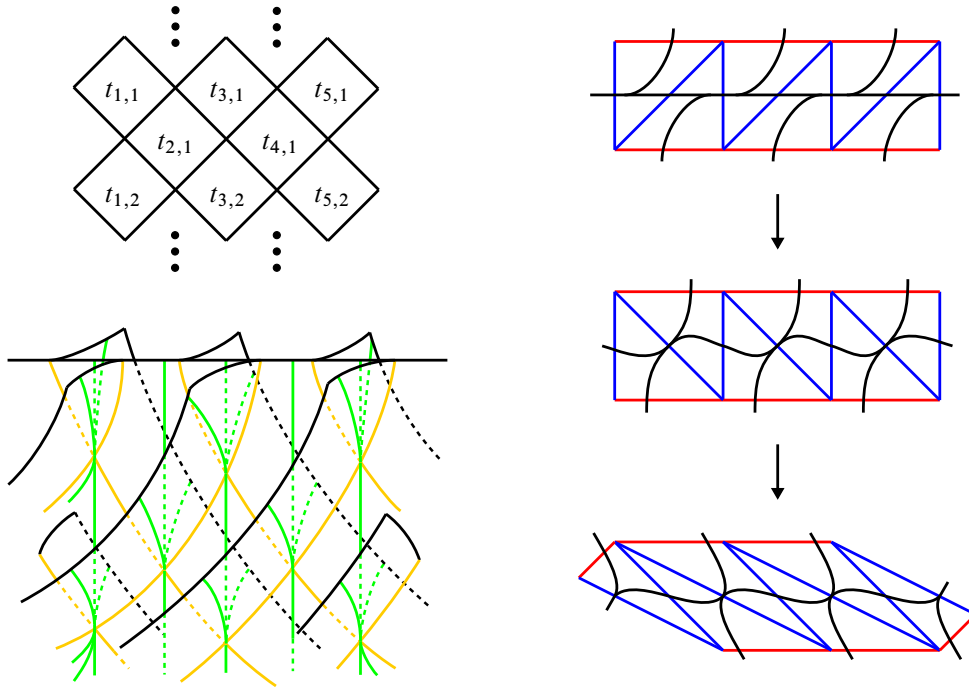


Figure 4: Illustrating a wall of width 4 from three viewpoints. Top left: tiling with quadrilaterals. Right: folding sequence of the unstable train track. Bottom left: unstable branched surface in a small neighborhood.

a tetrahedron can appear as $t_{1,j}$ or as $t_{w+1,j}$ for more than one value of j . However, a tetrahedron can only appear at most twice in walls, corresponding to the 2 side edges of the same color as the top edge. In particular a vertex of Φ can lie in at most 2 boundary cycles.

3.3 Classification of infinitesimal components

We begin our analysis of infinitesimal components of flow graphs. Throughout this subsection, fix a minimal set V of the flow graph Φ associated to a veering triangulation Δ . Recall that vertices of Φ are edges of Δ , so it makes sense to say whether e is in V for an edge e of Δ .

Each branch of the unstable train track τ in a face of Δ is dual to some edge of that face. Let τ' be the union of those branches dual to edges of Δ that lie in V . Up to rotation, there are 5 configurations for the portion of τ' lying on a face of Δ . We show and name these 5 types of faces in Figure 6.

We also consider the configurations for the portion of τ' lying on the boundary of a tetrahedron. Since V is minimal, and since the flow graph Φ contains the cycles formed by edges going from the top edge to the bottom edge of each tetrahedron, the top edge of a tetrahedron lies in V if and only if the bottom edge lies in V . Again by minimality, if a side edge of opposite color to the top edge is in V , then the bottom edge is in V , hence the top edge will be in V as reasoned above. From these restrictions, we can




Figure 5: Illustrating a wall of width 2 from the same three viewpoints as Figure 4.

enumerate 13 configurations for the portion of τ' on the boundary of a tetrahedron, up to rotation and reflection. We show and name these 13 types of tetrahedra in Figure 7.

Lemma 3.4 *Tetrahedra of type (VI)–(X) will not appear.*

Proof We use a double-counting argument,

$$\sum \# \text{ type (c) upper faces} - \# \text{ type (c) lower faces} = 0,$$

where the sum is over all tetrahedra of Δ , since each face belongs to the upper face of exactly one tetrahedron and the lower face of exactly one tetrahedron. By inspection of Figure 7, the number of type (c) upper faces is always greater or equal to the number of type (c) lower faces for each tetrahedron; for type (VI), (VIII), (IX) and (X) tetrahedra, strict inequality holds, so these in fact cannot occur.

A similar argument using type (i) faces eliminates the possibility of type (VII) tetrahedra (now that type (VIII) tetrahedra are eliminated). □

We now analyze what happens if we have a type (i) face. Let f be such a face, and e be the unique edge of f that lies in V .

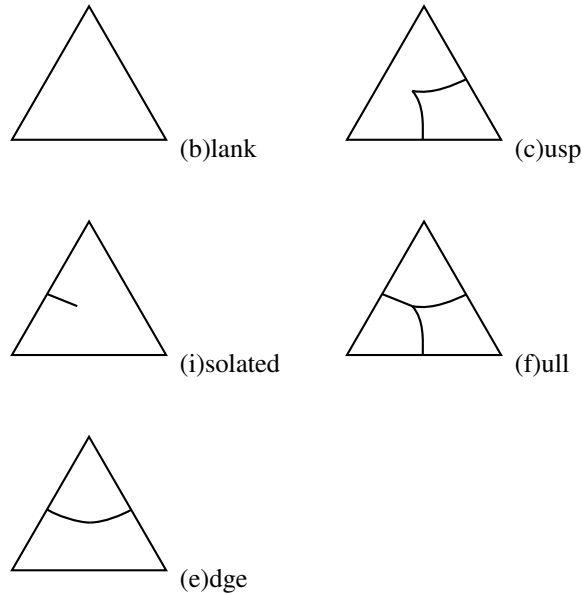


Figure 6: The 5 possible types of faces, labeled by letters (b), (i), (e), (c) and (f).

The edge e is the top edge of one tetrahedron, the bottom edge of one tetrahedron, and the side edges of two stacks of tetrahedra, one stack on each side of e . The face f determines a side of e , and we claim that the stack of tetrahedra on that side consists of one tetrahedron only.

Suppose otherwise. Then the top edge of the tetrahedron on the bottom of the stack is of opposite color to e (by Proposition 2.5). Looking at Figure 7, we claim that this tetrahedron must be of type (XI) or (XII). This is because these are the only types of tetrahedron with a side edge in V that is of opposite color to the top edge. As a consequence, the face of this tetrahedron which meet e must be of type (f). But again by looking at Figure 7, the face f being of type (i) forces the face below it and sharing the edge e to be of type (i) too (and the tetrahedron between them being of type (I) or (II)). In fact, the same is true for the face above f and sharing the edge e . Hence inducting upwards and downwards, the faces on the side of f that are adjacent to e must all be of type (i). This gives us a contradiction at the bottom of the stack of tetrahedra.

The tetrahedra on the top and bottom of e have to be of type (III) or (IV). In fact, they are either both of type (III) or both of type (IV), since by the reasoning above, a type (i) face on a side of e forces all of the faces in that stack with edge e to be of type (i).

If the tetrahedra on the top and bottom of e are of type (III), the stack on the side of e opposite to f also consists of one tetrahedron only by the reasoning two paragraphs above. In particular, e has degree 4 in this case. Note that there is still the freedom of the side tetrahedra being of type (I) or (II).

At this point, we recall the quadrilaterals that we considered in Section 3.2. By looking at the surface formed by the union of the quadrilaterals in the tetrahedra we have considered so far, this gives us a good

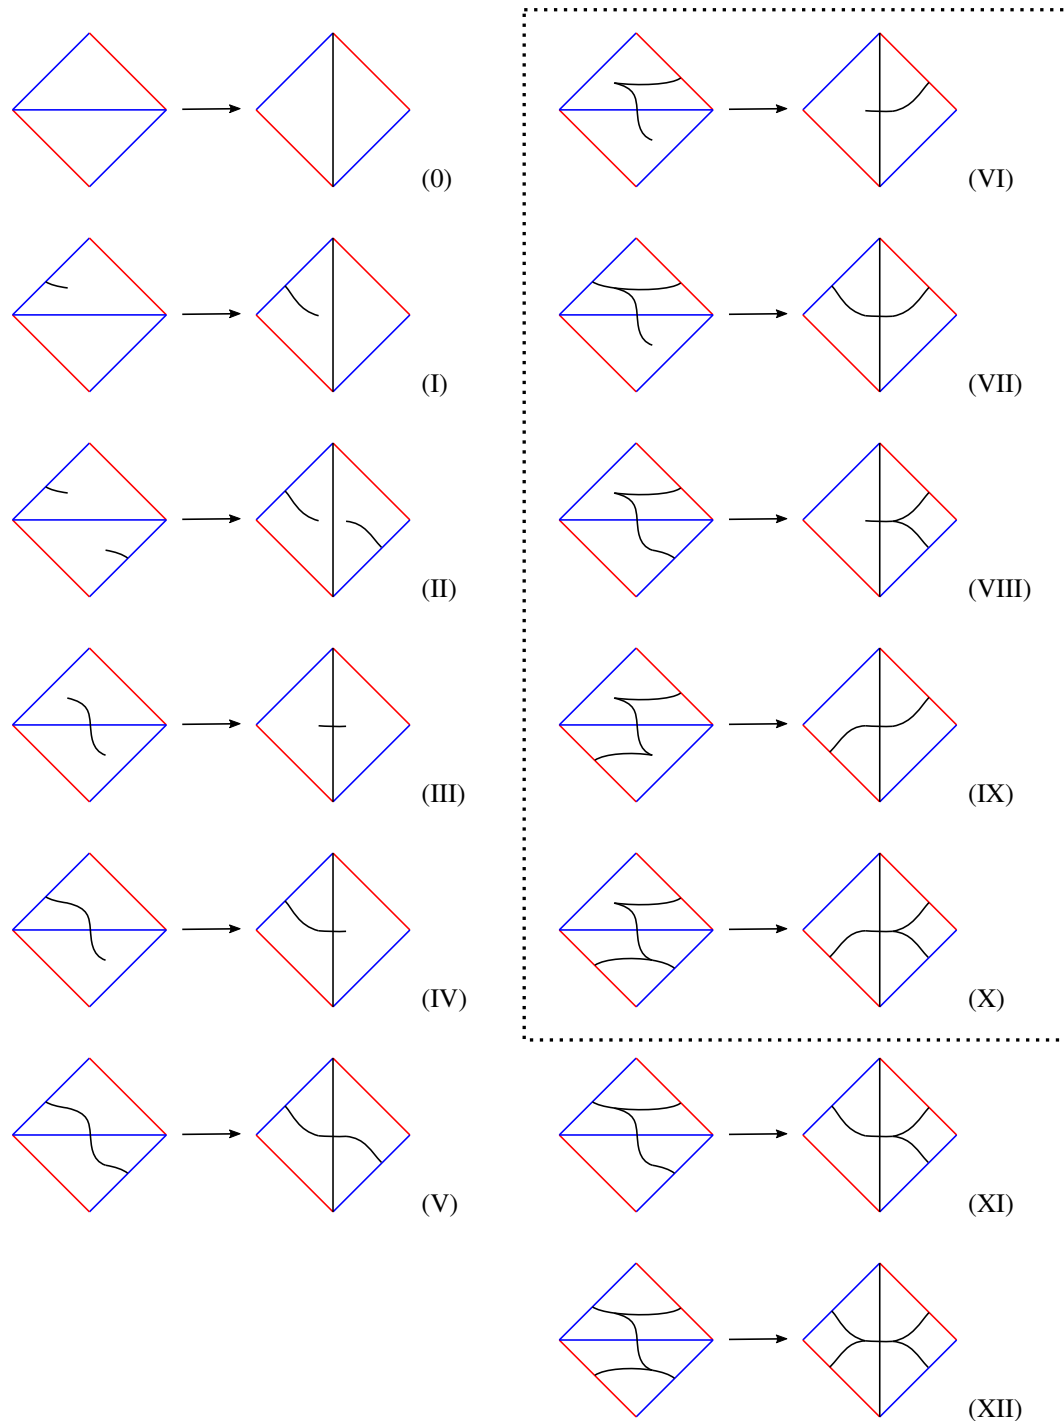


Figure 7: The 13 possible types of tetrahedra, labeled by Roman numerals (0)–(XII). We drew the tetrahedra by their top faces followed by their bottom faces. The tetrahedron of type (VI)–(X) inside the box are eliminated by Lemma 3.4.




Figure 8: The types of the tetrahedra around e split into three cases. The Roman numeral inside a quadrilateral is the type of the tetrahedron the quadrilateral belongs to.

way of keeping track of our argument. For example, in the situation of the paragraph above, we have 4 tetrahedra that share an edge e ; hence we have 4 quadrilaterals sharing a vertex at e . See Figure 8, left.

Now, our analysis can be extended in the vertical direction. The top edge of the tetrahedron above e is adjacent to a type (i) face. So we can repeat our arguments on there. This gives a group of 4 tetrahedra centered around that edge. In terms of the quadrilaterals, we are extending the tiling vertically. See Figure 9, left. A caveat of this picture, however, is that the top and bottom vertices of the quadrilaterals representing type (I) and (II) tetrahedra may not match up, since they might be representing different edges of the adjacent type (III) tetrahedron. See, for example, the situation of $t_{1,1}$ and $t_{1,2}$ in Figure 5. We denote this by making the quadrilaterals peel away slightly when the vertices do not match up. Since there are only finitely many tetrahedra, the vertical extension must loop back on itself eventually, giving us a width 2 wall whose infinitesimal cycle lies in V .

We still need to tackle the case when the tetrahedra on the top and bottom of e are of type (IV). We claim that the stack on the side of e opposite to f also consists of one tetrahedron only. Suppose otherwise; then arguing as before, the tetrahedron on the bottom of the stack is of type (XI) or (XII) and its faces that meet e are of type (f). But we know that the face at the bottom of the stack that meets e is of type (e), so this is not the case.

Looking at Figure 7, this single tetrahedron must be of type (IV) or (V), since these are the only types of tetrahedra that have a side edge in V meeting two type (e) faces. If the side tetrahedron is of type (IV), it has two type (i) faces, and the tetrahedra sharing those type (i) faces must be of type (I) or (II). In particular, we get a cluster of 6 tetrahedra. The corresponding 6 quadrilateral tile up a region as shown in Figure 8, center.

Now repeat the argument vertically as in the last case. We obtain a width 3 wall eventually, for which the infinitesimal cycles lie in V .

Finally, we tackle the case where e has type (IV) tetrahedra above and below and a type (V) tetrahedron on the side opposite to f . See Figure 8, right. Let e' be the top edge of this type (V) tetrahedron. By repeating our arguments up to this point vertically, we know that the tetrahedron on top of e' is of type (IV)

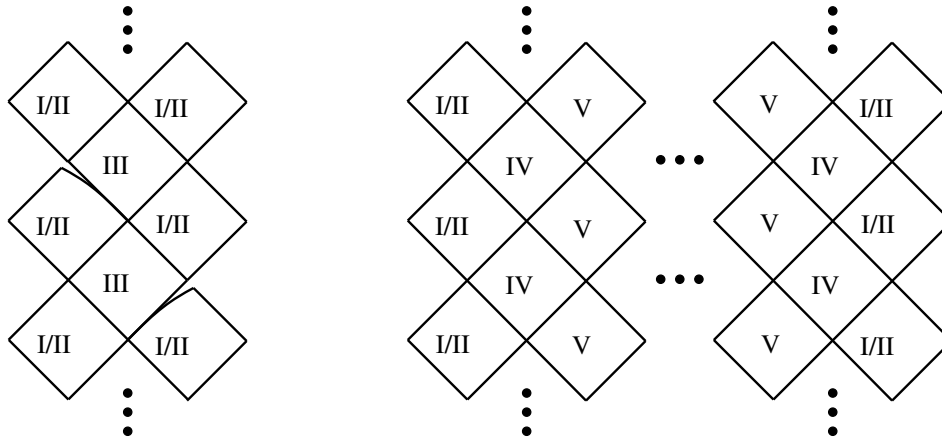


Figure 9: The patterns of quadrilaterals that can occur in our proof. Left: a width 2 wall. Right: a width ≥ 3 wall.

or (V). It is not of type (IV). Otherwise e' has a side where all the faces are of type (i), contradicting the fact that the type (V) tetrahedron below it has only type (e) faces. So the sides of e' have type (e) faces on the top and bottom of their stacks of tetrahedra. We have argued that this implies the stacks of tetrahedra on the sides of e' have one tetrahedron each, and that these side tetrahedra are of type (IV) or (V). The same argument also applies to the bottom edge of the original type (V) tetrahedron. In terms of our picture with quadrilaterals, this extends the tiling horizontally.

Now continue the analysis inductively, both horizontally and vertically. Again, the vertical extension must loop back on itself at some point. This implies that the horizontal extension has to stop at some point, otherwise it will loop back on itself and we will get a Klein bottle or a torus tiled by quadrilaterals from type (V) tetrahedra, which contradicts Δ having a strict angle structure (see [Hodgson et al. 2011, Theorem 1.5 and Corollary 3.11]). The horizontal extension stops by hitting a column of type (IV) tetrahedra, after which the tiling is “capped off” by type (I) or (II) tetrahedra. See Figure 9, right. Hence in every case, we get a wall with infinitesimal cycles contained in V .

We claim that if V is a proper subset of $V(\Phi)$, it consists solely of these infinitesimal cycles of walls. Suppose otherwise. Since these infinitesimal cycles have no incoming edges from vertices in V , we can remove all of them from V and we would still have a nonempty minimal set. Rename V as this new minimal set, and reclassify the faces and tetrahedra of Δ according to Figures 6 and 7 with respect to this new minimal set.

After this reclassification, there are no type (i) faces anymore, since those must be contained in walls by our analysis above, and we have eliminated all strongly connected components of V within walls. In particular, we can only have tetrahedra of type (0), (V), (XI) and (XII).

Suppose we have a type (e) face f . Let e be the unique edge of f not lying in V . Consider the stack of tetrahedra on the side of e determined by f . Looking at Figure 7, with the limited types of tetrahedra

we can have, as we move downwards the stack starting from f , we can only get to faces of type (e). So the stack must end on a type (e) face, but then the tetrahedron directly below this face, which is also the tetrahedron below e , cannot be any of the types in Figure 7, since none of those have a type (e) top face and have a top edge that does not lie in V . This contradiction shows that there are no type (e) faces anymore.

And so the only types of tetrahedra we can have are type (0) and (XII). In this case, it is clear that a tetrahedron sharing a face with a type (0) or a type (XII) tetrahedron must be of type (0) or type (XII) itself, respectively. Because M is connected, V is either empty or the entirety of $V(\Phi)$. This proves the claim and concludes our analysis of minimal sets of Φ , which we summarize as:

Theorem 3.5 *For a given veering triangulation Δ , the infinitesimal components of its flow graph Φ are exactly the infinitesimal cycles of walls. Furthermore, Φ/\sim is a rooted height 1 tree, ie it has a unique vertex v_0 such that there is a directed edge from v_0 to every other vertex.*

Proof We showed that any proper minimal set of Φ is a disjoint union of infinitesimal cycles of walls. This implies that every proper minimal set of Φ/\sim is a disjoint union of vertices. Furthermore, each infinitesimal cycle of a wall has incoming edges; hence each vertex of Φ/\sim which is an infinitesimal component has an incoming edge. This is enough to imply the second statement of the theorem. \square

We use this theorem to justify our statement earlier that for a finite cover $\tilde{\Delta}$ of a veering triangulation Δ , $\tilde{\Delta}$ has a strongly connected flow graph if and only if Δ has a strongly connected flow graph. The flow graph of $\tilde{\Delta}$ is a covering of that of Δ , and so the forward implication is easy. For the converse, if the flow graph of $\tilde{\Delta}$ is not strongly connected, it contains infinitesimal cycles of a wall. Such a wall projects down to a wall of Δ by Definition 3.3, thus Δ contains infinitesimal cycles as well.

Remark 3.6 With this theorem, we can also discuss how walls are related to the material in [Landry et al. 2023].

We first recall some terminology from [Landry et al. 2023]. The *dual graph* to a veering triangulation Δ is defined to be the 1-skeleton, ie the branch locus, of the unstable branched surface B , with the edges oriented by the coorientations on $\Delta^{(2)}$. An *AB cycle* is then defined to be a loop carried by the dual graph which only makes *antibranching turns*, ie does not follow components of the branch locus at each vertex of B . The lift of an AB cycle to the universal cover \tilde{M} determines a properly embedded plane carried by \tilde{B} , by taking the union over all sectors lying below the lifted AB cycle. Such a plane is called a *dynamic plane*. The total number of lifted AB cycles in a dynamic plane is defined to be the *width* of the dynamic plane. If there are 2 or more such lifted AB cycles, the region bounded in between them is called an *AB region*. For more details, we refer the reader to [Landry et al. 2023, Sections 2.2 and 3.1].

Now one can count that there are $\lceil \frac{1}{2}w \rceil$ or w adjacent AB cycles within a wall of width w , depending on whether the wall is twisted or not. These lift to w lifted AB cycles in a dynamic plane, which bound AB

regions in between. Conversely, if in the universal cover \tilde{B} carries a dynamic plane containing an AB region, then the portion of the flow graph within the AB region carries a line which only has incoming edges, hence quotients to an infinitesimal cycle of a wall.

3.4 Reduced flow graph

Given Theorem 3.5, a way of arranging for strong connectivity is to simply throw away all the infinitesimal cycles.

Definition 3.7 The *reduced flow graph* Φ_{red} of a veering triangulation is the flow graph with all infinitesimal cycles in walls and the edges that enter the cycles removed.

Thus by Theorem 3.5, Φ_{red} is strongly connected. Because of the simple nature of infinitesimal cycles, Φ_{red} inherits some of the properties of Φ . We end this section with two examples of this, which will be useful in Section 5.

Lemma 3.8 For every cycle c of Φ , c or c^2 is homotopic to a cycle of Φ_{red} in M .

Proof The cycles of Φ are those of Φ_{red} and the infinitesimal cycles of walls in Φ . Each infinitesimal cycle of a wall is homotopic (isotopic, even) to a boundary cycle in M , unless the wall has even width and is twisted, in which case one has to double the infinitesimal cycle $c_{w/2}$ before it is homotopic to a boundary cycle in M . \square

Lemma 3.9 The components of $B \setminus \Phi_{\text{red}}$ are annulus or Möbius bands with tongues.

Moreover, the arcs of attachment of the triangular sectors crisscross along the annulus/Möbius band. More precisely, the arcs on the annulus/Möbius band lift to $y = \pm x + 2i/w$ for $i \in \mathbb{Z}$ in the universal cover $[0, 1] \times \mathbb{R}$, for some $w \geq 1$. These arcs are subintervals of the branch locus of B and are oriented downwards (ie decreasing y in the above model). See Figure 10.

Proof By Lemma 2.10, it suffices to look at the components that contain infinitesimal cycles of some wall. These components are obtained by gluing components of $B \setminus \Phi$ along the infinitesimal cycles and the edges that enter them.

To describe these gluings more precisely, let us introduce some terminology. For a component J of $B \setminus \Phi$, we call the vertex of a tongue not lying on the annulus/Möbius band the *tip* of the tongue. If one cuts open ∂J along all the tips, the resulting set has a natural structure as an oriented train track, where the orientation is induced from that of Φ . Under this structure, each component of this set will be a cycle with some entering edges. We call each such component a *crown*.

To obtain a component of $B \setminus \Phi_{\text{red}}$ containing infinitesimal cycles, certain components of $B \setminus \Phi$ are glued together by identifying crowns in their boundaries. Under such a gluing, the union of annuli/Möbius




Figure 10: The complementary regions of Φ_{red} in B are annuli/Möbius bands with tongues. The attaching arcs form a crisscross pattern on the annulus/Möbius band. Here we show the situation in the neighborhood of the width 4 wall from Figure 4 (top) and the width 2 wall from Figure 5 (bottom).

bands is an annulus/Möbius band, while collections of tongues that share a tip are glued together to form individual tongues, showing that such a component of $B \setminus \Phi_{\text{red}}$ is an annulus/Möbius band with tongues as well.

The crisscross pattern of the arcs of attachment follows from the zigzag pattern for components of $B \setminus \Phi$ and from the description of the gluing.

We remark that this argument actually shows that all the infinitesimal cycles contained in a component of $B \setminus \Phi_{\text{red}}$ belong to a single wall. In this case, w in the statement of the lemma is the width of the wall. \square

Remark 3.10 We remark that if $w \geq 3$ for a component of $B \setminus \Phi_{\text{red}}$ as in the lemma, the tongues attached along the $y = x + 2i/(N + 1)$ arcs must all lie on one side and the tongues attached along the $y = -x + 2i/(N + 1)$ arcs must all lie on the opposite side, in the universal cover. In contrast, the sides of attachment do not have fixed patterns for $w = 1, 2$.

4 Finiteness of layered veering triangulations

In this section, we fix the gap in the proof of [Agol 2011, Theorem 6.2]. For completeness, we explain the setup of the original theorem then proceed to show a way of fixing its proof, albeit with a weaker bound.

We first have to explain layered veering triangulations.

Let $\phi: S_{g,n} \rightarrow S_{g,n}$ be a pseudo-Anosov homeomorphism on a finite type surface with

$$\chi(S_{g,n}) = 2 - 2g - n < 0.$$

S° will denote the surface obtained by removing the singularities of the stable and unstable foliations for ϕ , and ϕ° will denote the restriction of ϕ to S° . Write $T(\phi^\circ)$ for the mapping torus of ϕ° .

In [Agol 2011], it is shown that there exists a periodic folding sequence of train tracks $\tau_0 \rightsquigarrow \dots \rightsquigarrow \tau_N$, ie train tracks such that τ_{i+1} is obtained from τ_i using a folding move and $\phi(\tau_N) = \tau_0$, on S° . The sequence of ideal triangulations δ_i of S° dual to τ_i are then related to one another by diagonal switches in quadrilaterals, and ϕ° sends δ_N to δ_0 .

A veering triangulation can be constructed in the following way: Start with δ_0 on S° and attach a flat tetrahedron to the bottom that effects the diagonal switch from δ_0 to δ_1 . The bottom boundary of the complex can be identified with (S°, δ_1) . We inductively add tetrahedra to the bottom, until the bottom boundary of the complex can be identified with (S°, δ_N) . Finally, we glue this bottom boundary to the top boundary to get a triangulation Δ of $T(\phi^\circ)$, using the fact that ϕ° sends δ_N to δ_0 . We color an edge of δ_i that is dual to a small branch according to the direction of smoothing at the endpoints, using Figure 11. It can be checked that this determines a well-defined edge coloring of Δ which makes it into a veering triangulation.

We construct a directed graph according to this description. Start with a set of e vertices given by the set of edges in δ_0 . Notice that δ_1 and δ_0 differ by one edge exactly, say $\delta_1 \setminus e_1 = \delta_0 \setminus e_0$. Add a vertex corresponding to e_1 , and add three edges going from the three elements of $E(\delta_0)$ dual to the three branches of τ_0 which fold onto the branch of τ_1 dual to e_1 . One of these edges goes from e_0 to e_1 , call this edge *vertical*, and call the other two edges *slanted*. Inductively, for each i , add a vertex that corresponds to




Figure 11: Defining colors of edges of δ_i that are dual to small branches. This determines a well-defined edge coloring of Δ which makes it into a veering triangulation.




Figure 12: The cut open flow graph for a layered veering triangulation. There is a natural identification between adjacent layers $E(\delta_i)$ and $E(\delta_{i+1})$, which is represented in the figure by moving one element directly downwards. This divides the cut open flow graph into columns. The cut open reduced flow graph for the veering cellulation can be obtained by deleting some columns.

the new edge in δ_i and add three edges according to the folding move $\tau_i \rightsquigarrow \tau_{i+1}$, one vertical and two slanted. Call the resulting directed graph the *cut open flow graph* $\Phi \setminus S$. See Figure 12, where we draw vertical and slanted edges as vertical and slanted respectively.

The flow graph of Δ can be obtained from $\Phi \setminus S$ by identifying the vertices at the bottom corresponding to $E(\delta_N)$ to those at the top for $E(\delta_0)$ according to how ϕ° sends δ_N to δ_0 .

Now define another directed graph G by setting the set of vertices to be the set of edges in δ_0 , and placing an edge from i to j for every directed edge path in the cut open flow graph $\Phi \setminus S$ which starts at i and ends at $(\phi^\circ)^{-1}(j)$. Let A be the adjacency matrix of G (which we defined in Definition 3.1).

We note that A can also be defined as the transition matrix describing how weights on the branches of τ_0 distribute under the sequence of folding moves $\tau_0 \rightsquigarrow \dots \rightsquigarrow \tau_N$ and the return map ϕ° .

Meanwhile, let $\lambda = \lambda(\phi) > 1$ be the dilatation of ϕ . The train tracks τ_i in fact carry the unstable measured lamination of ϕ , hence the transverse measure on the leaves of the foliation collapse down to weights on the branches of τ_0 . These in turn define an eigenvector (w_i) of A with eigenvalue λ .

The assertion of [Agol 2011, Theorem 6.2] is that if $\lambda^{2g-2+\frac{2}{3}n} \leq P$, then the veering triangulation Δ has at most $\frac{1}{2}(P^9 - 1)$ tetrahedra.

We recall the proof presented in [Agol 2011], but phrased in the language here. We first claim that G has at least $2N + e$ edges. This is because for every slanted edge in $\Phi \setminus S$, we can take an edge path from $E(\delta_0)$ to $E(\delta_N)$ by concatenating paths of vertical edges to the back and front of the slanted edge. Also, for every vertex i in $E(\delta_0)$, there is an edge path from i to $E(\delta_N)$ that consists entirely of vertical edges. Since there are $2N$ slanted edges and e vertices in $E(\delta_0)$, we find $2N + e$ edges in G .

Also notice that by an index calculation, we have $e \leq 9(2g - 2 + \frac{2}{3}n)$. See [Agol 2011, Lemma 6.1] for details of this.

Now assume for the moment that G is strongly connected. Then [Agol 2011] proceeds by using an estimate of Ham and Song [2007]. We repeat Ham and Song's argument here since it is a good warm-up for the proof which we will present later. Fix a vertex i of G and fix a spanning tree T of G rooted at i , ie T is a subgraph of G which is a tree, and for every vertex j of G there is an edge path from j to i within T . Then there are at least $2N + 1$ edges in $G \setminus T$, and for every edge d in $G \setminus T$, we can find a path in G ending at i and of length e by prepending to the path traveling across d then to i within T an arbitrary path at the front. These paths will be distinct; hence using the fact that (w_i) is an eigenvector of A^e of eigenvalue λ^e , and summing over the paths, we have $\lambda^e w_i \geq \sum w_j \geq (2N + 1) \min_j w_j$. Taking the minimum over the left hand side, $\lambda^e \min w_i \geq (2N + 1) \min w_j$; hence $N \leq \frac{1}{2}(\lambda^e - 1) \leq \frac{1}{2}(P^9 - 1)$.

The problem with this argument however, is that G is not always strongly connected.

In fact, G is strongly connected if and only if the flow graph Φ is strongly connected. For if Φ is strongly connected, then for every pair of vertices (i, j) in $E(\delta_0)$, their images can be connected by an edge path α in Φ . The preimage of such an edge path under $\Phi \setminus S \rightarrow \Phi$ is a collection of paths $\alpha_1, \dots, \alpha_s$, for which the ending point of α_i , which lies in $E(\delta_N)$, is sent by ϕ to the starting point of α_{i+1} , which lies in $E(\delta_0)$. Hence the α_i define an edge path α' in G connecting i to j .

Conversely, if Φ is not strongly connected, then we can find a infinitesimal cycle c of a wall. As above, c lifts to a collection of paths c_1, \dots, c_s in $\Phi \setminus S$ and determines a cycle c' in G . In fact, c_i will consist entirely of vertical edges by the definition of the infinitesimal cycles of a wall. The vertices of G that lie in c' are exactly those elements of $E(\delta_0)$ that have image in c . Moreover, c' has no outgoing edges in G' , for otherwise there is an outgoing path of c in Φ .

This motivates us to consider instead the full subgraph of G obtained by restricting to the set of vertices that have image in the reduced flow graph Φ_{red} in Φ . We call this subgraph G_{red} . By the argument above, G_{red} can also be obtained by removing c' for all infinitesimal cycles c of Φ .

We also set the *cut open reduced flow graph*, $\Phi_{\text{red}} \setminus S$, to be the preimage of $\Phi_{\text{red}} \subset \Phi$ in $\Phi \setminus S$. Let $E_{\text{red}}(\delta_i)$ be the subset of $E(\delta_i) \subset \Phi \setminus S$ that lie in $\Phi_{\text{red}} \setminus S$. G_{red} can be obtained by gluing $E_{\text{red}}(\delta_N)$ at the bottom of $\Phi_{\text{red}} \setminus S$ to $E_{\text{red}}(\delta_0)$ at the top. Hence arguing as above, we can see that G_{red} is strongly connected.

The strategy now is to apply Ham and Song's argument on G_{red} to bound the number of vertices in Φ_{red} , then bound the number of vertices in $\Phi \setminus \Phi_{\text{red}}$.

We first set up some notation. Let N' be the number of vertices in Φ_{red} . Let e' be the number of elements in $E_{\text{red}}(\delta_0)$.

We claim that G_{red} has at least $2N' + e'$ edges. This follows from the same argument as for G : we can find an edge path in $\Phi_{\text{red}} \setminus S$ from $E_{\text{red}}(\delta_0)$ to $E_{\text{red}}(\delta_N)$ for every slanted edge in $\Phi_{\text{red}} \setminus S$ and for every vertex in $E_{\text{red}}(\delta_0)$.

Meanwhile, recall that $G \setminus G_{\text{red}}$ is the union of c' for the infinitesimal cycles c of Φ , where each c' has no outgoing edges. Hence the adjacency matrix A_{red} of G_{red} is a submatrix of A where all other entries in the same row are zeros. As a consequence, $(w_i)_{i \in G_{\text{red}}}$ is an eigenvector of A_{red} with eigenvalue λ .

Hence repeating Ham and Song's argument, we have $N' \leq \frac{1}{2}(\lambda^{e'} - 1)$.

Now suppose we have a width w wall. For simplicity, first suppose that the wall is untwisted. Then there are $w - 1$ infinitesimal cycles and 2 boundary cycles of the wall, all of the same length h . The $(w - 1)h$ vertices in the infinitesimal cycles are discarded in Φ_{red} , but the $2h$ vertices in the boundary cycles remain. Similarly, if the wall is twisted, there are $(w - 1)h$ vertices discarded while $2h$ vertices remain, for an appropriate h . In other words, if we let W be the maximum width of a wall in the veering triangulation Δ , then for every vertex in a boundary cycle, there are at most $\frac{1}{2}(W - 1)$ vertices in the infinitesimal cycles, which we discard when passing to Φ_{red} .

Meanwhile, as discussed in Section 3.2, each vertex of Φ can appear at most twice in the collection of all boundary cycles. Hence we conclude that the number of discarded vertices $N - N'$ is at most $N'(W - 1)$, or $N \leq N'W$. Thus it remains to bound W .

To that end, let c_2, \dots, c_W be the infinitesimal cycles in a width W wall, and let c_1 and c_{W+1} be the boundary cycles of the wall. Again, for simplicity, first suppose that the wall is untwisted. Then there are corresponding disjoint cycles c'_1, \dots, c'_{W+1} in G . The length of each cycle c'_i is given by the number of paths in the lift of c_i in the cut open flow graph. But this number is also equal to the intersection number of c_i with S° in $T(\phi^\circ)$. Since the c_i are parallel to each other, this number is of the same value L for each $i = 1, \dots, W + 1$ and hence at most $e/(W + 1)$ by the pigeonhole principle. Similarly, if the wall is twisted, the infinitesimal and boundary cycles determine cycles in G of lengths at most $2e/(W + 1)$.

Now pick a vertex of Φ_{red} that lies in c_1 . One of its preimages in $\Phi_{\text{red}} \setminus S$ has three incoming edges. The cycle c_1 passes through one of these edges, and for each of the remaining two, we can construct a path from $E_{\text{red}}(\delta_0)$ to $E_{\text{red}}(\delta_N)$ by concatenating vertical edges to the back and front of the edge. These two paths determine two edges that enter c'_1 at a vertex i in G_{red} . Since G_{red} is strongly connected, we can locate two edge paths in G_{red} of length L that end at the two edges respectively. Suppose these two paths start at vertices j and k . Then using the fact that (w_i) is an eigenvector of A_{red}^L with eigenvalue λ^L , we have $\lambda^L w_i \geq w_i + w_j + w_k$, which implies

$$w_i \geq \frac{2}{\lambda^L - 1} \min_j w_j.$$

But for every two vertices i and j in G_{red} , by strong connectivity, there is an edge path of length $\leq e'$ from i to j ; hence $w_i \geq \lambda^{-e'} w_j$. So $\max_j w_j / \min_i w_i \leq \lambda^{e'}$. Applying this to the above inequality,

$$\lambda^{e'} \geq \frac{2}{\lambda^L - 1}, \quad \lambda^L \geq 2\lambda^{-e'} + 1.$$

Hence,

$$\frac{2e}{W + 1} \geq L \geq \frac{\log(2\lambda^{-e'} + 1)}{\log \lambda}, \quad W \leq \frac{2 \log \lambda^e}{\log(2\lambda^{-e'} + 1)} - 1.$$

Putting everything together,

$$N \leq \frac{\lambda^{e'} - 1}{2} \left(\frac{2 \log \lambda^e}{\log(2\lambda^{-e'} + 1)} - 1 \right) \leq \frac{\lambda^e - 1}{2} \left(\frac{2 \log \lambda^e}{\log(2\lambda^{-e} + 1)} - 1 \right) \leq \frac{P^9 - 1}{2} \left(\frac{2 \log P^9}{\log(2P^{-9} + 1)} - 1 \right).$$

We record this as a theorem.

Theorem 4.1 *If M is the punctured mapping torus of a pseudo-Anosov homeomorphism $\phi: S_{g,n} \rightarrow S_{g,n}$, where the normalized dilatation $\lambda(\phi)^{2g-2+\frac{2}{3}n} \leq P$, then M has a veering triangulation with at most*

$$\frac{P^9 - 1}{2} \left(\frac{2 \log P^9}{\log(2P^{-9} + 1)} - 1 \right)$$

tetrahedra.

Remark 4.2 Note that

$$\frac{P^9 - 1}{2} \left(\frac{2 \log P^9}{\log(2P^{-9} + 1)} - 1 \right)$$

is asymptotically $\frac{9}{2} P^{18} \log P$ as $P \rightarrow \infty$, so we have worsened the exponent on the bound in [Agol 2011, Theorem 6.2] by a factor of $2 + \varepsilon$.

We remark that the original bound of $\frac{1}{2}(P^9 - 1)$ still holds for veering triangulations that have strongly connected flow graph. However, it is not clear if this is the case for “most” ϕ or if there is a way to tell if this is the case just from ϕ .

It also seems likely that there is room for improvement for our bound in the general case. For example, it should be possible to obtain better bounds on W by finding more paths that enter c'_1 , which should be easy when L is large. When W or L is large, it should also be possible to bound e' more effectively than just using e , which is what we have done here. See Section 6 for another discussion on how one might improve the bound.

5 Pseudo-Anosov flows

In this section we will reprove Schleimer and Segerman’s result that a veering triangulation induces a pseudo-Anosov flow without perfect fits on suitable Dehn fillings. We use the following notation to simplify the statement. If $s = (s_i)$ is a collection of slopes on each boundary component of \overline{M} , we write $M(s)$ to mean the closed 3-manifold obtained by Dehn filling \overline{M} along the slopes recorded by s . If $a = (a_i)$ and $b = (b_i)$ are collections of slopes on each boundary component of \overline{M} , then by $|\langle a, b \rangle| \geq n$ we mean that the geometric intersection numbers between a_i and b_i on each boundary component is at least n .

Theorem 5.1 *Suppose M admits a veering triangulation. Let $l = (l_i)$ denote the collection of ladderpole curves on each boundary component. Then $M(s)$ admits a transitive pseudo-Anosov flow ϕ if $|\langle s, l \rangle| \geq 2$. Furthermore, there are closed orbits c_i isotopic to the cores of the filling solid tori, such that each c_i is $|\langle s_i, l_i \rangle|$ -pronged, and ϕ is without perfect fits relative to $\{c_i\}$.*

We will recall the definitions of pseudo-Anosov flows, transitivity, and no perfect fits in Section 5.4.

A subtle point of the theorem is that there are actually two common notions of pseudo-Anosov flows on 3-manifolds in the literature, which we differentiate by calling them topological pseudo-Anosov flows and smooth pseudo-Anosov flows. Theorem 5.1 holds for both notions, due to Theorem 5.11. We will explain this technicality more in Sections 5.4 and 5.5.

The proof can be outlined as follows. We first thicken up the reduced flow graph Φ_{red} in M to $N(\Phi_{\text{red}})$ by replacing its edges with flow boxes. This set can be considered as a subset of a neighborhood of the unstable branched surface, $N(B)$, naturally. By understanding the complement of Φ_{red} in B , we are able to glue faces of $N(\Phi_{\text{red}})$ across its complementary regions in $N(B)$. Similarly, by understanding the complement of B in M , we are able to glue faces of $N(B)$ across its complementary regions in $M(s)$. These gluings preserve the (singular) 1-dimensional foliation on the flow boxes, hence that descends down to an (honest) 1-dimensional foliation on $M(s)$, which can be parametrized into a topological flow, for which we show is pseudo-Anosov, transitive, and without perfect fits (relative to the orbits $\{c_i\}$).

Remark 5.2 It is not too difficult to show that a veering triangulation induces a (topological) pseudo-Anosov flow on $M(s)$ for $|\langle s, l \rangle| \geq 2$, using the tool of dynamic pairs developed by Mosher [1996]. Specifically, one can apply the proof of [Mosher 1996, Proposition 2.6.2] to (B, Φ) to produce a dynamic pair in $M(s)$, which by [Mosher 1996, Theorem 3.4.1] gives rise to a pseudo-Anosov flow. The more challenging part however, at least from this approach, is to show that such a pseudo-Anosov flow is transitive and has no perfect fits.

In our proof, we use a lot of the same ideas as [Mosher 1996], but most notably we skip over constructing the “stable branched surface” in a dynamic pair, and instead construct a pseudo-Anosov flow directly from the “unstable branched surface” B and the “dynamic train track” Φ_{red} , to use the terminology from [Mosher 1996]. This allows us to analyze the pseudo-Anosov flow using the special properties of Φ_{red} and B , proving transitivity and no perfect fits.

5.1 Thickening up Φ_{red}

We know that Φ_{red} is strongly connected by construction. This is equivalent to its adjacency matrix $A \in \text{Hom}(\mathbb{R}^{V(\Phi_{\text{red}})}, \mathbb{R}^{V(\Phi_{\text{red}})})$ being irreducible. As such, by the Perron–Frobenius theorem, A has a positive eigenvector (w_v) with eigenvalue $\lambda \geq 1$. We know that $\lambda > 1$ since each vertex of Φ_{red} has three incoming edges. Meanwhile, A being irreducible implies A^T is irreducible as well, and so the latter has a positive eigenvector (w'_v) with the same eigenvalue $\lambda > 1$.

We remark in passing that each vertex of Φ_{red} having 3 incoming edges in fact implies that $\lambda = 3$ and $w_v = 1$ for all v . This knowledge, however, will play no role in the construction at all. We simply wish to point out that the value of λ has nothing to do with the dilatation factor of the monodromy when Δ is layered.




Figure 13: A flow box.

Recall that Φ_{red} naturally sits on B inside M . We will thicken up Φ_{red} by replacing each edge e of Φ_{red} going from v to w by a flow box

$$Z_e \cong \{(s, u, t) \in \mathbb{R}^3 : |s| \leq w_v \lambda^{t-1}, |u| \leq w'_w \lambda^{-t}, t \in [0, 1]\}.$$

See Figure 13. There are two 2–dimensional foliations on Z_e : the first one by leaves of the form $\{u = u_0 \lambda^{-t}\}_{u_0}$, which we will call the stable foliation, and the second one by leaves of the form $\{s = s_0 \lambda^{t-1}\}_{s_0}$, which we will call the unstable foliation. There is also an oriented 1–dimensional foliation on Z_e by curves $\{(s_0 \lambda^{t-1}, u_0 \lambda^{-t}, t) : t \in [0, 1]\}_{s_0, u_0}$, oriented by decreasing t . Notice that the leaves of the stable and unstable foliations intersect transversely along leaves of the oriented 1–dimensional foliation.

Give Z_e its Euclidean metric induced from \mathbb{R}^3 . There is a natural way of gluing up the Z_e across their top and bottom faces at the vertices of Φ_{red} , preserving the metric on those faces, by definition of w_v and w'_w . A priori there is a freedom for Z_e to twist along e , we get rid of this by requiring that the framing on Z_e induced from its u coordinate match up with the framing on e induced from B . See Figure 14. We call the resulting set $N(\Phi_{\text{red}})$, since it is a regular neighborhood of Φ_{red} in M .

Note that the oriented 1–dimensional foliations on Z_e piece together to give a decomposition of $N(\Phi_{\text{red}})$. This decomposition is almost an oriented 1–dimensional foliation, except the “leaves” are oriented 1–manifolds possibly with train track singularities branching off in forward and backward directions, arising from the multiple outgoing and incoming edges at the vertices. Despite this, for convenience we will still refer to this decomposition as an oriented 1–dimensional foliation.

Similarly, the stable and unstable foliations on Z_e each piece together to give a decomposition of $N(\Phi_{\text{red}})$ which is almost a 2–dimensional foliation, except that the “leaves” have branching. We will refer to these as the stable and unstable foliations on $N(\Phi_{\text{red}})$ respectively.

Denote $\partial^s Z_e = \{u = \pm w'_w \lambda^{-t}\}$ and $\partial^u Z_e = \{s = \pm w_v \lambda^{t-1}\}$. We will call $\partial^s N(\Phi_{\text{red}}) := \bigcup_e \partial^s Z_e$ the *stable boundary* of $N(\Phi_{\text{red}})$ and $\partial^u N(\Phi_{\text{red}}) := \bigcup_e \partial^u Z_e$ the *unstable boundary* of $N(\Phi_{\text{red}})$. We will also call the collection of the closures of intervals where the interior of the top and bottom faces of Z_e




Figure 14: Thickening up Φ_{red} to $N(\Phi_{\text{red}})$, which naturally embeds in $N(B)$.

meets $\partial^s N(\Phi_{\text{red}})$ or $\partial^u N(\Phi_{\text{red}})$ the branch locus. These are exactly the places where the leaves of the oriented 1–dimensional foliation on $N(\Phi_{\text{red}})$ have branching.

Now consider an I –fibered neighborhood $N(B)$ of B in the cusped model. This means that $N(B)$ is a closed regular neighborhood of B , with a map $M \rightarrow M$ restricting to a projection on $N(B) \rightarrow B$ with I –fibers and a homeomorphism outside of $N(B)$, and so that the boundary of $N(B)$ can be decomposed into surfaces with boundary, the interior on which the projection is a local homeomorphism, and the boundary curves correspond to the components of the branch locus of B . We will often conflate a component of the branch locus of B with the corresponding circle on $\partial N(B)$. We will also call the collection of circles in the latter setting the branch locus of $N(B)$.

$N(\Phi_{\text{red}})$ can be arranged to be a subset of $N(B)$, in such a way that $\partial^u N(\Phi_{\text{red}}) \subset \partial N(B)$, the branch locus of $N(\Phi_{\text{red}})$ is a subset of the branch locus of $N(B)$, and $N(\Phi_{\text{red}})$ is saturated with respect to the I –fibering of $N(B)$.

Note that we do not require the intervals $[-w_v \lambda^{t_0-1}, w_v \lambda^{t_0-1}] \times \{u_0\} \times \{t_0\} \subset Z_e \subset N(\Phi_{\text{red}})$ to coincide with the I –fibers of $N(B)$. Indeed, this is impossible since the branch locus of $N(\Phi_{\text{red}})$ are parallel along such intervals, while branch locus of $N(B)$ project to transverse curves on B . One can arrange for this property by “splitting” $N(\Phi_{\text{red}})$ slightly, but since this does not aid our construction, we will not do so.

5.2 Gluing along the stable boundary

The next step is to glue $N(\Phi_{\text{red}})$ along its stable boundary across its complementary regions in $N(B)$, so that the foliations on $N(\Phi_{\text{red}})$ descend to respective foliations on $N(B)$. To perform the gluing, we have to understand the complementary regions in question.




Figure 15: Complementary regions of $N(\Phi_{\text{red}})$ in $N(B)$. We construct an I -fibering conjugating the flows on the stable face, and collapse along the I -fibers.

By Lemma 3.9, the complementary regions of Φ_{red} in B are annuli or Möbius bands with tongues. The complementary regions of $N(\Phi_{\text{red}})$ in $N(B)$ are I -fibered neighborhoods of these; see Figure 15 (where the red intervals will come into play later). Fix one of these components K . The boundary of K is naturally divided into $\partial^s K \cup \partial^u K$, where $\partial^s K$ is identified to $\partial^s N(\Phi_{\text{red}})$ in M and $\partial^u K \subset \partial N(B)$. In particular $\partial^s K$ inherits the oriented 1-dimensional foliation on $\partial^s N(\Phi_{\text{red}})$. Furthermore, if we call the I -fibers over the tips of the tongues of $B \setminus \Phi_{\text{red}}$ the *cusps* of K , $\partial^s K$ can be divided along the cusps into two or one components (depending on whether the corresponding component of $B \setminus \Phi_{\text{red}}$ is an annulus or a Möbius band with tongues respectively). Call these components the *stable faces* of K .

Meanwhile, recall that $N(\Phi_{\text{red}})$, hence $\partial^s N(\Phi_{\text{red}})$, inherits the Euclidean metric of the flow box Z_e . The oriented 1-dimensional foliation contracts the metric in the transverse direction. Hence the oriented foliation on each stable face of K has exactly one S^1 leaf, and all other leaves enter through the cusps of K and spiral into the S^1 leaf.

More rigorously, parametrize the oriented 1-dimensional foliation on $\partial^s N(\Phi_{\text{red}})$ in some way, for example according to the t coordinate of each Z_e , so that we can consider it as a forward semiflow. Take the intervals $[-w_v, w_v] \times \{\pm w'_w \lambda^{-1}\} \times \{1\} \subset Z_e \subset N(\Phi_{\text{red}})$, with the Euclidean metrics, as sections of the flow, and note that the first return map contracts the metric by λ^{-1} , then apply Banach fixed point theorem.

We claim that there is an I -fibering of K , transverse to the stable faces and parallel to $\partial^u K$, for which the induced homeomorphism between the stable faces preserves the oriented 1-dimensional foliations on them, ie send leaves to leaves in an orientation preserving way.




Figure 16: The steps for constructing the I -fibering on each component of $N(B) \setminus N(\Phi_{\text{red}})$.

Suppose first that the corresponding component of $B \setminus \Phi_{\text{red}}$ is an annulus with tongues. Choose short local sections to the foliation near the unique S^1 leaf on each stable face of K , call them J and J' . The return maps of the oriented foliations on J and J' , which we call h and h' , are *contracting* by our analysis above, ie $h^k(J)$ is a strictly decreasing collection of subintervals with $\bigcap_{k=1}^{\infty} h^k(J)$ being a point, and similarly for $h'^k(J')$. Note that these points could lie on endpoints of J and J' if one side of the annulus has no tongues attached.

Construct the I -fibering on K in the following steps:

- (1) Define the I -fibering on the cusps of K , where the I -fibering is uniquely determined but degenerate, in the sense that the I -fibers are just points.
- (2) Extend the I -fibering to the triangular faces of $\partial^u K$, so that the base of the triangles, which are among the branch locus of $N(B)$, are I -fibers.
- (3) Extend the I -fibering across subintervals of the leaves of the foliation on the stable faces which start on the branch locus and end at the interior of J or J' .
- (4) Extend the I -fibering across subintervals of the leaves which start on the interior of the cusps of K and end at the interior of J or J' , using the fact that the union of these subintervals is a finite union of rectangles foliated as products.
- (5) Apply Lemma 5.3 below by taking f to be the homeomorphism induced by portion of the I -fibering already defined. Then use the extended f given by the lemma to construct the I -fibering between J and J' , and complete the construction by extending the I -fibering along the leaves as they go around the stable faces.

See Figure 16 for a graphical summary of these steps.

Lemma 5.3 Let $h: J \rightarrow J$ and $h': J' \rightarrow J'$ be injective contracting maps on intervals. Let

$$f: J \setminus h(J) \rightarrow J' \setminus h'(J')$$

be a given homeomorphism. Then there is a unique way of extending f to a homeomorphism $J \rightarrow J'$ such that $h'f = fh$.

Proof Write $J^{(k)} = h^k(J)$ and $J'^{(k)} = h'^k(J')$. By assumption, $J^{(k)}$ is a shrinking subinterval of J which contracts to a point s and $J'^{(k)}$ is a shrinking subinterval of J' which contracts to a point s' . Extend f by defining it to be $h'^k f(h^k)^{-1}$ on $J^{(k)} \setminus J^{(k+1)}$. Finally define $f(s) = s'$. \square

If the corresponding component of $B \setminus \Phi_{\text{red}}$ is a Möbius band with tongues instead, we choose two local sections \bar{J}_1 and \bar{J}_2 on the one stable face near the S^1 leaf, so that we have return maps $\bar{J}_1 \rightarrow \bar{J}_2$ and $\bar{J}_2 \rightarrow \bar{J}_1$ for the oriented foliation. (We can pick \bar{J}_1 first, then pick \bar{J}_2 near \bar{J}_1 using the fact that the return map $\bar{J}_1 \rightarrow \bar{J}_1$ is contracting.) Let s_1 and s_2 be where the S^1 leaf meets \bar{J}_1 and \bar{J}_2 . Then s_1 and s_2 have to be in the interior of \bar{J}_1 and \bar{J}_2 in this setting. Let J_i be the subinterval of \bar{J}_i to the left of s_i , and J'_i be the subinterval of \bar{J}_{i+1} to the right of s_{i+1} (indices taken mod 2, and left/right measured relative to the direction of the flow and a fixed orientation on the stable face). First construct the I -fibering on subintervals of leaves that start on cusps of K and end at the interior of \bar{J}_1 or \bar{J}_2 , then apply Lemma 5.4 below (with $m = 2$) to construct the I -fibering across a neighborhood of the periodic trajectory as above.

Lemma 5.4 Let $h_i: J_i \rightarrow J_{i+1}$ and $h'_i: J'_i \rightarrow J'_{i+1}$ be injective contracting maps on intervals J_1, \dots, J_m and J'_1, \dots, J'_m (indices taken mod m). Here, contracting means $\{h_{mk+i-1} \cdots h_i(J_i)\}_k$ is a decreasing collection of subintervals and $\bigcap_{k=1}^{\infty} h_{mk+i-1} \cdots h_i(J_i)$ is a single point for each i , and similarly for h'_i . Let $f_i: J_i \setminus h_{i-1}(J_{i-1}) \rightarrow J'_i \setminus h'_{i-1}(J'_{i-1})$ be given homeomorphisms. Then there is a unique way of extending f_i to homeomorphisms $J_i \rightarrow J'_i$ such that $h'_i f_i = f_{i+1} h_i$ for all i .

Proof This can be proved as in the previous lemma (which is the “baby case” when $m = 1$), with more annoying bookkeeping. We remark that we stated this lemma for all $m \geq 1$ because we will need this generality in the next subsection. \square

Remark 5.5 The homeomorphism across the stable faces induced by such an I -fibering cannot be made to preserve the Euclidean structure on $\partial^s N(\Phi_{\text{red}})$ in general. In particular, the homeomorphism will not, in general, take an interval of the form $[-w_v \lambda^{t_0-1}, w_v \lambda^{t_0-1}] \times \{\pm w'_w \lambda^{-t_0}\} \times \{t_0\} \subset Z_e \subset N(\Phi_{\text{red}})$ to another interval of this form.

Also, the homeomorphism cannot be chosen to be a diffeomorphism in general (with respect to the smooth structures induced by the Euclidean structures). This arises from the fact that one cannot replace “homeomorphisms” by “diffeomorphisms” in Lemma 5.4 in general by Kopell’s lemma [1970].

Now collapse across all such components K along the I -fiberings. This collapses $N(\Phi_{\text{red}})$ onto $N(B)$. The oriented 1-dimensional foliation on $N(\Phi_{\text{red}})$ descends to a decomposition of $N(B)$. The “leaves” of which are still not all 1-manifolds, but now we only have branching in the backwards direction, since all

the forward branching occurs along some cusp of some K , and we have collapsed all forward trajectories starting from those points. As before, we will still call this decomposition an oriented 1–dimensional foliation for convenience of notation.

Similarly, the stable and unstable foliations on $N(\Phi_{\text{red}})$ descend to decompositions of $N(B)$, which we call the stable and unstable foliations on $N(B)$ respectively. The leaves of the stable foliation on $N(B)$ are in fact 2–manifolds by the same argument as the last paragraph, but those of the unstable foliation on $N(B)$ have branching.

Remark 5.6 For the purposes of the next step of the construction, it is worth noting that when restricted to the unstable boundary, the collapse takes $\partial^u N(\Phi_{\text{red}})$ onto $\partial N(B)$, under which the restriction of the oriented 1–dimensional foliation of $N(\Phi_{\text{red}})$ to $\partial^u N(\Phi_{\text{red}})$ descends to the restriction of the oriented 1–dimensional foliation of $N(B)$ to $\partial N(B)$.

In particular, we can recover the latter in the following way: take the image of Φ_{red} on the boundary of $M \setminus B$, thicken it up within $\partial(M \setminus B)$ by replacing each edge by a flow box $\{(u, t) : |u| \leq w'_w \lambda^{-t}, t \in [0, 1]\}$ foliated by $\{(u_0 \lambda^{-t}, t) : t \in [0, 1]\}_{u_0}$ and piecing together their top and bottom edges at the vertices of Φ_{red} . The I –fiberings on the components K constructed above, when restricted to $\partial^u K$, will determine I –fiberings of the complementary regions, and we collapse along these I –fibers to get the oriented foliation on $\partial N(B)$.

Before we move on, we will construct some local sections to the foliation which will be used to prove no perfect fits in Section 5.4. Fix a component c of the branch locus of B . Consider the subset of double points of the branch locus of B that lie on c . These are each dual to some veering tetrahedron in Δ . Let V_c be the subset of vertices of Φ_{red} which are the bottom edges of one of these veering tetrahedra. Cyclically order $V_c = \{v_1, \dots, v_s\}$ according to the order in which c meets the corresponding double points.

Meanwhile, for a vertex v of Φ_{red} , let R_v be the union of the top faces of the flow boxes Z_e as e varies over the edges of Φ_{red} that exit v . Equivalently, this is also the union of the bottom faces of Z_e as e varies over the edges of Φ_{red} that enter v . These are embedded rectangles transverse to the foliation on $N(\Phi_{\text{red}})$.

For a collection $V_c = \{v_1, \dots, v_s\}$ as above, each $R_{v_i} \subset N(\Phi_{\text{red}})$ contains a subinterval of c . Here we remind the reader that we are using the same name for a component of the branch locus of B and the corresponding component of the branch locus of $N(B)$. Let $\partial^- R_{v_i}$ be the side of R_{v_i} which c enters, and let $\partial^+ R_{v_i}$ be the side of R_{v_i} which c exits. Note that $\partial^+ R_{v_i}$ and $\partial^- R_{v_{i+1}}$ lie on stable faces of the same component of $N(B) \setminus N(\Phi_{\text{red}})$ (indices taken mod s). We drew one instance of these as red intervals in Figure 15.

After collapsing along the components of $N(B) \setminus N(\Phi_{\text{red}})$, the image of the union $\bigcup_{i=1}^s R_{v_i}$ in $N(B)$ contains c , since for each component K of $N(B) \setminus N(\Phi_{\text{red}})$, the bases of the triangular faces of $\partial^u K$ are among the I –fibers that we chose. The images of $\partial^+ R_{v_i}$ and $\partial^- R_{v_{i+1}}$ may not match up away




Figure 17: The surfaces P_c are unions of images of R_{v_i} and connecting surfaces Q_i .

from c . However, what we will show is that we can at least find surfaces Q_i which are unions of finite subintervals of leaves of the 1–dimensional foliation on $N(B)$ which connect the image of $\partial^+ R_{v_i}$ to a subinterval of that of $\partial^- R_{v_{i+1}}$. See Figure 17, top.

Let K be the component of $N(B) \setminus N(\Phi_{\text{red}})$ which contains both $\partial^+ R_{v_i}$ and $\partial^- R_{v_{i+1}}$. To construct Q_i , consider the universal cover \tilde{K} of K . Note that there is a notion of height on \tilde{K} , with respect to the y –coordinate in the model described in Lemma 3.9, as opposed to K where that coordinate is circular. The oriented 1–dimensional foliations on the stable faces of K lift to oriented 1–dimensional foliations on the stable faces of \tilde{K} . Similarly, the I –fibering on K lifts to an I –fibering on \tilde{K} . A component of the branch locus of K is a subinterval of c connecting $\partial^+ R_{v_i}$ and $\partial^- R_{v_{i+1}}$. Lift $\partial^+ R_{v_i}$ and $\partial^- R_{v_{i+1}}$ together with the connecting subinterval to \tilde{K} . We abuse notation and still call the lifts of $\partial^+ R_{v_i}$ and $\partial^- R_{v_{i+1}}$ by the same name.

We claim that the crisscross property stated in Lemma 3.9 implies that the component of $\partial^\mu \tilde{K}$ containing an endpoint of $\partial^+ R_{v_i}$ does not lie below that containing the endpoint of $\partial^- R_{v_{i+1}}$ on the same side. To check this, let J be the component of $B \setminus \Phi_{\text{red}}$ corresponding to K and consider the universal cover \tilde{J} of J . Without loss of generality, we can assume that the lifted connecting subinterval corresponds to $y = -x$ in the branch locus of \tilde{J} , under the model described in Lemma 3.9. Notice that in this case $\partial^+ R_{v_i}$ lies along the stable face of \tilde{K} corresponding to $x = 0$ while $\partial^- R_{v_{i+1}}$ lies along that corresponding to $x = 1$.

Meanwhile, there are two sides to $[0, 1] \times \mathbb{R} \subset \tilde{J}$ for which its tongues are attached along. Correspondingly, the components of $\partial^\mu \tilde{K}$ also lie on two sides, and we need to check our claim on the two sides separately.

On the side for which the tongue is attached along $y = -x$, the component of $\partial^u \tilde{K}$ containing the endpoint of $\partial^+ R_{v_i}$ is the same as that containing the endpoint of $\partial^- R_{v_{i+1}}$, unless the tongue attached along $y = x - 2$ is done so on this same side (hence $w = 1$ or 2), in which case the former is higher up than the latter.

On the other side, the component of $\partial^u \tilde{K}$ containing the endpoint of $\partial^+ R_{v_i}$ is strictly higher up than that containing the endpoint of $\partial^- R_{v_{i+1}}$, unless $w = 1$ and the tongue attached along $y = x - 2$ is done so on the same side as that along $y = -x$, in which case the two component coincide.

This claim implies that if we transfer $\partial^+ R_{v_i}$ to the other stable face of \tilde{K} using the homeomorphism induced by the I -fibering, and call the image $\partial^{++} R_{v_i}$ temporarily, then the leaves of the 1-dimensional foliation passing through $\partial^{++} R_{v_i}$ will also pass through $\partial^- R_{v_{i+1}}$. Hence we can take the union of subintervals of leaves going between $\partial^{++} R_{v_i}$ and $\partial^- R_{v_{i+1}}$, project down to K then collapse to $N(B)$ to get Q_i .

Here we caution that $\partial^{++} R_{v_i}$ may not lie strictly above $\partial^- R_{v_{i+1}}$, that is, some of the leaves that form Q_i may be oriented from $\partial^- R_{v_{i+1}}$ to $\partial^{++} R_{v_i}$, as indicated in Figure 17. The important feature here, rather, is that $Q_i \cap \partial^+ R_{v_i} = \partial^+ R_{v_i}$ and $Q_i \cap \partial^- R_{v_{i+1}} \subset \partial^- R_{v_{i+1}}$.

Let P_c be the union of the images of R_{v_i} and Q_i . Notice that P_c is an immersed annulus with corners in $N(B)$, which contains c and intersects the same leaves (of the oriented 1-dimensional foliation) as the images of R_{v_i} . P_c is not strictly speaking a local section, since the Q_i are tangent along the flow. One can perturb P_c so that it becomes transverse to the flow everywhere, but since we ultimately only use the structure of P_c after projecting to the orbit space, this is not necessary for our arguments. We will point out more features of P_c in Section 5.4.

5.3 Gluing along the unstable boundary

We will follow the same strategy to glue $N(B)$ along its boundary across its complementary regions in $M(s)$, so that the foliations on $N(B)$ descend to respective foliations on $M(s)$, for suitable s .

Lemma 5.7 *The components of $M \setminus B$ are (once-punctured cusped polygons) $\times S^1$. Fix one of these components T ; ∂T is naturally decomposed into a number of annulus faces meeting along cusp circles. Φ_{red} on each annulus face consists of one or two oriented circles along with some branches attached inductively which exit through the cusp circles. Each oriented circle of $\Phi_{\text{red}} \cap \partial T$ has isotopy class equal to the ladderpole slope, in particular they are parallel. See Figure 18, left.*

Proof The first sentence is Proposition 2.9, and the second sentence follows easily. For the third sentence, note that Φ_{red} on each annulus face A of ∂T is an oriented train track with branches exiting through the boundary and with only diverging switches. This forces $\Phi_{\text{red}} \cap A$ to be a union of parallel circles with branches attached inductively which exit through the cusp circles. There cannot be more than 2




Figure 18: Collapsing across complementary regions of $M(s) \setminus N(B)$.

circles on each annulus or else the inner circles will have no branches attached. These will then contradict Φ_{red} being strongly connected, since these circles can have no outgoing edges. Meanwhile, each cusp circle meets Φ_{red} , otherwise there will be a component of $B \setminus \Phi_{\text{red}}$ carrying a circle in its branch locus, contradicting Lemma 3.9. These two observations imply the third sentence.

To show the fourth sentence, it suffices to show that the oriented circles of $\Phi_{\text{red}} \cap \partial T$ are oriented coherently. We will show this using [Landry et al. 2024, Lemma 5.4]. (Again, we caution the reader that the flow graph in [Landry et al. 2024] is oriented in the opposite direction compared to this paper.) Their lemma implies that each cycle of $\Phi \cap \partial T$ has slope equal to the ladderpole slope. Since Φ_{red} is a subgraph of Φ , the same statement holds for $\Phi_{\text{red}} \cap \partial T$. We remark that the third sentence of our lemma follows from their lemma as well. \square

Recall that we can recover the oriented 1–dimensional foliation on the boundary of each component of $M \setminus N(B)$ from the portion of Φ_{red} on the boundary of the corresponding component of $M \setminus B$, as described in Remark 5.6. From Lemma 5.7, it follows that the foliation on each annulus will consist of a band of S^1 leaves, with all other leaves spiraling out of the band and exiting through the cusp circles. Using the transverse measure du on the rectangles $\{(u, t) : |u| \leq w'_w \lambda^{-t}, t \in [0, 1]\}$, one can see that the band of S^1 leaves actually just consists of one leaf. For if the band has nonzero width, pick a local section to the foliation within the band and observe that the backward return map of the foliation has to be contracting, providing a contradiction. See Figure 18, middle.

Now fix a component T of $M \setminus N(B)$, fix an identification $T \cong (\text{punctured cusped } k\text{-gon}) \times S^1$ and take the orientation on the punctured cusped k -gon to be the one induced from that of $T \subset M$ and the ladderpole slope. Label the annulus faces of ∂T as $A^{(1)}, \dots, A^{(k)}$ cyclically according to the orientations we have chosen. Suppose a slope s is given on the torus end of T , whose geometric intersection number with the ladderpole class l is at least 1. We want to produce a “pronged l -fibering” of $T(s)$ preserving the leaves on $A^{(j)}$. By this we mean a decomposition of $T(s)$ into intervals and $|\langle s, l \rangle|$ -prongs, such that the

endpoints of the prongs lie along the unique S^1 leaves on the $A^{(j)}$, and such that the homeomorphisms on the halves of $A^{(j)}$ induced from the interval fibers preserve the leaves of the oriented 1–dimensional foliations.

Let m be the meridian of T in the description $T \cong (\text{punctured cusped } k\text{-gon}) \times S^1$, ie the isotopy class of $\partial(\text{punctured cusped } k\text{-gon}) \times \{t_0\}$. We pick $|\langle s, l \rangle|/k$ local sections $J_1^{(j)}, \dots, J_{|\langle s, l \rangle|/k}^{(j)}$ near the S^1 leaf on $A^{(j)}$, so that we have backward return maps $J_i^{(j)} \rightarrow J_{i+1}^{(j)}$ (subindices taken mod $|\langle s, l \rangle|/k$) of the oriented foliation. Let $s_i^{(j)}$ be where $J_i^{(j)}$ meets the S^1 leaf on $A^{(j)}$. First define the fibering on the cusp circles of $T(s)$, where it is uniquely determined and degenerate, then extend the fibering over subintervals of leaves that start at the interior of $J_i^{(j)}$ and end on the cusp circles, using the fact that the union of these trajectories are unions of rectangles foliated as products. Finally apply Lemma 5.4 to I_i , the subinterval of $J_i^{(j)}$ to the left of $s_i^{(j)}$, and I'_i , the subinterval of $J_{i \pm \langle s, m \rangle}^{(j+1)}$ to the right of $s_{i \pm \langle s, m \rangle}^{(j+1)}$ for each $j \pmod k$ where the \pm is taken to be the sign of $\langle s, l \rangle$, to produce the pronged I –fibering in the remaining subset of $T(s)$.

Now given a collection of slopes s on each torus end of M with $|\langle s, l \rangle| \geq 1$, collapse along the fibers of the pronged I –fibering in each corresponding component of $M(s) \setminus N(B)$. See Figure 18, right, for an illustration of a collapsed component of $M(s) \setminus N(B)$ with $|\langle s, l \rangle| = 3$. The oriented 1–dimensional foliation on $N(B)$ descends to a decomposition of $M(s)$, which is a genuine oriented 1–dimensional foliation, since all the branching of the leaves in $N(B)$ occurs along the branch locus of $N(B)$, and we have collapsed all backward trajectories ending at those points. In particular this oriented 1–dimensional foliation can be continuously parametrized into a topological flow.

For each torus end of M , call the image of the S^1 leaves on $A^{(j)}$ after collapsing the *core orbit* of the corresponding end.

5.4 Showing pseudo-Anosovity and no perfect fits

We have constructed our flow at this point. It remains to check that it satisfies the properties described in Theorem 5.1. We first recall the relevant definitions.

Definition 5.8 Consider the map $\begin{pmatrix} \lambda^{-1} & 0 \\ 0 & \lambda \end{pmatrix}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$. This preserves the foliations of \mathbb{R}^2 by horizontal and vertical lines. Let $\phi_{n,0,\lambda}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the lift of this map over $z \mapsto z^{n/2}$ that preserves the lifts of the quadrants. (When n is odd, one has to choose a branch of $z \mapsto z^{n/2}$ but it is easy to see that the result is independent of the choice.) Let $\phi_{n,k,\lambda}: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be the composition of $\phi_{n,0,\lambda}$ and rotating by $2\pi k/n$ counterclockwise. Also pull back the foliations of \mathbb{R}^2 by horizontal and vertical lines. The resulting two singular foliations are preserved by $\phi_{n,k,\lambda}$, call them l^s and l^u respectively. Let $\Phi_{n,k,\lambda}$ be the mapping torus of $\phi_{n,k,\lambda}$. Let Λ^s and Λ^u be suspensions of l^s and l^u respectively, and consider the suspension flow on $\Phi_{n,k,\lambda}$. Call the suspension of the origin the *pseudohyperbolic orbit* of $\Phi_{n,k,\lambda}$. The local behavior of the flow near the pseudohyperbolic orbit serves as the local model for singular orbits in a pseudo-Anosov flow.

Definition 5.9 A smooth pseudo-Anosov flow on a closed smooth 3–manifold N is a continuous flow ϕ^t , along with a path metric d on N , which is induced from a Riemannian metric g away from the singular orbits, satisfying:

- (S1) There is a finite collection of closed orbits $\{\gamma_1, \dots, \gamma_s\}$, called the *singular orbits* such that ϕ^t is smooth away from the singular orbits.
- (S2) Away from the singular orbits, there is a splitting of the tangent bundle into three ϕ^t –invariant line bundles $TM = E^s \oplus E^u \oplus T\phi^t$ such that

$$|d\phi^t(v)| < C\lambda^{-t}|v|$$

for every $v \in E^s, t > 0$, and

$$|d\phi^t(v)| < C\lambda^t|v|$$

for every $v \in E^u, t < 0$, for some $C, \lambda > 1$.

- (S3) Each singular orbit γ_i has a neighborhood N_i and a map f_i sending N_i to a neighborhood of the pseudohyperbolic orbit in $\Phi_{n_i, k_i, \lambda}$, for some $n_i \geq 3$, such that
 - f_i is bi-Lipschitz on N_i and smooth away from γ_i ,
 - f_i preserves the orbits, and
 - f_i sends E^s and E^u to line bundles tangent to Λ^s and Λ^u respectively.

In this case, we say that γ_i is n_i –pronged.

A topological pseudo-Anosov flow on a closed 3–manifold N is a continuous flow ϕ^t satisfying:

- (T1) There is a finite collection of closed orbits $\{\gamma_1, \dots, \gamma_s\}$, called the *singular orbits*, and two singular 2–dimensional foliations Λ^s and Λ^u , called the stable and unstable foliations respectively, which are nonsingular away from the singular orbits.
- (T2) Away from the singular orbits, every point has a neighborhood which is a *flow box*, ie a set of the form $I_s \times I_u \times I_t$ such that the flow lines are the lines with constant s and u , the stable and unstable foliations are the foliations by planes with constant u or s coordinate respectively.
- (T3) There is a *Markov partition*, ie there is a finite collection of flow boxes $\{I_s^{(i)} \times I_u^{(i)} \times I_t\}_i$ covering N with disjoint interiors, such that

$$(I_s^{(i)} \times I_u^{(i)} \times \{1\}) \cap (I_s^{(j)} \times I_u^{(j)} \times \{0\}) = \bigcup_k J_s^{(ij,k)} \times I_u^{(i)} \times \{1\} = \bigcup_k I_s^{(j)} \times J_u^{(ji,k)} \times \{0\}$$

for some finite collection of subintervals $J_s^{(ij,k)} \subset I_s^{(i)}$ and $J_u^{(ji,k)} \subset I_u^{(j)}$.

- (T4) Pick a path metric d on M . For every p and q on the same stable leaf, there exists an orientation-preserving homeomorphism $T: (-\infty, \infty) \rightarrow (-\infty, \infty)$ such that

$$\lim_{t \rightarrow \infty} d_N(\phi^t(p), \phi^{T(t)}(q)) = 0.$$

Respectively, for every p and q on the same unstable leaf, there exists an orientation-preserving homeomorphism $T: (-\infty, \infty) \rightarrow (-\infty, \infty)$ such that

$$\lim_{t \rightarrow -\infty} d_N(\phi^t(p), \phi^{T(t)}(q)) = 0.$$

(T5) Each singular orbit γ_i has a neighborhood N_i and a continuous map f_i sending N_i to a neighborhood of the pseudohyperbolic orbit in $\Phi_{n_i, k_i, \lambda}$, for some $n_i \geq 3, \lambda > 1$, such that

- f_i preserves the orbits, and
- f_i preserves Λ^s and Λ^u on the two sets.

In this case, we say that γ_i is n_i -pronged.

Definition 5.10 A (smooth/topological) pseudo-Anosov flow is *transitive* if it has a dense orbit.

It has long been a folklore fact that the notions of smooth and topological pseudo-Anosov flows are essentially equivalent in the case when the flows are transitive. The easier direction is that a transitive smooth pseudo-Anosov flow is a topological pseudo-Anosov flow. Indeed, the only nontrivial facts to check are that $E^s \oplus T\phi^t$ and $E^u \oplus T\phi^t$ are integrable (away from singular orbits), which follows from stable manifold theory (see for example [Katok and Hasselblatt 1995, Chapter 17.4]), and that there exists a Markov partition, which follows from the arguments in [Ratner 1973, Section 2]. For the other direction, we have:

Theorem 5.11 *Given a transitive topological pseudo-Anosov flow ϕ^t on a closed 3-manifold N , there exists a homeomorphism $F: N \rightarrow N$ and a smooth pseudo-Anosov flow $\hat{\phi}^t$ (with respect to some smooth structure on N), such that F maps the trajectories of ϕ^t to that of $\hat{\phi}^t$ (preserving their orientations), ie ϕ^t is C^0 -orbit equivalent to $\hat{\phi}^t$.*

This theorem has been proven recently by Shannon [2021] in the case of transitive Anosov flows, ie when there are no singular orbits. His methods in fact generalize immediately to the general case of transitive pseudo-Anosov flows. We will explain more carefully how to apply his arguments to prove Theorem 5.11 in Section 5.5. We remark that it is still open whether Theorem 5.11 is true without the hypothesis of transitivity. Also see [Mosher 1996, Section 3.1] for a related discussion.

Before we recall the definition of no perfect fits, we need a nontrivial fact: Suppose a closed 3-manifold N admits a (smooth/topological) pseudo-Anosov flow. Lift everything up to the universal cover \tilde{N} . It is shown in [Fenley and Mosher 2001, Proposition 4.1] that the orbit space \mathbb{C} of the flow on \tilde{N} is homeomorphic to \mathbb{R}^2 , and inherits two (possibly singular) 1-dimensional foliations \mathbb{C}^s and \mathbb{C}^u .

Remark 5.12 Fenley and Mosher [2001] only dealt with smooth pseudo-Anosov flows, but given Theorem 5.11, the facts stated above holds for topological pseudo-Anosov flows as well. Alternatively, the proof of [Fenley and Mosher 2001, Proposition 4.1] is valid verbatim for topological pseudo-Anosov flows.




Figure 19: A perfect fit rectangle.

We will in fact generalize the definition of no perfect fits slightly, compared to the usual definition found in, for example, [Fenley 2012].

Definition 5.13 Let ϕ be a pseudo-Anosov flow on a closed 3-manifold N , and let $\{c_1, \dots, c_k\}$ be a collection of closed orbits of ϕ . Lift these up to a flow $\tilde{\phi}$ on the universal cover \tilde{N} together with a collection of orbits $\{\tilde{c}_i\}$ which are the preimages of $\{c_i\}$.

A *perfect fit rectangle* is a rectangle-with-one-ideal-vertex properly embedded in \mathbb{O} such that 2 opposite sides of the rectangle lie along leaves of \mathbb{O}^s and the remaining 2 opposite sides lie along leaves of \mathbb{O}^u , and such that the restrictions of \mathbb{O}^s and \mathbb{O}^u to the rectangle foliate it as a product, ie conjugate to the foliations of $[0, 1]^2 \setminus \{(1, 1)\}$ by vertical and horizontal lines. See Figure 19.

The collection of orbits $\{\tilde{c}_i\}$ determines a set \mathcal{C} in \mathbb{O} . We will say that ϕ has no perfect fits relative to $\{c_1, \dots, c_k\}$ if there are no perfect fit rectangles in \mathbb{O} disjoint from \mathcal{C} .

Given a perfect fit rectangle with sides F and H along leaves of \mathbb{O}^s , sides G and K along leaves of \mathbb{O}^u , and sides F and G adjacent to the ideal vertex, notice that we can always choose a smaller perfect fit rectangle with sides F, G, H' and K' , where H' is closer to F than H , and K' is closer to G than K . We will frequently use this observation when analyzing perfect fit rectangles.

A pseudo-Anosov flow is without perfect fits according to the definition in [Fenley 2012] if and only if it is without perfect fits relative to \emptyset in Definition 5.13. This is in turn equivalent to the pseudo-Anosov flow having no perfect fits relative to the set of its singular orbits, since by definition a perfect fit rectangle cannot contain a singular point in its interior, and we can always choose a smaller perfect fit rectangle that does not contain any singular points on the boundary as well.

It is immediate from their definitions that the property of transitivity and no perfect fits relative to a collection of orbits is preserved under C^0 -orbit equivalence. Hence in view of Theorem 5.11, it suffices for us to show Theorem 5.1 with “pseudo-Anosov flow” taken to mean topological pseudo-Anosov flow.

We first verify that our flow on $M(s)$ is a topological pseudo-Anosov flow. For (T1), take the collection of singular orbits to be the core orbits of the ends of M with $|\langle s, l \rangle| \geq 3$, and take Λ^s and Λ^u to be the stable and unstable foliations on $M(s)$ respectively.

For (T2), one can construct flow box neighborhoods (away from the singular orbits) in $M(s)$ by piecing together those in $N(\Phi_{\text{red}})$.

We will call the images of the flow boxes Z_e in $M(s)$ after collapsing $N(B) \setminus N(\Phi_{\text{red}})$ and $M(s) \setminus N(B)$ by the same name. These form a Markov partition for the flow by definition, verifying (T3).

Take $d_{M(s)}$ to be the path metric on $M(s)$ induced by the Euclidean metrics $d_{Z_e}^{\text{Eucl}}$ on Z_e . This means that we define the length of a path by summing over the $d_{Z_e}^{\text{Eucl}}$ -lengths of the portions within each Z_e , and define the distance between two points to be the infimum length of paths between them. Here for a path lying on more than one Z_e , one is allowed to calculate its length by the Euclidean metric on any of them. Note that since each Z_e is compact, the identity map on Z_e is uniformly continuous with respect to $d_{Z_e}^{\text{Eucl}}$ and $d_{M(s)}|_{Z_e}$.

Let us prove (T4) in the case when p and q lie on the same stable leaf; the proof for when p and q lie on the same unstable leaf is symmetric. Notice that it suffices to prove this when p and q are close to each other, so we can assume that p and q lie in the same flow box Z_e . We parametrize the orbits through p and q so that they go through each flow box in unit time. Then we would have

$$d_{Z_e}^{\text{Eucl}}(\phi^t(p), \phi^t(q)) < C\lambda^{-t}$$

whenever $\phi^t(p)$ and $\phi^t(q)$ lie in the flow box Z_e (where C depends on p, q). So (T4) follows from uniform continuity between $d_{Z_e}^{\text{Eucl}}$ and $d_{M(s)}$.

For (T5), the neighborhoods N_i can be constructed by piecing up the flow box neighborhoods in $N(\Phi_{\text{red}})$ again. We then have to find conjugations f_i between N_i and neighborhoods of the pseudohyperbolic orbits in $\Phi_{n_i, k_i, \lambda}$ for $n_i = |\langle s_i, l_i \rangle|$. These can be constructed by choosing local sections to the flow in N_i and in $\Phi_{n_i, k_i, \lambda}$, then applying Lemma 5.4 to the prongs that are the intersections of these local sections with the singular stable and unstable leaves, taking f there to be an arbitrary homeomorphism. This constructs f_i on the singular stable and unstable leaves on the local sections, then f_i is uniquely determined on the local sections (after possibly shrinking them), by requiring that it preserves the stable and unstable foliations. Then we complete the construction of f_i by following the trajectories as they go around N_i . Observe that $|\langle s, l \rangle| \geq 2$ is required here, otherwise we will have 1-pronged singular orbits. (Meanwhile “2-pronged singular orbits” are just nonsingular orbits.)

We then show that the flow is transitive. Pick an infinite path in Φ_{red} which contains all finite paths in Φ_{red} as subpaths. Such a path can be constructed using the fact that Φ_{red} is strongly connected. Write the path as $v_1 \xrightarrow{e_1} v_2 \xrightarrow{e_2} \dots$. We claim that there is an infinite flow line in $M(s)$ passing through Z_{e_1}, Z_{e_2}, \dots in that order. To show this claim, consider the subset in Z_{e_1} of points whose forward trajectories flow through Z_{e_1}, \dots, Z_{e_k} . These subsets will be nonempty decreasing subflow boxes of Z_{e_1} ; hence taking the intersection as $k \rightarrow \infty$, we get points in Z_{e_1} whose flow lines have the required property. Pick one of these flow lines and denote it by l . Now for any point $y \in M(s)$, suppose the forward and backward trajectories of y pass through $Z_{d_0}, Z_{d_1}, Z_{d_2}, \dots$ and $Z_{d_0}, Z_{d_{-1}}, Z_{d_{-2}}, \dots$ respectively. Note that these

sequences may not be unique, due to the fact that leaves of the 1–dimensional foliation on $N(\Phi_{\text{red}})$ have forward and backward branching, but we can always just make a choice. Now fix some $N > 0$, the flow line l passes through $Z_{d_{-N}}, \dots, Z_{d_N}$ in that order at some point, by construction. Since both l and the backward flow line through y pass through $Z_{d_{-N}}, \dots, Z_{d_0}$, they intersect Z_{d_0} in flow lines not more than $w_{v_{-N}}\lambda^{-N}$ apart in the $d_{Z_{d_0}}^{\text{Eucl}}$ metric. Similarly, looking in the forward direction, l and the flow line through y are not more than $w'_{w_N}\lambda^{-N}$ apart in the $d_{Z_{d_0}}^{\text{Eucl}}$ metric. Hence we conclude that l passes through a $(\max w_v + \max w'_v)\lambda^{-N} d_{Z_{d_0}}^{\text{Eucl}}$ -neighborhood of y . Letting $N \rightarrow \infty$ and using uniform continuity between $d_{Z_{d_0}}^{\text{Eucl}}$ and $d_{M(s)}$, this proves transitivity.

Finally we prove that the flow has no perfect fits relative to the core orbits, which we denote by c_i . To do this we will make use of the surfaces P_c constructed in Section 5.3. There, we defined these as subsets of $N(B)$, but here we will consider their images in $M(s)$ after collapsing $M \setminus N(B)$ and still refer to them by the same name. Consider the lifts of P_c to the universal cover $\widetilde{M}(s)$. We denote these by $\widetilde{P}_{\tilde{c}}$, where \tilde{c} ranges over components of the branch locus of \widetilde{B} that are lifts of c , which we will conflate with the corresponding components of the branch locus of $\widetilde{N}(B)$. After collapsing $\widetilde{M}(s) \setminus \widetilde{N}(B)$, each $\tilde{c} \subset \widetilde{M}(s)$ lies on a leaf of $\widetilde{\Lambda}^u$ and is transverse to the flow, meeting exactly those flow lines which spiral out of the lift of the core orbit produced by collapsing the component of $\widetilde{M}(s) \setminus \widetilde{N}(B)$ which \tilde{c} opens up towards. Moreover, as one goes along \tilde{c} in its orientation lifted from c , we meet flow lines that spiral closer to the lift of the core orbit. Hence \tilde{c} projects to an open half leaf of \mathbb{O}^u on \mathbb{O} , which under the orientation of \tilde{c} , limits to the lift of the core orbit (as a point in \mathbb{O}) in the forwards direction and escapes to infinity in \mathbb{O} in the backwards direction. Since the projection is a homeomorphism, we will call this open half leaf \tilde{c} as well for convenience.

Meanwhile, each $\widetilde{P}_{\tilde{c}}$ projected to \mathbb{O} will be a surface containing \tilde{c} . More precisely, recall that P_c is the union of the images of R_{v_i} and Q_i . Let $\widetilde{R}_{v_i} \subset \widetilde{\Phi}_{\text{red}}$ be the lifts of R_{v_i} that meet \tilde{c} and $\widetilde{Q}_i \subset \widetilde{N}(B)$ the lifts of Q_i that meet \tilde{c} . Then $\widetilde{P}_{\tilde{c}}$ is the union of the images of \widetilde{R}_{v_i} and \widetilde{Q}_i after collapsing $\widetilde{N}(B) \setminus \widetilde{N}(\Phi_{\text{red}})$ and $\widetilde{M}(s) \setminus \widetilde{N}(B)$. The image of each \widetilde{R}_{v_i} will be projected homeomorphically to a rectangle in \mathbb{O} along \tilde{c} since these are transverse to the flow, while the image of each \widetilde{Q}_i is tangent to the flow, hence will be projected into sides of the rectangles.

The way these rectangles fit together is described by the following lemma.

Lemma 5.14 *The projection of each $\widetilde{P}_{\tilde{c}}$ is homeomorphic to*

$$\{(x, y) \in (-\infty, 0) \times \mathbb{R} : r_1(x) \leq y \leq r_2(x)\}$$

where r_1 (resp. r_2) is a nonincreasing (resp. nondecreasing) negative (resp. positive) piecewise constant function, \tilde{c} is sent to $(-\infty, 0) \times \{0\}$ in the direction of increasing x , and the restriction of Λ^s (resp. Λ^u) are sent to the foliation by vertical (resp. horizontal) lines. See Figure 17.

Proof The rectangles in \mathbb{O} that are the projections of \widetilde{R}_{v_i} have nonoverlapping interiors and cover \tilde{c} , so the only thing that remains to be shown is the monotonicity of r_1 and r_2 .

To show this, recall that $Q_i \cap \partial^+ R_{v_i} = \partial^+ R_{v_i}$ and $Q_i \cap \partial^- R_{v_{i+1}} \subset \partial^- R_{v_{i+1}}$. This implies that the line that is the image of \tilde{Q}_i is the right vertical side of the rectangle that is the projection of \tilde{R}_{v_i} and a subset of the left vertical side of the rectangle that is the projection of $\tilde{R}_{v_{i+1}}$. In simpler terms, the rectangles get taller or are of equal height as one goes along \tilde{c} , which establishes the lemma. \square

For convenience, from now on let us write $\tilde{P}_{\tilde{c}}$ and \tilde{R}_{v_i} for their images in \mathbb{O} as well.

To establish no perfect fits, we will basically argue that if there is a perfect fit rectangle in \mathbb{O} disjoint from the images of \tilde{c}_i , then we can reduce it to sit inside one of these $\tilde{P}_{\tilde{c}}$ with the ideal vertex at a corner of the rectangle with larger x coordinate. But from the monotonicity of r_1 and r_2 , such a perfect fit rectangle cannot be properly embedded in $\tilde{P}_{\tilde{c}}$, giving us a contradiction.

So suppose we have a perfect fit rectangle in \mathbb{O} disjoint from the images of \tilde{c}_i with sides F and H along leaves of \mathbb{O}^s , sides G and K along leaves of \mathbb{O}^u , and sides F and G adjacent to the ideal vertex. Lift H to $\widetilde{M}(s)$. By flowing forward (and possibly rechoosing K to be closer to G) we can assume that the lift of H is short enough to be contained within the lift of some flow box, which we denote by \tilde{Z}_e . In particular we can take the lift of H to lie on the image of some \tilde{R}_v . If the lift of H does not meet the image of the branch locus of $\widetilde{N}(B)$ in its interior, ie it is fully contained in the bottom face of some \tilde{Z}_e , then we can rechoose the lift of H to be on the top face of \tilde{Z}_e . When we do so, the length of the lift (with respect to the Euclidean metric on the \tilde{Z}_e 's) is increased by a factor of λ , so we cannot continue this operation indefinitely. In other words, we can choose a lift of H which lies on the image of some \tilde{R}_v and whose interior meets the image of the branch locus of $\widetilde{N}(B)$.

Recall that the components of the branch locus of $\widetilde{N}(B)$ carry orientations induced from those of B . We split into two cases depending on how the lift of H meets the image of the branch locus of $\widetilde{N}(B)$ with respect to this orientation.

Case 1 is if the lift of H meets the image of a component \tilde{c} of the branch locus oriented in the same direction as K leaving H . In this case, by moving K closer to G , we can suppose that K lies within \tilde{c} on \mathbb{O} . This uses the assumption that the perfect fit rectangle is disjoint from the images of \tilde{c}_i , otherwise K could go beyond \tilde{c} on the leaf of \mathbb{O}^s . Now by construction, $\tilde{P}_{\tilde{c}}$ in \mathbb{O} contains both H and K . In the description of Lemma 5.14, the direction at which K leaves H is the direction of increasing x , so the perfect fit rectangle must sit inside $\tilde{P}_{\tilde{c}}$ and with the ideal vertex at a corner with larger x coordinate. See Figure 20. As pointed out, this gives us a contradiction.

Case 2 is if the lift of H meets a component \tilde{c} of the branch locus oriented in the different direction as K leaving H . As above, we can assume that K lies within \tilde{c} on \mathbb{O} . Recall that $\tilde{P}_{\tilde{c}}$ consists of a union of rectangles \tilde{R}_{v_i} along \tilde{c} , and H is contained in one of these rectangles. Since K is a finite interval, it only goes through finitely many \tilde{R}_{v_i} , call this finite number the *length* of K . Notice that it is sometimes possible to reduce the length of K by moving H closer to F and from $\tilde{R}_{v_{i+1}}$ to \tilde{R}_{v_i} while keeping it inside $\tilde{P}_{\tilde{c}}$. In any case, we can assume that the length of K is the minimum possible value among all perfect fit rectangles that are disjoint from the images of \tilde{c}_i in \mathbb{O} .




Figure 20: In case 1, the ideal vertex of the perfect fit rectangle must lie inside $\tilde{P}_{\tilde{c}}$, giving us a contradiction.

If the length of K is 1, then we have a perfect fit rectangle with H and K within a single rectangle \tilde{R}_v . But then the perfect fit rectangle would not be properly embedded, so this is a contradiction.

Hence the remaining case is if we cannot move H from $\tilde{R}_{v_{i+1}}$ to \tilde{R}_{v_i} for some consecutive rectangles \tilde{R}_{v_i} and $\tilde{R}_{v_{i+1}}$ along \tilde{c} . Let $\partial^+ \tilde{R}_{v_i}$ be the side of \tilde{R}_{v_i} which \tilde{c} exits, and $\partial^- \tilde{R}_{v_{i+1}}$ be the side of $\tilde{R}_{v_{i+1}}$ which \tilde{c} enters. (This is consistent with the notation used in Section 5.2.) We can at least move H onto the image of $\partial^- \tilde{R}_{v_{i+1}}$ in \mathbb{C} in this case. Then we can pick a lift of H which lies in the image of the stable face of some component \tilde{K} of $\widetilde{N(B)} \setminus \widetilde{N(\Phi_{\text{red}})}$ which contains $\partial^- \tilde{R}_{v_{i+1}}$. The assumption that we cannot move H to \tilde{R}_{v_i} means that an endpoint of $\partial^+ \tilde{R}_{v_i}$ lies on a component of $\partial^\mu \tilde{K}$ which meets the other stable face at a point whose image in $\widetilde{M(s)}$ has forward trajectory meeting this lift of H in its interior. See Figure 21 for an illustration of the situation.

We claim that the crisscross property stated in Lemma 3.9 implies that at least one of the endpoints of $\partial^+ \tilde{R}_{v_i}$ lies on a component of $\partial^\mu \tilde{K}$ whose branch locus L is oriented in the different direction as \tilde{c} and meets the stable face of \tilde{K} containing $\partial^- \tilde{R}_{v_{i+1}}$ at a point z on or above $\partial^- \tilde{R}_{v_{i+1}}$. To check this, we let J be the component of $B \setminus \Phi_{\text{red}}$ corresponding to K , let \tilde{J} be the universal cover of J , and suppose that \tilde{c} lies along $y = -x$ in the branch locus of \tilde{J} under the model described in Lemma 3.9.

If $w \geq 2$, then the endpoint of $\partial^+ \tilde{R}_{v_i}$ away from the side on which the tongue is attached along $y = -x$ to \tilde{J} lies on the component of $\partial^\mu \tilde{K}$ whose branch locus corresponds to $y = x - 2/w$, hence is oriented in the different direction as \tilde{c} and meets the stable face of \tilde{K} containing $\partial^- \tilde{R}_{v_{i+1}}$ at a point strictly above $\partial^- \tilde{R}_{v_{i+1}}$. This is the case that is illustrated in Figure 21.

If $w = 1$ and if the tongue attached along $y = x - 2$ is done so in the opposite side as that along $y = -x$, we have a similar conclusion: the endpoint of $\partial^+ \tilde{R}_{v_i}$ away from the side on which the tongue is attached along $y = -x$ lies on the component of $\partial^\mu \tilde{K}$ whose branch locus corresponds to $y = x - 2$, hence is oriented in the different direction as \tilde{c} , but now meeting the stable face of \tilde{K} containing $\partial^- \tilde{R}_{v_{i+1}}$ at a point on $\partial^- \tilde{R}_{v_{i+1}}$.

Finally, if $w = 1$ and if the tongue attached along $y = x - 2$ is done so in the same side as that along $y = -x$, then the endpoint of $\partial^+ \tilde{R}_{v_i}$ on this common side lies on the component of $\partial^\mu \tilde{K}$ whose branch

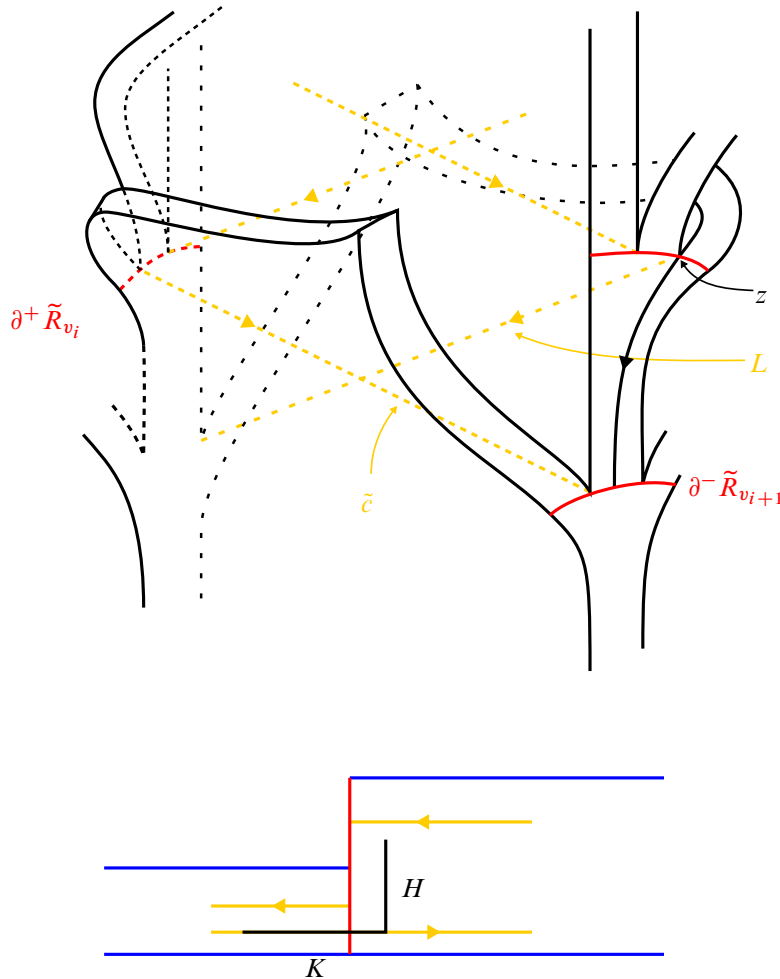


Figure 21: In case 2, up to possibly rechoosing a lift of H that is higher up, we can assume that the lift of H meets a component of the branch locus of $\widehat{N}(B)$ oriented in the same direction as K leaving H .

locus corresponds to $y = x - 2$, hence is oriented in the different direction as \tilde{c} and meets the stable face of \tilde{K} containing $\partial^- \tilde{R}_{v_{i+1}}$ at a point on $\partial^- \tilde{R}_{v_{i+1}}$.

Now, if, in the claim, the point z lies on $\partial^- \tilde{R}_{v_{i+1}}$ then we are in case 1 since the component of the branch locus on which L in the claim lies on will meet the lift of H and is oriented in the same direction as K leaving H .

If z is strictly above $\partial^- \tilde{R}_{v_{i+1}}$, then up to moving K closer to G , we can rechoose a lift of H on the image of a higher \tilde{R}_v so that it meets a component of the branch locus of $\widehat{N}(B)$ oriented in the different direction as \tilde{c} , hence the same direction as K leaving H . See Figure 21. Then we have reduced to case 1 again. Hence in any case, we see that the perfect fit rectangle cannot exist.

5.5 Proof of Theorem 5.11

By [Brunella 1995] and the standard fact that pseudo-Anosov flows are expansive (which can in turn be proved via a symbolic dynamics argument making use of a Markov partition), ϕ^t admits a Birkhoff section. Recall that this means there is an embedding of an oriented surface with boundary $\Sigma \hookrightarrow N$, such that $\partial\Sigma$ is a union of orbits, $\text{int } \Sigma$ is positively transverse to N , and every point intersects Σ in finite forward time. We briefly sketch an argument here.

For every point $x \in N$, take a local section to the flow near x . The local section is divided into $2n$ regions by the stable and unstable leaves containing x (where $n = 2$ unless x lies on a singular orbit). Name them $A_1, B_1, \dots, A_n, B_n$ in a cyclic order. It is standard that the set of periodic orbits in a transitive pseudo-Anosov flow is dense. For example, this can again be proved using a symbolic dynamics argument, similar to how we proved transitivity before. Use this fact to find nonsingular periodic orbits p_i passing through A_i . We can further assume that the stable and unstable leaves which p_i lie on are orientable, since the set of periodic orbits satisfying this is dense as well. Encode the closed orbits p_i by periodic symbolic sequences using a Markov partition, then concatenate these sequences to get a long periodic sequence which corresponds to a closed orbit q in N “shadowing” p_1, \dots, p_n , in that cyclic order. Now construct a $2n$ -gon on the local section with vertices at where $p_1, q, p_2, q, \dots, p_n, q$ meet the local section. Here the multiple listings of q refer to the multiple times q meets the local section as it shadows the p_i 's. Connect up the edges $[q, p_k]$ and $[p_k, q]$ on the local section for each k by following the flow. After perturbation, we get an immersed local Birkhoff section. See Figure 22. Since N is compact, we can take the union of finitely many such surfaces and have all points meet the union in finite forward time. Then we can perform some surgeries along the self-intersections to get a genuine Birkhoff section Σ .

Our argument here is a generalization of an argument of Fried [1983] to pseudo-Anosov flows. We remark that there is an alternative argument to construct q in [Brunella 1995] without using symbolic dynamics. We can further assume that p_i are orientation preserving paths on the stable and unstable leaves they lie on, since the set of orientation preserving periodic orbits is dense as well.

Let Σ° be Σ without its boundary components, and let $\bar{\Sigma}$ be Σ with all its boundary components collapsed to points. The first return map on Σ restricts to a pseudo-Anosov homeomorphism h° on Σ° , with the intersections of Λ^s and Λ^u with Σ° acting as the stable and unstable foliations. This in turn induces a pseudo-Anosov homeomorphism \bar{h} on $\bar{\Sigma}$. (Note that a map obtained this way on a general Birkhoff section may have 1-pronged singularities, but this is eliminated by requiring that the components of $\partial\Sigma$ lie on orientable stable and unstable leaves.) By classical results, there is a smooth structure on $\bar{\Sigma}$ such that \bar{h} is smooth away from the singular points. Hence the mapping torus of $(\bar{\Sigma}, \bar{h})$ carries a smooth structure such that the suspension flow is smooth away from the singular points. Moreover, for each closed orbit of the suspension flow, there is a neighborhood diffeomorphic to a neighborhood of the pseudohyperbolic orbit of some $\Phi_{n,k,\lambda}$ in Definition 5.8. In particular this is true for the closed orbits which are suspensions of points of $\bar{\Sigma} \setminus \Sigma^\circ$. We call these closed orbits the surgery orbits for ease of notation.

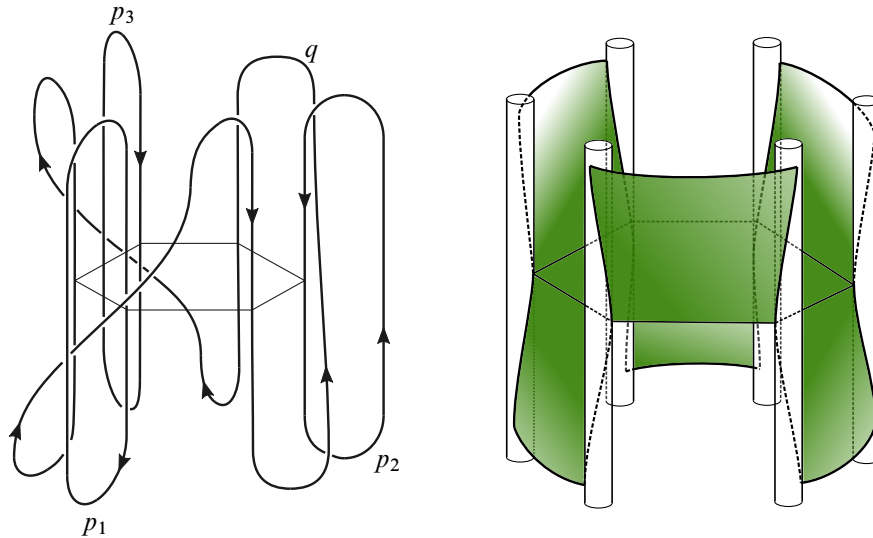


Figure 22: Left: pick a long periodic orbit q which weaves around the shorter periodic orbits p_1, \dots, p_n . Right: take a $2n$ -gon transverse to the flow with vertices at p_1, q, \dots, p_n, q and connect up its sides to produce an immersed local Birkhoff section.

Meanwhile, note that $N \setminus \partial\Sigma$ with the restricted ϕ^t flow is C^0 -orbit equivalent to the suspension flow on the mapping torus (Σ°, h°) . In particular, there is a vector field on $N \setminus \partial\Sigma$ (with respect to some smooth structure) whose trajectories are equal to the flow lines of ϕ^t . Using the orbit equivalence, we can transfer neighborhoods of the surgery orbits, minus the surgery orbits themselves, to neighborhoods of the ends of $N \setminus \partial\Sigma$, then fill back in the components of $\partial\Sigma$ to get neighborhoods of $\partial\Sigma$ in N . It remains to perform Dehn surgery on these neighborhoods, which we will do by cutting out small special shaped subneighborhoods and gluing similarly shaped regions back in, and arguing that the result is orbit equivalent to ϕ^t on the original manifold N . The first step is a restatement of the arguments in [Goodman 1983], while the second step relies on [Shannon 2021, Theorem 3.9].

Within each such neighborhood in N , locate and excise a small *cross-shaped neighborhood* of $\partial\Sigma$ using the orbit equivalence map between $N \setminus \partial\Sigma$ and the mapping torus of (Σ°, h°) . Here we provide an illustration of these cross-shaped neighborhoods in Figure 23 and refer to [Shannon 2021, Section 5.2] for their precise definition. The key feature of these neighborhoods is that their boundaries have alternating annulus faces, the flow transverse to half of the faces and tangent to the remaining faces.

Then we glue some cross-shaped neighborhoods back in by matching up the vector fields at the boundary, so that the glued vector field determines a smooth flow (away from the singular orbits). See [Shannon 2021, Section 5.3] for an explicit description of this gluing map. The key observation that makes this work is that there is a freedom in this gluing: one can do Dehn twists suitably along annulus faces which the flow is transverse to, and still have the vector fields match up. So far there is no difference between the Anosov case and the general pseudo-Anosov case.




Figure 23: A cross-shaped neighborhood and the portion of the Birkhoff section constructed in it.

Then one uses a cone field argument to find the stable and unstable line bundles, E^s and E^u , and show that the glued up flow is pseudo-Anosov, provided that the neighborhood we excised is small enough. This is done in the Anosov case in [Shannon 2021, Section 5.4], and is in fact more simple in the pseudo-Anosov case, since one only needs to construct E^s and E^u away from the singular orbits. On the singular orbits, the definition only requires an orbit equivalence between a neighborhood in N and a neighborhood in $\Phi_{n,k,\lambda}$, and these are provided by the orbit equivalence between $N \setminus \partial\Sigma$ and the mapping torus of (Σ°, h°) , since by construction the orbits of $\partial\Sigma$ are nonsingular. As remarked above, this method of constructing (pseudo-)Anosov flows on Dehn surgeries appeared back in [Goodman 1983].

Finally, we construct a Birkhoff section on this glued flow by taking the portion of the original Birkhoff section Σ outside of the excised cross-shaped neighborhood, and extending it inside the glued cross-shaped neighborhood by taking a union of straight lines (with respect to the Euclidean metric on a canonical set of coordinates) towards the pseudohyperbolic orbit. Call this surface Σ' . We illustrate the portion of Σ' within each cross-shaped neighborhood in Figure 23 and refer to [Shannon 2021, Section 5.5] for precise formulas, where it is also shown that this indeed defines a Birkhoff section, and that the return map h' is pseudo-Anosov. Moreover, h and h' induce the same automorphism on $\pi_1(\Sigma) \cong \pi_1(\Sigma')$, preserving the peripheral subgroups. Hence using [Shannon 2021, Theorem 3.9], which applies to pseudo-Anosov flows as well, as stated there, we get a C^0 -orbit equivalence between ϕ^t on N and the smooth pseudo-Anosov glued flow.

5.6 Properties of the construction

We point out two properties of the pseudo-Anosov flows we constructed. We will phrase our proofs in terms of the initial topological pseudo-Anosov flow we constructed, but thanks to Theorem 5.11, one can just take an orbit equivalence map to prove the properties for the smooth pseudo-Anosov flow as well.

Proposition 5.15 *The unstable lamination of the pseudo-Anosov flow, obtained by blowing air into the leaves of the unstable foliation that contain the core orbits, is carried by $B \subset M(s)$.*

Proof The unstable lamination can also be obtained by taking the unstable foliation on $N(B)$ and blowing air into the nonmanifold leaves. There is a projection $N(B) \rightarrow B$ by definition, and this projects the lamination to B . \square

Definition 5.16 Suppose a pseudo-Anosov flow ϕ has a Markov partition as in Definition 5.9. Define a directed graph G by letting the set of vertices be the flow boxes, and putting an edge from $(I_s^{(j)} \times I_u^{(j)} \times [0, 1]_t)$ to $(I_s^{(i)} \times I_u^{(i)} \times [0, 1]_t)$ for every $J_s^{(ij,k)}$.

Notice that G has a natural embedding in N by placing the vertices in the interior of the corresponding flow box and placing the edges through the corresponding intersections $J_s^{(ij,k)} \times I_u^{(i)} \times \{1\}$. G together with this embedding in N is said to *encode the Markov partition*.

Proposition 5.17 *The pseudo-Anosov flow constructed in Theorem 5.1 admits a Markov partition encoded by the reduced flow graph Φ_{red} .*

Proof The flow boxes Z_e form a Markov partition by construction, but notice that the graph that encodes this Markov partition is not Φ_{red} . Indeed, the flow boxes Z_e correspond to edges, not vertices, of Φ_{red} ; hence the graph encoding this Markov partition is instead a “dual” of Φ_{red} .

What we will do is extract from this an alternate Markov partition that is encoded by Φ_{red} . For each vertex v of Φ_{red} , recall the rectangle R_v considered in Section 5.2. Let Z'_v be the closure of the union of trajectories in $M(s)$ that start in the image of the interior of R_v and end when it meets the image of the interior of another (possibly the same) $R_{v'}$. These Z'_v cover $M(s)$ with disjoint interiors, but they might not be flow boxes. Instead, they are in general homeomorphic to $\{(s, u, t) : s \in I_s, u \in I_u, t \in [f(s, u), 1]\}$ with a function f with discontinuities along finitely many lines $u = u_0$. Each of these intervals of discontinuity correspond to an adjacent pair of flow boxes leaving R_v .

Consider the total number of intervals of discontinuity at the bottom of all the Z'_v . If this number is 0, then we have a genuine flow box decomposition which gives a Markov partition encoded by Φ_{red} .

If this number is positive, we claim that we can modify the Z'_v to reduce it. Fix an interval of discontinuity J of Z'_v . Suppose the two sides to it on the bottom of the Z'_v lie on the image of R_{v_1} and R_{v_2} . We can isotope R_{v_1} near J , moving each point along a trajectory of the flow, so that R_{v_1} matches up with R_{v_2} along J , eliminating the discontinuity. See Figure 24.

One might worry that this produces new intervals of discontinuity, but we argue that this will not happen. Using the same notation as above, since we only modify R_{v_1} near J , the operation only possibly creates new intervals of discontinuity that lie on the same sides of R_{v_1} and R_{v_2} as J and share an endpoint with J . Suppose K is such an interval. We illustrate one example of such a K in Figure 24. $J \cap K$ is a point x that lies on the image of a component c of the branch locus of $N(B)$ that passes through R_{v_1}




Figure 24: Modifying R_v to eliminate the intervals of discontinuity, so that we obtain a decomposition by flow boxes.

and R_{v_2} . In particular, R_{v_1} and R_{v_2} match up along x already. (In Figure 24, we illustrate c as the yellow line; compare with Figure 17, top.) So we can in fact isotope R_{v_1} while fixing K , ensuring that we will not create a new interval of discontinuity at K . Doing this inductively, we can reduce the number of intervals of discontinuity to 0. \square

Corollary 5.18 *Given a cycle l of Φ in $M(s)$, there is a periodic orbit c of the flow such that l is homotopic to c . Conversely, given a periodic orbit c of the flow, there exists a cycle l of Φ such that l is homotopic to some multiple c^k for $k \geq 1$.*

Proof This is mostly a straightforward consequence of Proposition 5.17 and Lemma 3.8.

Given a cycle l of Φ , we first assume that there is a homotopic cycle of Φ_{red} , which we also write as l . Then we can proceed as in the proof of transitivity to construct orbits of the flow passing through the sequence of flow boxes corresponding to the sequence of vertices of Φ_{red} passed through by l' . A similar argument as in that proof shows that such an orbit will be unique; hence the orbit c will be periodic, since the sequence of flow boxes is periodic. A homotopy between l and c can now be constructed using a straight line homotopy within each flow box.

We caution that c may not be a primitive closed orbit. Indeed, the above construction in fact gives a S^1 leaf c' of the singular 1-dimensional foliation on $N(\Phi_{\text{red}})$. But when we quotient $N(\Phi_{\text{red}})$ to $M(s)$, c' may cover an orbit multiple times.

If l is not homotopic to a cycle of Φ_{red} , then by the proof of Lemma 3.8, l must be an infinitesimal cycle within a twisted wall. Let l' be the boundary cycle of that wall. Applying the argument above for this l' , we find a circular leaf c' in $N(\Phi_{\text{red}})$ lying on the stable boundary $\partial_s N(\Phi_{\text{red}})$ which is homotopic to l'

in $M(s)$. When we quotient $N(\Phi_{\text{red}})$ to $N(B)$, c' double covers its image c'' , which is hence homotopic to l . The image of c'' after quotienting to $M(s)$ will be the desired periodic orbit c in this case. Again, we remark that c may not be primitive.

Conversely, given a periodic orbit c of the flow on $M(s)$, consider its preimage in $N(\Phi_{\text{red}})$. This will be a disjoint union of circular leaves which cover c possibly multiple times. Pick one of these circles c' . The periodic sequence of flow boxes which it meets will determine a cycle l of $\Phi_{\text{red}} \subset \Phi$. A homotopy between l and c can now be constructed using a straight line homotopy within each box as before. \square

Combining [Landry 2022, Theorem A] and [Landry et al. 2024, Proposition 5.7], it is known that when $|\langle s, l \rangle| \geq 3$, the homology classes of the cycles of Φ in $M(s)$ span a cone dual to a face of the Thurston unit ball, namely the face determined by the negative of the Euler class of the veering triangulation

$$-e(\Delta) = \sum \frac{1}{2}(\#\text{prongs}(\gamma) - 2) \cdot [\gamma] \in H_1(M(s)).$$

Hence Corollary 5.18 implies that the homology classes of the periodic orbits of the pseudo-Anosov flow in Theorem 5.1 span the same dual cone.

Remark 5.19 Corollary 5.18 and its proof also allows one to count the number of periodic orbits of the pseudo-Anosov flow using Φ . The precise statements one can get are essentially the same as those shown in [Landry et al. 2023, Theorem 6.1], but again we remind the reader that our setting in this paper is opposite to that in [Landry et al. 2023].

We elaborate on this a little more. The proof of Corollary 5.18 shows that for each cycle l of Φ , there is a unique primitive periodic orbit c of the flow homotopic to l , unless if l determines a circular leaf that lies on the boundary of $N(\Phi_{\text{red}})$ and multiple covers a periodic orbit of the flow under the quotient maps. If such a circular leaf lies on $\partial^s N(\Phi_{\text{red}}) \setminus \partial^u N(\Phi_{\text{red}})$, then l lies on the boundary of a component of $B \setminus \Phi_{\text{red}}$, and if the circular leaf multiple covers a periodic orbit, the component must be a Möbius band with tongues. If such a circular leaf lies on $\partial^u N(\Phi_{\text{red}})$, then l can be homotoped into a torus end of M .

Conversely, the proofs of Corollary 5.18 and Lemma 3.8 show that for each primitive periodic orbit c of the flow, there is a unique cycle l of Φ homotopic to c , unless the preimage of c in $N(\Phi_{\text{red}})$ multiple covers c or if l is homotopic to a boundary cycle of a wall. The former can only happen when the preimage lies in the boundary of $N(\Phi_{\text{red}})$. If the preimage lies in $\partial^s N(\Phi_{\text{red}}) \setminus \partial^u N(\Phi_{\text{red}})$, then c is homotopic to the core of a component of $B \setminus \Phi_{\text{red}}$ which is a Möbius band with tongues as above. If the preimage lies in $\partial^u N(\Phi_{\text{red}})$ then c is a core orbit. In the latter case, l is homotopic to an AB cycle, by the discussion in Remark 3.6.

6 Discussion and further questions

We conclude with some questions coming out of this paper.

In Section 3, we analyzed *how* the flow graph of a veering triangulation can fail to be strongly connected. An equally interesting question is *when* does the flow graph fail to be strongly connected. For example,

for layered veering triangulations, one could ask if there is a topological criterion on the monodromy that determines whether the flow graph is strongly connected. Another interesting question is if having a strongly connected flow graph is a generic property among all veering triangulations.

We already pointed out in Section 4 that in our argument for Theorem 4.1, there is a lot of room for improvement for the bounds we used. A more interesting approach to try to obtain better bounds overall is to study the flow graphs themselves. It is shown in [McMullen 2015] how to recover the dilatation of the pseudo-Anosov monodromy from the Teichmüller polynomial of the associated fibered face. The Teichmüller polynomial is in turn related to the veering polynomial of the layered veering triangulation, defined in [Landry et al. 2024] as the inclusion of the Perron polynomial of the flow graph in the 3–manifold. Now, the combinatorial rigidity of veering triangulations might impose restrictions on the graph isomorphism type of flow graphs, which might in turn have consequences for the behavior of their Perron polynomials, possibly giving some bounds on dilatations. We remark that Parlak [2024] computed some data for veering polynomials of veering triangulations in the census. This data might be helpful for extrapolating patterns for the approach above.

One motivation for improving the bound is to use Theorem 4.1 to solve the minimal dilatation problem by computation. Starting with an upper bound, we can use Theorem 4.1 to reduce the search to a finite list. However, it has been conjectured that one cannot do better than $P = \left(\frac{1}{2}(3 + \sqrt{5})\right)^2$ (see [Hironaka 2010], where examples that asymptotically attain this bound are also shown). Plugging this into our Theorem 4.1, we have to look at all 3–manifolds triangulated by at most 9.7×10^{15} veering tetrahedra. In comparison, the veering triangulation census only covers triangulations up to 16 veering tetrahedra for now, so huge advancements would have to be made before this approach is computationally feasible.

Finally, there are plenty of natural questions one could ask about our construction of the pseudo-Anosov flow in Section 5.

We have been using the unstable branched surface B throughout this paper, but there is also a stable branched surface associated to a veering triangulation, which can be defined as the unstable branched surface associated to the veering triangulation with reversed coorientation on the faces. It is natural to ask whether the stable lamination of our pseudo-Anosov flow is carried by the stable branched surface.

Also, one can repeat our entire construction with the stable branched surface replacing the unstable branched surface. Then one can ask whether the pseudo-Anosov flows obtained using the stable and unstable branched surfaces are C^0 –orbit equivalent. A related development is that Parlak [2024] has found veering triangulations whose veering polynomial is different from that of the veering triangulation with reversed coorientation on the faces.

We mentioned in the introduction that the construction of a pseudo-Anosov flow without perfect fits from a veering triangulation was first done by Schleimer and Segerman. It is an interesting question to ask whether the flow constructed by our approach ends up being equivalent to the one constructed by Schleimer and Segerman.

One could also ask if our construction is inverse, in either direction, to the construction by Agol and Guéritaud, which produces a veering triangulation out of a pseudo-Anosov flow without perfect fits. More precisely, if one starts with a pseudo-Anosov flow without perfect fits relative to orbits $\{c_i\}$ on a closed orientable 3-manifold N and constructs a veering triangulation on N with $\{c_i\}$ drilled out, then constructs a pseudo-Anosov flow by filling in slopes that recover N , is the final pseudo-Anosov flow C^0 -orbit equivalent to the initial pseudo-Anosov flow? Conversely, if one starts with a veering triangulation and constructs a pseudo-Anosov flow by filling along slopes s with $|\langle s, l \rangle| \geq 2$, then constructs a veering triangulation with the core orbits drilled out, is the final veering triangulation isomorphic to the initial veering triangulation? If one could show that our construction is equivalent to that of Schleimer and Segerman, then the answer to both questions would be “yes”, since as mentioned in the introduction, this is a feature of their construction.

References

- [Agol 2011] **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 MR Zbl
- [Birman and Williams 1983a] **J S Birman, R F Williams**, *Knotted periodic orbits in dynamical systems, I: Lorenz’s equations*, Topology 22 (1983) 47–82 MR Zbl
- [Birman and Williams 1983b] **J S Birman, R F Williams**, *Knotted periodic orbits in dynamical system, II: Knot holders for fibered knots*, from “Low-dimensional topology” (S J Lomonaco, Jr, editor), Contemp. Math. 20, Amer. Math. Soc., Providence, RI (1983) 1–60 MR Zbl
- [Brunella 1995] **M Brunella**, *Surfaces of section for expansive flows on three-manifolds*, J. Math. Soc. Japan 47 (1995) 491–501 MR Zbl
- [Farb et al. 2011] **B Farb, C J Leininger, D Margalit**, *Small dilatation pseudo-Anosov homeomorphisms and 3-manifolds*, Adv. Math. 228 (2011) 1466–1502 MR Zbl
- [Fenley 2012] **S Fenley**, *Ideal boundaries of pseudo-Anosov flows and uniform convergence groups with connections and applications to large scale geometry*, Geom. Topol. 16 (2012) 1–110 MR Zbl
- [Fenley and Mosher 2001] **S Fenley, L Mosher**, *Quasigeodesic flows in hyperbolic 3-manifolds*, Topology 40 (2001) 503–537 MR Zbl
- [Frankel et al. 2019] **S Frankel, S Schleimer, H Segerman**, *From veering triangulations to link spaces and back again*, preprint (2019) arXiv 1911.00006
- [Fried 1983] **D Fried**, *Transitive Anosov flows and pseudo-Anosov maps*, Topology 22 (1983) 299–303 MR Zbl
- [Futer and Guéritaud 2013] **D Futer, F Guéritaud**, *Explicit angle structures for veering triangulations*, Algebr. Geom. Topol. 13 (2013) 205–235 MR Zbl
- [Giannopolous et al. 2019] **A Giannopolous, S Schleimer, H Segerman**, *A census of veering triangulations*, electronic reference (2019) Available at <https://math.okstate.edu/people/segerman/veering.html>
- [Goodman 1983] **S Goodman**, *Dehn surgery on Anosov flows*, from “Geometric dynamics” (J Palis, Jr, editor), Lecture Notes in Math. 1007, Springer (1983) 300–307 MR Zbl

- [Guéritaud 2016] **F Guéritaud**, *Veering triangulations and Cannon–Thurston maps*, *J. Topol.* 9 (2016) 957–983 MR Zbl
- [Ham and Song 2007] **J-Y Ham, W T Song**, *The minimum dilatation of pseudo-Anosov 5–braids*, *Exp. Math.* 16 (2007) 167–179 MR Zbl
- [Hironaka 2010] **E Hironaka**, *Small dilatation mapping classes coming from the simplest hyperbolic braid*, *Algebr. Geom. Topol.* 10 (2010) 2041–2060 MR Zbl
- [Hodgson et al. 2011] **C D Hodgson, J H Rubinstein, H Segerman, S Tillmann**, *Veering triangulations admit strict angle structures*, *Geom. Topol.* 15 (2011) 2073–2089 MR Zbl
- [Hodgson et al. 2016] **C D Hodgson, A Issa, H Segerman**, *Non-geometric veering triangulations*, *Exp. Math.* 25 (2016) 17–45 MR Zbl
- [Katok and Hasselblatt 1995] **A Katok, B Hasselblatt**, *Introduction to the modern theory of dynamical systems*, *Encycl. Math. Appl.* 54, Cambridge Univ. Press (1995) MR Zbl
- [Kopell 1970] **N Kopell**, *Commuting diffeomorphisms*, from “Global analysis” (S-s Chern, S Smale, editors), *Proc. Sympos. Pure Math.* 14, Amer. Math. Soc., Providence, RI (1970) 165–184 MR Zbl
- [Landry 2018] **M Landry**, *Taut branched surfaces from veering triangulations*, *Algebr. Geom. Topol.* 18 (2018) 1089–1114 MR Zbl
- [Landry 2022] **MP Landry**, *Veering triangulations and the Thurston norm: homology to isotopy*, *Adv. Math.* 396 (2022) art. id. 108102 MR Zbl
- [Landry et al. 2023] **MP Landry, Y N Minsky, S J Taylor**, *Flows, growth rates, and the veering polynomial*, *Ergodic Theory Dynam. Systems* 43 (2023) 3026–3107 MR Zbl
- [Landry et al. 2024] **MP Landry, Y N Minsky, S J Taylor**, *A polynomial invariant for veering triangulations*, *J. Eur. Math. Soc. (JEMS)* 26 (2024) 731–788 MR Zbl
- [McMullen 2015] **C T McMullen**, *Entropy and the clique polynomial*, *J. Topol.* 8 (2015) 184–212 MR Zbl
- [Mosher 1996] **L Mosher**, *Laminations and flows transverse to finite depth foliations*, preprint (1996)
- [Parlak 2024] **A Parlak**, *Computation of the taut, the veering and the Teichmüller polynomials*, *Exp. Math.* 33 (2024) 1–26 MR Zbl
- [Ratner 1973] **M Ratner**, *Markov partitions for Anosov flows on n -dimensional manifolds*, *Israel J. Math.* 15 (1973) 92–114 MR Zbl
- [Schleimer and Segerman 2020] **S Schleimer, H Segerman**, *Essential loops in taut ideal triangulations*, *Algebr. Geom. Topol.* 20 (2020) 487–501 MR Zbl
- [Schleimer and Segerman 2021] **S Schleimer, H Segerman**, *From loom spaces to veering triangulations*, preprint (2021) arXiv 2108.10264
- [Schleimer and Segerman \geq 2024] **S Schleimer, H Segerman**, *From veering triangulations to pseudo-Anosov flows and back again*, in preparation
- [Shannon 2021] **M Shannon**, *Hyperbolic models for transitive topological Anosov flows in dimension three*, preprint (2021) arXiv 2108.12000

Department of Mathematics, University of California, Berkeley
Berkeley, CA, United States

Département de mathématiques, Université du Québec à Montréal
Montreal, QC, Canada

ianagol@math.berkeley.edu, tsang.chi_cheuk@uqam.ca

Received: 24 January 2022 Revised: 9 November 2022

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Robert Lipshitz	University of Oregon lipshitz@uoregon.edu
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Robert Oliver	Université Paris 13 bobol@math.univ-paris13.fr
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
David Futer	Temple University dfuter@temple.edu	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
Ian Hambleton	McMaster University ian@math.mcmaster.ca	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Chris Wendl	Humboldt-Universität zu Berlin wendl@math.hu-berlin.de
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Thomas Koberda	University of Virginia thomas.koberda@virginia.edu	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Markus Land	LMU München markus.land@math.lmu.de		


See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2024 is US \$705/year for the electronic version, and \$1040/year (+\$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

AGT peer review and production are managed by EditFlow[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<https://msp.org/>

© 2024 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 24 Issue 6 (pages 2971–3570) 2024

Definition of the cord algebra of knots using Morse theory	2971
ANDREAS PETRAK	
An analogue of Milnor's invariants for knots in 3-manifolds	3043
MIRIAM KUZBARY	
Wall-crossing from Lagrangian cobordisms	3069
JEFF HICKS	
Foliated open books	3139
JOAN E LICATA and VERA VÉRTESI	
Algebraic and Giroux torsion in higher-dimensional contact manifolds	3199
AGUSTIN MORENO	
Locally equivalent Floer complexes and unoriented link cobordisms	3235
ALBERTO CAVALLO	
Strongly shortcut spaces	3291
NIMA HODA	
Extendable periodic automorphisms of closed surfaces over the 3-sphere	3327
CHAO WANG and WEIBIAO WANG	
Bounding the Kirby–Thompson invariant of spun knots	3363
ROMÁN ARANDA, PUTTIPONG PONGTANAPAIAN, SCOTT A TAYLOR and SUIXIN (CINDY) ZHANG	
Dynamics of veering triangulations: infinitesimal components of their flow graphs and applications	3401
IAN AGOL and CHI CHEUK TSANG	
L-spaces, taut foliations and the Whitehead link	3455
DIEGO SANTORO	
Horizontal decompositions, I	3503
PAOLO LISCA and ANDREA PARMA	
The homology of a Temperley–Lieb algebra on an odd number of strands	3527
ROBIN J SROKA	
Hyperbolic homology 3-spheres from drum polyhedra	3543
RAQUEL DÍAZ and JOSÉ L ESTÉVEZ	