



Michael P. Landry · Yair N. Minsky · Samuel J. Taylor

A polynomial invariant for veering triangulations

Received September 1, 2020; revised January 1, 2023

Abstract. We introduce a polynomial invariant $V_\tau \in \mathbb{Z}[H_1(M)/\text{torsion}]$ associated to a veering triangulation τ of a 3-manifold M . In the special case where the triangulation is layered, i.e. comes from a fibration, V_τ recovers the Teichmüller polynomial of the fibered face canonically associated to τ . Via Dehn filling, this gives a combinatorial description of the Teichmüller polynomial for any hyperbolic fibered 3-manifold.

For a general veering triangulation τ , we show that the surfaces carried by τ determine a cone in homology that is dual to its cone of positive closed transversals. Moreover, we prove that this is *equal* to the cone over a (generally non-fibered) face of the Thurston norm ball, and that τ computes the norm on this cone in a precise sense. We also give a combinatorial description of V_τ in terms of the *flow graph* for τ and its Perron polynomial. This perspective allows us to characterize when a veering triangulation comes from a fibration, and more generally to compute the face of the Thurston norm ball determined by τ .

Keywords. Teichmüller polynomial, pseudo-Anosov flow, fibered face, Thurston norm, veering triangulation

Contents

1. Introduction	2
2. Background	6
3. The veering and taut polynomials	12
4. The flow graph and its Perron polynomial	14
5. Carried classes, homology directions, and the Thurston norm	21
6. Relating the polynomials V_τ and Θ_τ	37
7. The fibered case and the Teichmüller polynomial	50
References	56

Michael P. Landry: Department of Mathematics and Statistics, Saint Louis University,
St. Louis, MO 63103, USA; michael.landry@slu.edu

Yair N. Minsky: Department of Mathematics, Yale University, New Haven, CT 06511, USA;
yair.minsky@yale.edu

Samuel J. Taylor: Department of Mathematics, Temple University, Philadelphia, PA 19122, USA;
samuel.taylor@temple.edu

Mathematics Subject Classification (2020): Primary 57K32; Secondary 57K35, 37D20, 57M60

1. Introduction

Veering triangulations of cusped oriented hyperbolic 3-manifolds were introduced by Agol as a means to canonically triangulate certain pseudo-Anosov mapping tori [2]. In particular, Agol showed that for any pseudo-Anosov homeomorphism f on a surface S , the surface S may be punctured along the (invariant) singularities of f so that the resulting mapping torus M of the punctured surface admits a particular type of *layered ideal triangulation* τ , called the *veering triangulation*. This triangulation is layered in the sense that it is built by layering tetrahedra on an ideal triangulation of the punctured surface. The triangulation τ of M is in fact not only an invariant of the monodromy f , but also an invariant of the fibered face of the Thurston norm ball that it determines.

Among all layered triangulations of hyperbolic 3-manifolds, Agol characterized the veering triangulation τ in terms of a combinatorial condition which can be interpreted as a bicoloring of the edges of τ . Answering a question of Agol, Hodgson–Rubinstein–Segerman–Tillmann [25] showed that this combinatorial condition can be satisfied by *nonlayered* triangulations and the resulting class of ideal triangulations is now referred to as *veering*. Veering triangulations have since emerged as an important object in several areas of low-dimension geometry and have attracted much attention. In particular, they have connections to hyperbolic geometry [19, 23, 25] (although they are usually not geometric [20]), are used to study algorithmic problems in the mapping class group [7], encode the hierarchy of subsurface projections associated to their monodromies [35], and, most important for this paper, directly correspond to certain flows on the manifold M [27, 28, 42].

In this paper, we introduce a polynomial invariant of veering triangulations that generalizes McMullen’s Teichmüller polynomial from the fibered setting [33]. We remark that although veering triangulations occur only on manifolds with cusps, our construction recovers the Teichmüller polynomial in general via Dehn filling. Before giving the details, let us informally summarize what we see as the central points of our construction.

- First, in that it recovers the Teichmüller polynomial when the veering triangulation is layered, our construction gives a general, straightforward procedure to compute the Teichmüller polynomial directly from the data of the veering triangulation. Moreover, this data is readily available through either Mark Bell’s program `flipper` [6], which computes the veering triangulation in the fibered setting, or the veering census of Giannopoulos, Schleimer, and Segerman [21]. In fact such an algorithm has already been devised by Parlak [40], and implemented by Parlak, Schleimer, and Segerman (see [21]).

- Second, we show that any veering triangulation τ is naturally associated to a face \mathbf{F}_τ of the Thurston norm ball of M . Indeed, an integral homology class is contained in the cone over \mathbf{F}_τ if and only if it is represented by a surface carried by the underlying branched surface of τ . Moreover, the face \mathbf{F}_τ is fibered *exactly* when τ is layered; so in particular the surfaces carried by a nonlayered veering triangulation determine a non-fibered face of the Thurston norm ball. In this sense, our polynomial invariant provides a generalization of the Teichmüller polynomial as requested by both McMullen [33] and Calegari [10, Question 3.2]. In general, the polynomial is a quotient of the Perron polynomial of a

certain directed graph (the *flow graph* Φ_τ) associated to τ . As in Fried's theory of homology directions [17], the cone of directed cycles of this graph is dual to \mathbf{F}_τ , and therefore \mathbf{F}_τ can also be directly computed from the Perron polynomial of Φ_τ .

- Third is the connection to certain flows, which we investigate in a sequel paper [30]. Following unpublished work of Agol–Guéritaud and work-in-progress of Schleimer–Segerman ([42] and a forthcoming sequel), veering triangulations of hyperbolic manifolds precisely correspond to *pseudo-Anosov flows without perfect fits* as introduced and studied by Fenley [13, Definition 4.2]. These form an important class of flows on 3-manifolds that generalize the suspension flow of a pseudo-Anosov homeomorphism on its mapping torus. When our veering triangulation comes from such a flow φ , we will show that

- the (combinatorially defined) flow graph Φ_τ codes φ 's orbits, in a manner similar to a Markov partition for φ ,
- the Perron polynomial of Φ_τ packages growth rates of closed orbits of φ , after cutting M along certain surfaces transverse to the flow.

We next turn to giving a more formal explanation of our results.

1.1. The veering and taut polynomials

Let M be a 3-manifold with veering triangulation τ , and let $G = H_1(M)/\text{torsion}$. In Section 3, we define the *veering polynomial* V_τ which is an invariant of τ contained in the group ring $\mathbb{Z}[G]$. Essentially by its construction, V_τ comes with a canonical factor Θ_τ , defined up to multiplication by a unit $\pm g \in \mathbb{Z}[G]$, which we call the *taut polynomial*. These polynomials are invariants of modules defined by relations among the edges of the veering triangulation on the universal free abelian cover of M , which are determined by the tetrahedra and faces of τ , respectively. Informally, the tetrahedron relations for the veering polynomial impose conditions modeled on a train track fold, while the face relations for the taut polynomial impose conditions modeled on the switch conditions of a train track.

One main result, which in particular is needed for the explicit connection to the Teichmüller polynomial in the fibered setting (see Theorem B), is the precise relation between these polynomials. For its statement, we note that there is a canonical collection of directed cycles c_1, \dots, c_n in M , which we call *AB-cycles*, whose homology classes are denoted by $g_i = [c_i] \in H_1(M)$. See Section 1.3 for more on their significance.

Theorem A (Factorization). *Suppose $\text{rank}(H_1(M)) > 1$. Then up to multiplication by a unit $\pm g \in \mathbb{Z}[G]$,*

$$V_\tau = \Theta_\tau \cdot \prod_{i=1}^n (1 \pm g_i).$$

We note the possibility that some *AB-cycles* may be trivial in $H_1(M)$ and refer the reader to Theorem 6.1 for a more precise formulation, including an explicit recipe for the signs in the factorization formula.

Since V_τ is always computed as a determinant of a square matrix, we will see that Theorem A can be interpreted as an extension of McMullen’s determinant formula for the Teichmüller polynomial (see Section 7).

1.2. Connection to the Teichmüller polynomial

Next suppose that τ is layered. In this case, it corresponds to a fibered face $\mathbf{F} = \mathbf{F}_\tau$ of the Thurston norm ball and has an associate Teichmüller polynomial $\Theta_{\mathbf{F}} \in \mathbb{Z}[G]$, defined up to a unit. In Section 7, we prove

Theorem B (Teichmüller = taut). *The Teichmüller polynomial $\Theta_{\mathbf{F}}$ agrees with the taut polynomial Θ_τ :*

$$\Theta_{\mathbf{F}} = \Theta_\tau \quad \text{up to a unit } \pm g \in \mathbb{Z}[G].$$

Moreover, if N is any hyperbolic 3-manifold with fibered face \mathbf{F}_N and M is obtained by puncturing N along the singular orbits of its suspension flow, then

$$\Theta_{\mathbf{F}_N} = i_*(\Theta_\tau),$$

where i_ is induced by the inclusion $i: M \rightarrow N$ and τ is the veering triangulation associated to M .*

The above statement combines Theorem 7.1 and Proposition 7.2. Note that Theorem B, together with Theorem A, gives a way to compute the Teichmüller polynomial for any fibered hyperbolic 3-manifold by first puncturing along singular orbits. In the layered (i.e. fibered) setting, no AB -cycle can be trivial by Theorem E below.

In fact, using our work Parlak [40] describes an algorithm to compute the veering and taut polynomials given a (possibly nonlayered) veering triangulation. In separate work [41], she also relates the taut polynomial defined here to the Alexander polynomial, thereby generalizing a result of McMullen for the Teichmüller polynomial [33] via Theorem B.

We conclude by noting that there has been recent interest in developing algorithms to compute the Teichmüller polynomial and several special cases were previously considered in [5, 8, 31].

1.3. Faces of the Thurston norm ball

Moving beyond the fibered setting, we obtain results for general veering triangulations, which we show determine (generally nonfibered) faces of the Thurston norm ball.

For this, we first describe an alternative construction of V_τ . In Section 4.3, we define a graph $\Phi = \Phi_\tau$ in M associated to τ , which we call the *flow graph*. Let $P_\Phi \in \mathbb{Z}[H_1(\Phi)]$ be the Perron polynomial of Φ_τ which is defined as $P_\Phi = \det(I - A_\Phi)$, where A_Φ is an “adjacency matrix” for Φ (see Section 4.3). In Theorem 4.8, we show

Theorem C (From flow graph to veering polynomial). *Let $i_*: \mathbb{Z}[H_1(\Phi)] \rightarrow \mathbb{Z}[G]$ be induced by the inclusion $i: \Phi_\tau \rightarrow M$. Then*

$$V_\tau = i_*(P_\Phi).$$

Next, we recall that the 2-skeleton $\tau^{(2)}$ of τ is a transversely oriented branched surface which can carry surfaces similar to the way a train track on a surface can carry curves (see Section 2.1). We let $\text{cone}_2(\tau)$ be the closed cone in $H_2(M, \partial M)$ positively generated by classes that are represented by the surfaces that τ carries. We call $\text{cone}_2(\tau)$ the *cone of carried classes* and note that it can be explicitly computed as the homology classes of nonnegative solutions to the *switch conditions* for $\tau^{(2)}$.

We show that $\text{cone}_2(\tau)$ is dual to the cone in $H_1(M)$ generated by closed positive transversals to $\tau^{(2)}$ in M , which we call the *cone of homology directions* (see Section 5.1). The cone of homology directions is in turn generated by the support of P_Φ (see Theorem 5.1 and Lemma 5.14):

Theorem D (Duality of cones). *For $\alpha \in H_2(M, \partial M)$, the following are equivalent:*

- (1) $\alpha \in \text{cone}_2(\tau)$,
- (2) $\langle \gamma, \alpha \rangle \geq 0$ for all closed positive transversals γ to τ ,
- (3) $\langle i(c), \alpha \rangle \geq 0$ for each $c \in H_1(\Phi_\tau)$ in the support of P_Φ .

For the connection to faces of the Thurston norm ball, let x denote the Thurston norm on $H_2(M, \partial M)$ and let $B_x(M)$ be its unit ball. In Section 5.2, we associate to τ a combinatorial *Euler class* $e_\tau \in H^2(M, \partial M)$, one definition of which is

$$e_\tau = -\frac{1}{2} \left\langle \prod_{i=1}^n g_i, \cdot \right\rangle \in H^2(M, \partial M),$$

where g_1, \dots, g_n are classes represented by the AB -cycles as in Theorem A. (This is related to the class $[G]$ from [26, Sec. 3] by a multiple of $-\frac{1}{2}$). In Theorem 5.15, we prove

Theorem E (τ determines a face). *The cone of carried classes $\text{cone}_2(\tau)$ is equal to the cone over a (possibly empty) face \mathbf{F}_τ of the Thurston norm ball $B_x(M)$. This cone is characterized by the property that it is the subset of $H_2(M, \partial M)$ on which $-e_\tau = x$.*

Furthermore, the following are equivalent:

- (i) $i_*(\text{supp}(P_{\Phi_\tau}))$ lies in an open half-space of $H_1(M; \mathbb{R})$,
- (ii) there exists $\eta \in H^1(M)$ with $\eta([\gamma]) > 0$ for each closed τ -transversal γ ,
- (iii) τ is layered,
- (iv) \mathbf{F}_τ is a fibered face.

We emphasize that Theorem E also gives a characterization of layeredness of τ (and fiberedness of \mathbf{F}_τ) in terms of a cohomological positivity condition. This makes it a combinatorial analog of Fried's criterion [17, Theorem D] for a flow to be circular, i.e. admit a cross section.

As a last remark, we note that by [29] if \bar{M} is a closed manifold obtained from M by Dehn filling along slopes intersecting the ladderpoles (see Section 2.1.2) of ∂M enough times, the image of $\tau^{(2)}$ under the inclusion $i: M \hookrightarrow \bar{M}$ determines a face $\bar{\mathbf{F}}$ of the Thurston norm ball. Thus $i_*(V_\tau)$ is an object associated to $\bar{\mathbf{F}}$ generalizing the Teichmüller polynomial. Moreover, $\bar{\mathbf{F}}$ is fibered exactly when the image of $\text{supp}(P_{\Phi_\tau})$ lies in an open half-space of $H_1(\bar{M}; \mathbb{R})$.

1.4. Sequel paper: Flows, growth rates, and the veering polynomial

In a followup [30] to the current paper, we develop the connection between the combinatorial approach to the veering polynomial developed here and the pseudo-Anosov flow associated to the veering triangulation. Since this is relevant for motivating our constructions, we briefly summarize the main points of [30].

Suppose that φ is a *pseudo-Anosov flow without perfect fits* on a closed manifold \bar{M} . Then unpublished work of Agol–Guéritaud produces a veering triangulation τ on the manifold M obtained by puncturing \bar{M} along the singular orbits of φ .

First, the flow graph Φ_τ codes the orbits of φ in the following precise sense. There is a map from directed cycles of Φ_τ to closed orbits of φ so that each directed cycle is homotopic to its image. This map is uniformly bounded-to-one and for any closed orbit γ , either γ or γ^2 is in its image. These properties are similar to those of a Markov partition for φ , but we emphasize that Φ_τ is combinatorially defined (Section 4.3) and canonically associated to τ .

Second, from the connection between orbits of the flow and directed cycles of Φ_τ , we use P_Φ to compute the growth rates of orbits of φ . If S is a fiber surface carried by a layered triangulation τ , this recovers well-known properties that McMullen established for the Teichmüller polynomial (using Theorems A and B). If S is a transverse surface carried by a (possibly nonlayered) veering triangulation τ which is *not* a fiber in $\mathbb{R}_+\mathbf{F}_\tau$, then S necessarily misses closed orbits of the flow by Fried’s criterion, and so the usual counting results are not possible.

However, we can consider the manifold $M|S$ obtained by cutting M along S and ask whether there is a class $\xi \in H^1(M|S)$ that is positive on the surviving closed orbits. In this case, we show how the growth rate of closed orbits in $M|S$ with respect to ξ is recorded by the flow graph Φ_τ and its Perron polynomial P_Φ . This analysis includes the special case where S is in the boundary of a fibered cone $\mathbb{R}_+\mathbf{F}_\tau$, and our results are new even in this setting.

2. Background

Here we record some background that we will need throughout the paper. In what follows, all 3-manifolds are assumed to be connected and oriented.

2.1. Veering triangulations

We begin by defining a taut ideal triangulation following Lackenby [26] (see also [25]). Such triangulations are also called *transverse taut* by e.g. [19].

A *taut ideal tetrahedron* is an ideal tetrahedron (i.e. a tetrahedron minus its vertices) along with a coorientation on each face such that two of its faces point into the tetrahedron and two of its faces point out of the tetrahedron. The inward pointing faces are called its *bottom faces* and the outward faces are called its *top faces*. Each of its edges is then assigned angle π or 0 depending on whether the coorientations on the adjacent faces agree or disagree, respectively. See Figure 1.




Fig. 1. A (truncated) taut tetrahedron with its face coorientations and angles.

An ideal triangulation of M is *taut* if each of its faces has been cooriented so that each ideal tetrahedron is taut and the angle sum around each edge is 2π . The local structure around each edge e is as follows: e includes as a π -edge in two tetrahedra. For the other tetrahedra meeting e , e includes as a 0-edge and these tetrahedra form the *fan* of e . We observe that the fan of e has two *sides* each of which is linearly ordered by the coorientation on faces. See, for example, Figure 3, where the coorientation points upwards.

A *veering triangulation* τ of M is a taut ideal triangulation of M in which each edge has a consistent *veer*; that is, each edge is labeled to be either *right* or *left* veering such that each tetrahedron of τ admits an orientation preserving isomorphism to the model veering tetrahedron pictured in Figure 2, in which the veers of the 0-edges are specified: right veering edges have positive slope and left veering edges have negative slope. The π -edges can veer either way, as long as adjacent tetrahedra satisfy the same rule. In other words (cf. [23]), each oriented taut tetrahedron t of τ can be realized as a thickened rhombus in $\mathbb{R}^2 \times \mathbb{R}$ with angle π at its diagonal edges and angle 0 at its side edges such that its vertical diagonal lies above its horizontal diagonal and its right/left veering side edges have positive/negative slope, respectively.

Note that these conditions imply that the 0-edges around a tetrahedron t have alternating veers, and that each face has edges that veer in both directions. From our definition it also follows that for each edge e of τ , each side of e is nonempty (i.e. contains a tetrahedron). This was observed in [25, Lemma 2.3] and is also part of Agol's original definition [2, Definition 4.1]. To prove it, note that along a bottom face of a tetrahedron we encounter the π -edge followed by the left veering followed by the right veering edge in the counter-clockwise order, using the coorientation on faces. However, along a top face we encounter




Fig. 2. On the left is a model veering tetrahedron; right veering edges are red, left veering edges are blue. If this tetrahedron appears in a veering triangulation, then the veers of the top and bottom (π -) edges are determined by adjacent tetrahedra. In the center is an upward flat triangle, and on the right is a downward flat triangle.

the π -edge followed by the right veering followed by the left veering edge, and so two tetrahedra cannot be glued along faces so that their π -edges are identified.

A veering (or taut) triangulation is said to be *layered* if it can be built by stacking tetrahedra onto a triangulated surface and quotienting by a homeomorphism of the surface (for another, more formal definition, we refer the reader to [42, Definition 2.15]). Finally, we recall that the constructions of layered veering triangulations of Agol [2] and Guéritaud [23] start with a pseudo-Anosov homeomorphism f of a surface S and produces a veering triangulation τ on the mapping torus M of $f: S \setminus \{\text{singularities}\} \rightarrow S \setminus \{\text{singularities}\}$.

Remark 2.1 (Veering definitions). For us, a veering triangulation of M is a taut ideal triangulation such that each edge has a consistent veer. Elsewhere in the literature, this is known as a *transverse veering triangulation* [19, 25, 42]. A slightly more general definition can be given where the taut structure is replaced by a *taut angle structure*, which does not impose a coorientation on faces. The two conditions, however, are equivalent up to a double cover [19, Lemma 5.4] and agree for layered veering triangulations, which was the setting of Agol’s original definition [2].

2.1.1. The 2-skeleton $\tau^{(2)}$ as a branched surface. As observed by Lackenby [26], the taut structure of τ naturally gives its 2-skeleton $\tau^{(2)}$ the structure of a transversely oriented branched surface in M . (See [15, 39] for general facts about branched surfaces). The smooth structure on $\tau^{(2)}$ can be obtained by, within each tetrahedron, smoothing along the π -edges and pinching along the 0-edges, thus giving $\tau^{(2)}$ a well-defined tangent plane field at each of its points. See Figure 3.

With this structure, the branching locus of $\tau^{(2)}$ is the disjoint union of edges of τ , and the sectors (i.e. the complementary components of the branching locus) are the faces of τ .

The branched surface $\tau^{(2)}$ has a branched surface fibered neighborhood $N = N(\tau^{(2)})$ foliated by intervals such that collapsing N along its I -fibers recovers $\tau^{(2)}$. The transverse orientation on the faces of τ consistently orients the fibers of N , and a properly embedded oriented surface S in M is *carried* by $\tau^{(2)}$ if it is contained in N where it is positively




Fig. 3. The branched surface $\tau^{(2)}$ near an edge of τ .

transverse to its I -fibers. We also say that S is carried by τ . Note that up to isotopy a surface may be carried by τ in different ways.

A carried surface S embedded in N transverse to the fibers defines a nonnegative integral weight on each face of τ given by the number of times the I -fibers over that face intersect S . These weights satisfy the *matching (or switch) conditions* stating that the sum of weights on one side of an edge match the sum of weights on the other side. Conversely, a collection of nonnegative integral weights satisfying the matching conditions gives rise to a surface embedded in N transverse to the fibers in the usual way.

More generally, any collection of nonnegative weights on faces of τ satisfying the matching conditions defines a nonnegative relative cycle giving an element of $H_2(M, \partial M; \mathbb{R})$ and we say that a class is *carried* by $\tau^{(2)}$ if it can be realized by such a nonnegative cycle. Just as with surfaces, a carried class can be represented by more than one nonnegative cycle on faces.

We conclude by observing the following:

Lemma 2.2. *If $\alpha \in H_2(M, \partial M; \mathbb{Z})$ is carried by $\tau^{(2)}$, then it is realized by a nonnegative integral cycle on $\tau^{(2)}$ and hence an embedded surface carried by $\tau^{(2)}$.*

In fact, the lemma holds for any transversely oriented branched surface.

Proof. Let $W = \mathbb{R}^F$ be the weight space of $\tau^{(2)}$ and let $Z \subset W$ be the subspace of relative 2-cycles, i.e. weights satisfying the branching conditions. Let $A: Z \rightarrow H_2(M, \partial M; \mathbb{R})$ be the map to relative homology. We can choose rational bases for Z and H_2 so that A is represented by an integer matrix.

Let $P \subset Z$ be the set of solutions to $Ax = \alpha$. Because α is carried by $\tau^{(2)}$, P contains a point m lying in the nonnegative orthant $\mathbb{R}_{\geq 0}^F$. Let Q denote the face of the nonnegative orthant containing m in its relative interior. Since $P \cap Q$ is nonempty and cut out by integer equations, there is a rational point $w \in P \cap Q$. (Here we are using the fact that for any rational linear map $L: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and rational $w \in \mathbb{R}^m$, if the equation $Lx = w$ has solutions then rational solutions exist and are dense among all solutions.)

Clearing denominators, there is an integer $n \geq 0$ such that $n\alpha$ is represented by the integral cycle nw , giving rise to an embedded surface S . By [45, Lemma 1], S is a union of n surfaces, each representing α and carried by $\tau^{(2)}$. ■

2.1.2. The veering triangulation as seen from ∂M . If we truncate the tips of the tetrahedra of τ , as in Figure 1, we obtain a compact manifold whose boundary components are tori. For details, see [25, 26]. We denote this manifold \mathring{M} and continue to use τ to refer to the modified veering structure.

In what follows, we will often write ∂M to mean $\partial \mathring{M}$ and $(M, \partial M)$ to mean $(\mathring{M}, \partial \mathring{M})$. This simplifies our notation and should cause no confusion.

The intersection $\partial\tau = \tau^{(2)} \cap \partial M$ is a cooriented train track with each complementary component a *flat triangle* (i.e. a bigon with three branches of $\partial\tau$ in its boundary) that corresponds to the tip of truncated taut tetrahedron. See the right-hand side of Figure 2. The veering structure of τ determines (and is determined by) the structure of these induced train tracks. We will recall some facts here that are needed for Section 5.1.1, but we refer the reader to [19, 29] for a more detailed analysis.

Each flat triangle T on ∂M complementary to the track $\partial\tau$ has two vertices at cusps, corresponding to 0-edges of the associated tetrahedron, and one smooth vertex, corresponding to the π -edge. If the coorientation at the smooth vertex points out of T , then T is called an *upward* triangle, and otherwise it is called *downward*. Note that the veer of the τ -edges corresponding to the 0-vertices (i.e. cusps) of T are determined by whether T is upward or downward, as in Figure 2.

The flat triangles of ∂M are organized into *upward and downward ladders* as follows (see [19, Observation 2.8]): For each component of ∂M , the union of all upward flat triangles is a disjoint collection of annuli such that each annulus A in the collection, called an *upward ladder*, has ∂A carried by $\partial\tau$ and all other branches of $\partial\tau$ meeting A (called *rungs*) join its two boundary components. We define a *downward ladder* similarly, and observe that each component of ∂M is an alternating union of upward and downward ladders meeting along their common boundary (called *ladder poles*). See Figure 4.




Fig. 4. The tessellation of ∂M by flat triangles organized into upward and downward ladders, colored green and cyan respectively for compatibility with Figure 1.

2.2. Thurston norm

Suppose that M is a finite volume hyperbolic manifold. Then the *Thurston norm* x is a norm on the vector space $H_2(M, \partial M; \mathbb{R})$ which extends the formula

$$x(\alpha) = \min\{-\chi(S)\},$$

where $\alpha \in H_2(M, \partial M; \mathbb{Z})$ and the minimum is over all properly embedded surfaces representing α without sphere or disk components [45]. At times it will be convenient to identify $H_2(M, \partial M; \mathbb{R})$ with $H^1(M; \mathbb{R})$ via Lefschetz duality.

The unit ball $B_x = B_x(M)$ in $H_2(M, \partial M; \mathbb{R})$ with respect to x is called the *Thurston norm ball* and it is a rational finite-sided polyhedron. There is a (possibly empty) collection of top dimension faces of B_x , called *fibred faces*, such that

- every class α representing a fiber S in a fibration of M over S^1 is contained in the interior of the cone $\mathbb{R}_+ \mathbf{F}$ over a fibred face \mathbf{F} of B_x ,
- every primitive integral class α in the interior of the cone over a fibred face represents a fiber in a fibration of M over S^1 [45].

According to [3] and [35, Proposition 2.7], the Agol–Guéritaud construction applied to any fiber in a fixed fibred face \mathbf{F} produces the same veering triangulation τ . In particular, a *layered* veering triangulation τ of M is canonically associated to a fibred face \mathbf{F}_τ of $B_x(M)$.

In general, we say that a properly embedded surface S is *taut* if no component of S is nullhomologous and S is Thurston norm minimizing, i.e. $-\chi(S) = x([S])$. For example, fibers are necessarily taut.

Returning to the situation of interest, Lackenby [26, Theorem 3] proved that for any taut ideal triangulation, the underlying branched surface is *taut* in the sense that every surface it carries is taut.

2.3. Polynomials and specialization

Fix a finitely generated free abelian group G and denote its group ring with integer coefficients by $\mathbb{Z}[G]$. Recall that $\mathbb{Z}[G]$ is a UFD.

Let $P \in \mathbb{Z}[G]$ and write $P = \sum_{g \in G} a_g \cdot g$. The *support* of P is

$$\text{supp}(P) = \{g \in G : a_g \neq 0\}.$$

For $P \in \mathbb{Z}[G]$ with $P = \sum_{g \in G} a_g \cdot g$ and $\alpha \in \text{Hom}(G, \mathbb{R})$ the *specialization* of P at α is the single variable expression

$$P(u^\alpha) = \sum_{g \in G} a_g \cdot u^{\alpha(g)} \in \mathbb{Z}[u^r : r \in \mathbb{R}].$$

If G has rank n and is written multiplicatively, we can identify $\mathbb{Z}[G]$ with $\mathbb{Z}[t_1^{\pm 1}, \dots, t_n^{\pm 1}]$. Making this identification, specialization is the image of P under the homomorphism $\alpha_*: \mathbb{Z}[t_1^{\pm 1}, \dots, t_n^{\pm 1}] \rightarrow \mathbb{Z}[u^r : r \in \mathbb{R}]$ obtained by replacing t_i with $u^{\alpha(t_i)}$.

3. The veering and taut polynomials

Let M be a 3-manifold with a veering triangulation τ . Denote the sets of edges, faces, and tetrahedra of τ by E , F , and T , respectively. Since $\chi(M) = 0$, we see that $|E| = |T| = \frac{1}{2}|F|$. Each tetrahedron has a unique *bottom* edge and this induces a bijection from tetrahedra to edges that we will use throughout. If \mathbf{e} is the bottom of the tetrahedron t , then we also say that t lies *above* \mathbf{e} .

As τ is veering, each of its edges is labeled to be either right or left veering. See Section 2.1 and Figure 2. Each abstract veering tetrahedron t gives a relation among its edges. Let \mathbf{b} and \mathbf{t} denote the bottom and top edges of t , respectively. Among the four side edges of t , let \mathbf{s}_1 and \mathbf{s}_2 be those which have the *opposite* veer from \mathbf{t} . Then we have the *tetrahedron relation*

$$\mathbf{b} = \mathbf{t} + \mathbf{s}_1 + \mathbf{s}_2, \quad (1)$$

associated to t . We also say that this tetrahedron relation is associated to the bottom edge \mathbf{b} of t . Informally, this relation mimics the map on a measured train track induced by a fold. See Figure 5.

Let $G = \text{Hom}(H^1(M), \mathbb{Z})$ be the first homology of M modulo torsion and let \tilde{M}^{ab} be the associated covering space. This is the universal free abelian cover of M , and its deck group is G . Let $\tilde{\tau}$ be the preimage of τ . Note that edges, faces, and tetrahedra of $\tilde{\tau}$ are in bijective correspondence with $E \times G$, $F \times G$, and $T \times G$. The correspondence is determined by any choice of lifts of simplices from τ .

We define the *edge module* $\mathcal{E}(\tilde{\tau})$ to be the free \mathbb{Z} -module on the edges of $\tilde{\tau}$ modulo the relations from equation (1) for each tetrahedron of $\tilde{\tau}$. The action of G on \tilde{M}^{ab} by deck transformations makes $\mathcal{E}(\tilde{\tau})$ into a module over the group ring $\mathbb{Z}[G]$ and we will henceforth consider $\mathcal{E}(\tilde{\tau})$ as a $\mathbb{Z}[G]$ -module. Now *choose* a lift of each edge of τ to $\tilde{\tau}$. Since the free \mathbb{Z} -module on the edges of $\tilde{\tau}$ is isomorphic to $\mathbb{Z}[G]^E$ as a $\mathbb{Z}[G]$ -module, $\mathcal{E}(\tilde{\tau})$ has the presentation

$$\mathbb{Z}[G]^E \xrightarrow{L} \mathbb{Z}[G]^E \rightarrow \mathcal{E}(\tilde{\tau}) \rightarrow 0, \quad (2)$$

where L maps \mathbf{b} to $\mathbf{b} - (\mathbf{t} + \mathbf{s}_1 + \mathbf{s}_2)$. That is, the image of each edge is determined by the tetrahedron relation for the tetrahedron lying above that edge.

We define the *veering polynomial* of τ to be the element

$$V = V_\tau = \det(L) \in \mathbb{Z}[G].$$

We remark that the map L can be written in the form $L = I - A$, where I is the identity matrix and A is a matrix with coefficients in G . In Section 4, we will see that A can be interpreted as the adjacency matrix for a directed graph (the *flow graph* of Section 4.3) associated to τ .

The following lemma shows that V_τ is well-defined:

Lemma 3.1. *The veering polynomial $V_\tau = \det(L) \in \mathbb{Z}[G]$ depends only on τ and not on the choice of lifts of edges to $\tilde{\tau}$.*

Proof. If e is an edge of τ and \tilde{e} is the chosen lift in $\tilde{\tau}$, then the effect of replacing \tilde{e} with $g \cdot \tilde{e}$ amounts to conjugating the matrix L by the $|E| \times |E|$ matrix which is obtained from the identity matrix by replacing the 1 in the diagonal entry corresponding to e with g . This does not affect the determinant. ■

3.1. The face module and Θ_τ

Our construction of the veering polynomial V_τ makes explicit use of the veering structure of τ . Here, we define a closely related polynomial, which we call the taut polynomial of τ , whose construction uses only the taut structure of τ (see Section 2.1). Although we begin to describe the connection between the two polynomials in this section, their precise relationship will be fully explained in Section 6. In Section 7 we will show that when τ is layered, its taut polynomial is *equal* to the Teichmüller polynomial of the associated fibered face of the Thurston norm ball. Taken together, these results will give the explicit connection between V_τ and the Teichmüller polynomial.

Any face f of the veering triangulation lies at the bottom of a unique tetrahedron t and has a distinguished edge \mathbf{b} which is the bottom edge of t . We often call \mathbf{b} the *bottom edge* of f . Let \mathbf{x} and \mathbf{y} be the other edges of f . Then the *face relation* associated to f is

$$\mathbf{b} = \mathbf{x} + \mathbf{y}. \quad (3)$$

Since there are two faces lying at the bottom of tetrahedron, t has two associated face relations.

Again, identifying $\mathbb{Z}[G]^E$ with the free module of edges of $\tilde{\tau}$, we quotient by the relations given in equation (3) for each face of $\tilde{\tau}$ to obtain the *face module* $\mathcal{E}^\Delta(\tilde{\tau})$. We have the presentation

$$\mathbb{Z}[G]^F \xrightarrow{L^\Delta} \mathbb{Z}[G]^E \rightarrow \mathcal{E}^\Delta(\tilde{\tau}) \rightarrow 0, \quad (4)$$

where L^Δ is determined by mapping a face f to its associated relation $\mathbf{b} - (\mathbf{x} + \mathbf{y})$.

The module $\mathcal{E}^\Delta(\tilde{\tau})$ has a well-defined Fitting ideal $I^\Delta \subset \mathbb{Z}[G]$ generated by determinants of $|E| \times |E|$ submatrices of L^Δ (see, for example, [38, Theorem 1]). The *taut polynomial* $\Theta_\tau \in \mathbb{Z}[G]$ is defined to be the greatest common divisor of the elements of I^Δ . Note that this only determines Θ_τ up to multiplication by a unit $\pm g \in \mathbb{Z}[G]$.

We next observe the following immediate consequence of our definitions and use it to relate the two polynomials.

Lemma 3.2. *For each bottom face f of a tetrahedron t , there is a unique top face f' of t such that the sum of the face relations for f and f' is equal to the tetrahedron relation for t .*

Further, f' is characterized by the property that it meets f within t along the edge that has the same veer as the top edge of t .

Proof. Consider the face relation (3) for f . Recall that the tetrahedron relation for t is $\mathbf{b} = \mathbf{t} + \mathbf{s}_1 + \mathbf{s}_2$, where \mathbf{t} is the top edge of t and $\mathbf{s}_1, \mathbf{s}_2$ are the side edges which have

opposite veer from that of \mathbf{t} . Let $\mathbf{r}_1, \mathbf{r}_2$ be the other two edges of t labeled so that $\mathbf{b}, \mathbf{s}_i, \mathbf{r}_i$ give the bottom faces of t for $i = 1, 2$.

After possibly swapping indices, we may suppose that $\mathbf{b}, \mathbf{s}_1, \mathbf{r}_1$ span the face f and so $\mathbf{b} - (\mathbf{s}_1 + \mathbf{r}_1) = 0$ is the face relation for f . By construction, \mathbf{r}_1 has the same veer as \mathbf{t} and so the edges $\mathbf{t}, \mathbf{r}_1, \mathbf{s}_2$ span a top face of t . Indeed, since the third edge in the face with \mathbf{t} and \mathbf{r}_1 must have opposite veer, this leaves only \mathbf{s}_1 and \mathbf{s}_2 as possibilities and $\mathbf{b}, \mathbf{s}_1, \mathbf{r}_1$ is already a face on the bottom of t . Hence, we see that $\mathbf{t}, \mathbf{r}_1, \mathbf{s}_2$ span a face f' at the top of t .

From this we can see that the tetrahedron glued to t along f' has \mathbf{r}_1 as its bottom edge. Otherwise, the bottom edge would be either \mathbf{t} or \mathbf{s}_2 . However, in the first case \mathbf{t} would have an empty side of its fan and in the second case the new tetrahedron would have four side edges all the same veer (that of \mathbf{t}), either giving a contradiction. We conclude that the face relation for f' is $\mathbf{r}_1 - (\mathbf{t} + \mathbf{s}_2) = 0$.

Finally, summing the face relations for the bottom face f and top face f' gives $\mathbf{b} - (\mathbf{t} + \mathbf{s}_1 + \mathbf{s}_2) = 0$, which is the tetrahedron relation for t . This completes the proof. ■

We record the following observation made in the proof of Lemma 3.2.

Fact 3.3. *Let \mathbf{e} be an edge of τ and let t be a tetrahedron in the fan of \mathbf{e} . Then the top edge of t has the same veer as \mathbf{e} if and only if t is topmost in its side of \mathbf{e} .*

Since the faces relations determine the tetrahedra relations, we have

Corollary 3.4. *There is a surjective $\mathbb{Z}[G]$ -module homomorphism*

$$\mathcal{E}(\tilde{\tau}) \rightarrow \mathcal{E}^\Delta(\tilde{\tau}) \rightarrow 0,$$

and so Θ_τ divides V_τ .

Proof. The surjection is immediately given by Lemma 3.2 along with the definitions of the modules. This implies that I^Δ contains the Fitting ideal for $\mathcal{E}(\tilde{\tau})$ (see e.g. [32, Appendix] or [38, Chapter 3]), which is principally generated by V_τ . This completes the proof. ■

In fact, the quotient polynomial V_τ/Θ_τ can be easily described as a product of polynomials of the form $1 \pm g$ for $g \in G$ (Theorem 6.1). The proof, however, is quite involved and postponed until Section 6.

4. The flow graph and its Perron polynomial

In this section, we reinterpret our construction in terms of the *flow graph* (Section 4.3) associated to the veering triangulation. This will have implications for identifying the cone of classes carried by the veering triangulation in Section 5.

For a directed graph D with n vertices, let A denote the matrix with entries

$$A_{ab} = \sum_{\partial e = b - a} e, \tag{5}$$

where the sum is over all edges e from the vertex a to the vertex b . This is similar to a standard adjacency matrix for a directed graph, but note that A lives in the matrix ring $M_{n \times n}(\mathbb{Z}[C_1(D)])$, where $C_1(D)$ is the group of simplicial 1-chains in D . We call A the *adjacency matrix* for D . The *Perron polynomial* of D is defined to be $P_D = \det(I - A)$.

Following McMullen [34], we define the *cycle complex* $\mathcal{C}(D)$ of D to be the graph whose vertices are directed simple cycles of D and whose edges correspond to disjoint cycles. We recall that P_D equals the *clique polynomial* of $\mathcal{C}(D)$, which in particular shows that $P_D \in \mathbb{Z}[H_1(D)]$ (see [34, Theorem 1.4 and Section 3]). Here, the clique polynomial associated to $\mathcal{C}(D)$ is

$$P_D = 1 + \sum_C (-1)^{|C|} C \in \mathbb{Z}[H_1(D)], \quad (6)$$

where the sum is over nonempty cliques C of the graph $\mathcal{C}(D)$ and $|C|$ is the number of vertices of C . One proves the above formula by writing the determinant as a sum over elements of the symmetric group S_n and relating the sign of a permutation to the number of cycles in its cycle decomposition.

We begin by introducing an alternative construction of (a variant of) the Perron polynomial which mimics the construction of the veering polynomial.

4.1. Vertex modules of labeled graphs

Let D be an arbitrary directed graph with vertex set V , and let $\alpha: H_1(D) \rightarrow G$ be a surjective group homomorphism. As before, we also use α to denote the induced ring homomorphism $\alpha: \mathbb{Z}[H_1(D)] \rightarrow \mathbb{Z}[G]$. We think of α as labeling the cycles of D , and hence labeling each clique in $\mathcal{C}(D)$ by the product in G of its constituent cycles.

Let \tilde{D}_α be the cover of D corresponding to $\ker(\alpha \circ \text{ab}: \pi_1(D) \rightarrow G)$ with deck group G . For example, if α is the identity homomorphism then \tilde{D}_α is the universal free abelian cover of D . Note that by making a choice of lift for each vertex of D and letting G act by deck transformations, we obtain a bijection between the vertex set of \tilde{D}_α and $V \times G$, and similarly for the edges.

Define the α -labeled vertex module $V(\tilde{D}_\alpha)$ of D to be the $\mathbb{Z}[G]$ -module obtained by taking the free module on the vertices of \tilde{D}_α modulo the following relations: for each vertex v of \tilde{D}_α if e_1, \dots, e_k are the edges of \tilde{D}_α with initial vertex v and terminal vertices w_1, \dots, w_n , respectively, then

$$v = w_1 + \dots + w_n.$$

This gives a presentation

$$\mathbb{Z}[G]^V \xrightarrow{L_{D,\alpha}} \mathbb{Z}[G]^V \rightarrow V(\tilde{D}_\alpha) \rightarrow 0, \quad (7)$$

where $L_{D,\alpha}$ is a square matrix determined by $L_{D,\alpha}(v) = v - (w_1 + \dots + w_n)$. We set $P_{D,\alpha} = \det(L_{D,\alpha}) \in \mathbb{Z}[G]$. The following lemma is proven in the same manner as Lemma 3.1.

Lemma 4.1. *The polynomial $P_{D,\alpha}$ depends only on D and α , and not on the choice of lifts of vertices to \tilde{D}_α .*

The key technical result of this section is the following:

Proposition 4.2. *With notation as above, $P_{D,\alpha} = \alpha(P_D) \in \mathbb{Z}[G]$.*

In words, the polynomial $P_{D,\alpha}$ is obtained by replacing the terms of the Perron polynomial with their images under α .

Proof. Choose a maximal tree $T \subset D$, and label each edge of T with the trivial element of G . For each directed edge e outside of T , there is a unique (not necessarily directed) cycle in $T \cup \{e\}$ traversing e in the positive direction. Label e with the image of this loop under α . The corresponding labeling of the vertices of $\mathcal{C}(D)$ (obtained via concatenation of edges) is given by the homomorphism α . This labeling gives D the structure of a G -labeled directed graph D^G . It has an adjacency matrix A^G with entries in the group ring $\mathbb{Z}[G]$, given by replacing the edges in the entries of A in (5) with their G -labels. A proof entirely similar to that of (6) gives $\det(I - A^G) = \alpha(P_D)$ (cf. [4, Theorem 2.14]). Hence it suffices to show that $\det(I - A^G) = P_{D,\alpha}$.

Lift T to a tree \tilde{T} in \tilde{D}_α , which we note contains an orbit representative of each vertex. By Lemma 4.1, we are free to choose our lifts of vertices to be the vertices of \tilde{T} . We claim that for this choice, $I - A^G = L_{D,\alpha}$. For simplicity we reuse the symbol v to denote the vertex of \tilde{T} lying above v . We can write $L_{D,\alpha}v_i = v_i - (g_{i,1}v_1 + \cdots + g_{i,m}v_m)$, where $g_{i,j} \in G$ and each $g_{i,j}v_j$ corresponds to a unique edge $e_{i,j}$ from v_i to v_j . Note that $g_{i,j} = 1 = \text{id}_G$ if and only if $e_{i,j}$ lies in \tilde{T} . If $e_{i,j}$ lies outside of \tilde{T} , then $g_{i,j}$ is the α -image of the loop in T corresponding to $e_{i,j}$. This is exactly the description of the matrix $I - A^G$, so $I - A^G = L_{D,\alpha}$. In conclusion,

$$P_{D,\alpha} = \det(L_{D,\alpha}) = \det(I - A^G) = \alpha(P_D). \quad \blacksquare$$

We next turn to applying these results to a particular directed graph associated to a veering triangulation.

4.2. The stable branched surface B^s and the dual graph Γ_τ

Here we define the *stable branched surface* B^s in M associated to the veering triangulation τ . It will play an important role in this section and the next. The branched surface B^s also appears in [42], where it is called the *upper branched surface in dual position*.

Topologically, B^s is the 2-skeleton of the dual complex of τ in M . For each tetrahedron t we define a smooth structure on $B_t^s = B^s \cap t$ as follows: if the top edge of t is left veering, then we smooth according to the left-hand side of Figure 5 and otherwise we smooth according the right-hand side. It remains to show that this defines a smooth structure globally, i.e. it agrees across the faces of τ . Note that for a face f of t , $f \cap B_t^s$ is a train track with three branches and a single interior switch. So $f \cap B_t^s$ has a single large branch.




Fig. 5. The stable branched surface B^s in a tetrahedron t depending on whether the top edge of t is left or right veering. The tetrahedron is a rotated version of the usual veering tetrahedron of Figure 2 in order to make the picture easier to draw.

Lemma 4.3. *Let f be a face of τ and let e be its bottom edge. For either of the two tetrahedra t containing f , $f \cap B_t^s$ has its large branch meeting e . Hence B^s has a well-defined smooth structure making it a branched surface.*

We note that the smooth structure on B^s is characterized by its intersection with each face, as given in the first sentence of the lemma.

Proof. We argue within a single tetrahedron t .

If f is a bottom face of t , then this is immediate from Figure 5. If f is a top face of t , then we also see from the figure that the large branch of $f \cap B_t^s$ meets the edge of f with the same veer as the top edge of t . Hence, it remains to observe that the bottom edge of the tetrahedron t' glued to t along f has the same veer as the top of t . This was established in the proof of Lemma 3.2 (cf. Fact 3.3). ■

We note that since it is topologically a spine for the ideal triangulation τ , B^s is a deformation retract of M . This can be seen directly in Figure 5.

The 1-skeleton of B^s is the graph Γ_τ dual to τ whose edges are directed by the coorientation on the faces of τ . We call Γ_τ (or simply Γ) the *dual graph*. Alternatively, Γ can be described as the graph with a vertex interior to each tetrahedron and a directed edge crossing each cooriented face from the vertex in the tetrahedron below the face to the vertex in the tetrahedron above the face. We will always view Γ as embedded in M .

Note that since Γ is the 1-skeleton of B^s onto which M deformation retracts, the map $\pi_1(\Gamma) \rightarrow \pi_1(M)$ is surjective.

4.3. The flow graph Φ_τ

For a veering triangulation τ , its *flow graph* is a directed graph denoted by Φ_τ (or simply Φ if there is no chance of confusion) and defined as follows: the vertices of Φ are in correspondence with τ -edges, and for each tetrahedron t with relation (1), there are directed edges from \mathbf{b} to each of \mathbf{t} , \mathbf{s}_1 , and \mathbf{s}_2 . That is, there are Φ -edges from the bottom τ -edge of each tetrahedron to its top τ -edge and the two side τ -edges whose veer is opposite that of the top τ -edge. Let \mathbf{e} be a τ -edge with corresponding Φ -vertex v . Then v has outgoing valence 3 and it is a consequence of the lemmas to follow that the incoming valence of v is $n - 3$, where n is the degree of \mathbf{e} . The name “flow graph” is motivated by




Fig. 6. On the left is a local picture of Φ (green) in *standard position*; Φ -edges are oriented out of the vertex lying in the tetrahedron's bottom τ -edge. On the right we have isotoped Φ upwards to lie in *dual position*; the vertex which used to lie on the bottom τ -edge is now equal to the Γ -vertex interior to the tetrahedron.

results in [30] showing that in the presence of a certain flow related to τ , directed cycles in Φ_τ correspond in a uniform way to closed orbits of the flow.

There is an embedding $i: \Phi \hookrightarrow M$ which maps each vertex of Φ into its corresponding τ -edge, and maps each Φ -edge into either a τ -face or a segment connecting the top and bottom τ -edges of a tetrahedron. This can be done in such a way that the interior of each edge is smoothly embedded in B^s , as depicted on the left-hand side of Figure 6. We call this the *standard position* for Φ . This embedding induces a homomorphism $i_*: H_1(\Phi) \rightarrow G = H_1(M)/\text{torsion}$ and hence a ring homomorphism $i_*: \mathbb{Z}[H_1(\Phi)] \rightarrow \mathbb{Z}[G]$. The following lemma shows that these homomorphisms are surjective.

Lemma 4.4. *The inclusion $i: \Phi \rightarrow M$ induces a surjection*

$$i_*: \pi_1(\Phi) \twoheadrightarrow \pi_1(M).$$

Before proving Lemma 4.4 we develop the combinatorics of B^s and its branch locus Γ , and discuss their interactions with Φ . We recall that Γ is the directed graph dual to τ .

A line segment, closed curve, or ray which is smoothly immersed in the branch locus of B^s is called a *branch segment*, *loop*, or *ray* respectively. A branch loop will also be called a *branch cycle* when we wish to think of it as a directed cycle of Γ . Let $p = (e_1, e_2, e_3, \dots)$ be a directed path in Γ , and let v_i be the terminal vertex of e_i . We say v_i is a *branching turn* of p if the concatenation $e_i * e_{i+1}$ is a branch segment, and an *anti-branching turn* otherwise.

Anti-branching turns will play an important role in Section 6, and they can be combinatorially distinguished from branching turns in the following way.

Lemma 4.5 (Veering characterization of branching and anti-branching turns). *Let t be a tetrahedron containing a Γ -vertex v . Let f_1 be a bottom face and let f_2 be a top face of t , and let e_1, e_2 be the corresponding Γ -edges. Let s be the unique τ -edge with the property that f_1 and f_2 are incident to s on the same side, with f_2 lying immediately above f_1 . Then v is a branching turn of (e_1, e_2) if and only if s and the top edge of t have opposite veer. In this case, when Φ is in standard position, there is a Φ -edge lying in f_1 from the bottom τ -edge of t to s .*

The proof of Lemma 4.5 follows from inspecting the two cases pictured in Figure 5.

Each sector A of B^s is a topological disk pierced by a single τ -edge, as in Figure 7. The Γ -edges bounding A are oriented so that exactly one vertex is a source, which we call the *bottom* of A , and one is a sink, which we call the *top* of A . The top and bottom divide the boundary of A into two oriented Γ -paths called *sides*. Each side has at least two Γ -edges because the τ -edge piercing A has a nonempty fan on each side. The following lemma says that if you remove the last edge in any side of any sector of B^s , the resulting path is a branch segment, and that the entire side is never a branch segment.

Lemma 4.6 (Sectors and turns). *Let A be a sector of B^s and let p be a side of A considered as a directed path in Γ from the bottom to the top of A . The last turn of p is anti-branching, and all other turns are branching.*

Proof. Let \mathbf{e} be the τ -edge piercing A . The side p corresponds to a side of a fan of \mathbf{e} , and each turn of p lies interior to one of the corresponding tetrahedra. As we observed in Fact 3.3, the top τ -edge of one of these tetrahedra will have the same veer as \mathbf{e} if and only if it is topmost in the side of the fan. By Lemma 4.5, the lemma is proved. ■

By isotoping Φ upward from standard position we can arrange for the vertex sets of Γ and Φ to coincide, and for Φ to still lie in B^s with the interior of each Φ -edge smoothly embedded. This isotopy pushes a Φ -vertex in a τ -edge e onto the Γ -vertex interior to the tetrahedron above e . See the right-hand side of Figure 6. We call this the *dual position* for Φ . Note that in dual position, the interior of each Φ -edge is disjoint from Γ and positively transverse to $\tau^{(2)}$. From this point forward we assume Φ is in dual position unless otherwise stated. Since the two inclusions of Φ are homotopic, this will have no effect on results that do not explicitly deal with the position of Φ in relation to other objects in M , for example Lemma 4.4.

The following lemma describes the intersection of Φ with a sector of B^s .

Lemma 4.7 (Sectors and Φ). *Let A and p be as in Lemma 4.6. Then if e is a Γ -edge in p , then the bottom vertex of e is connected by a Φ -edge in A to the top of A if and only if e is not topmost in A .*

Proof. It is clear that there is a Φ -edge running from the bottom to the top of A . Further, Lemma 4.5 tells us that there will be a Φ -edge in A from a turn of p to the top of A if and only if the turn is branching. Applying Lemma 4.6 finishes the proof. ■

The picture described by Lemma 4.7 is shown in Figure 7. Note that some sectors may have sides with two edges. However, we see in the following argument that for every Γ -edge e , there is a sector A of B^s such that e lies in a side of A with at least three edges.

Proof of Lemma 4.4. Let e be a Γ -edge with terminal point v . We claim there is a sector S_e and a side p of S_e such that e lies in p and neither endpoint of e is the last turn of p . This can be seen by considering the branching around v : of the (locally) six sectors meeting at v , the sector S_e is determined uniquely by the property that its boundary contains both e and the next edge in the branch cycle containing e (the reader may find




Fig. 7. A sector of B^s with Φ in dual position. The intersection with $\tau^{(2)}$ is shown in gray and Φ -edges are green.

it helpful to consult Figure 5). Since v is thus a branching turn for a side of S_e , we have proved our claim by Lemma 4.6. It follows from Lemma 4.7 that each Γ -edge cobounds a disk in B^s with two Φ -edges. Since $\pi_1(\Gamma)$ surjects $\pi_1(M)$, this proves the lemma. ■

Theorem 4.8. *Let (M, τ) be a veering triangulation and Φ_τ its flow graph. Then*

$$V_\tau = i_*(P_\Phi) = 1 + \sum_{\sigma} (-1)^{|\sigma|} i_*(\sigma),$$

where P_Φ is the Perron polynomial of Φ_τ .

In other words, the veering polynomial is obtained from the Perron polynomial of its flow graph by replacing its directed cycles with the corresponding homology classes in M . We warn the reader that some of these classes may be trivial, as is the case when τ represents a top-dimensional face that is not fibered. This is discussed in Section 5.3.

Proof. Let $\tilde{M}^{\text{ab}} \rightarrow M$ be the universal free abelian cover of M , which has deck group G , and let $\tilde{\Phi}$ be the preimage of Φ . Since $\Phi \rightarrow M$ is π_1 -surjective by Lemma 4.4, the restriction $\tilde{\Phi} \rightarrow \Phi$ is the covering map associated to $\ker(i_* \circ \text{ab}: \pi_1(\Phi) \rightarrow G)$.

Choose a lift of each Φ -vertex to \tilde{M}^{ab} , determining a presentation for $V(\tilde{\Phi}_{i_*})$ and thus a matrix L_{Φ, i_*} as in (7). By Proposition 4.2, $i_*(P_\Phi)$ is equal to $P_{\Phi, i_*} = \det(L_{\Phi, i_*})$. By the correspondence between Φ -vertices and τ -edges, the chosen lifts of Φ -vertices give a family of lifted τ -edges and thus a presentation for $\mathcal{E}(\tilde{\tau})$ with an associated matrix L as in (2). For our choices of generators, the definition of Φ gives that these two presentation matrices are equal. Thus

$$V_\tau = \det(L) = \det(L_{\Phi, i_*}) = i_*(P_\Phi),$$

completing the proof. ■

5. Carried classes, homology directions, and the Thurston norm

Fix a veering triangulation τ of M , let $\Gamma = \Gamma_\tau$ be its dual graph, and let $\Phi = \Phi_\tau$ denote its flow graph. In this section, we first show that various cones in $H_1(M; \mathbb{R})$ naturally associated to these directed graphs are equal (Theorem 5.1). Then we show that the dual of these cones in $H^1(M; \mathbb{R}) = H_2(M, \partial M; \mathbb{R})$ is precisely the cone over a face of the Thurston norm unit ball (Theorem 5.12) and describe a connection to the veering polynomial (Theorem 5.15).

5.1. Homology directions and carried classes

Let V be a finite-dimensional real vector space and let $C \subset V$ be a *convex polyhedral cone*, which by definition is the nonnegative span of finitely many vectors in V . The *dimension* of C is the dimension of the subspace generated by C . The *dual cone* to C is

$$C^\vee = \{u \in V^* : u(v) \geq 0 \text{ for all } v \in C\} \subset V^*.$$

One sees that $(C^\vee)^\vee = C$. Let $d(C)$ be the dimension of the largest linear subspace contained in C . Then the dimension of C^\vee is $\dim(V) - d(C)$. If $d(C) = 0$, we say C is *strongly convex*. Thus if C is a top-dimensional strongly convex cone (e.g. a cone over a fibered face of $B_x(M)$), then C^\vee is also a top-dimensional strongly convex cone. For a reference on convex polyhedral cones see [18, Section 1.2].

Let D be a directed graph embedded in M . Define $\text{cone}_1(D) \subset H_1(M; \mathbb{R})$ to be the nonnegative span of the images of directed cycles in D under inclusion. We will write $\text{cone}_1^\vee(D) := (\text{cone}_1(D))^\vee$. For the dual graph Γ , we call $\text{cone}_1(\Gamma)$ the *cone of homology directions* of τ . Note that any directed cycle in Γ gives a closed curve in M which is positively transverse to $\tau^{(2)}$, and conversely any closed curve transverse to $\tau^{(2)}$ is homotopic to a directed cycle in Γ . Hence the cone of homology directions of τ is equal to the cone generated by all closed curves which are positively transverse to $\tau^{(2)}$ at each point of intersection.

We define the *cone of carried classes* of τ , denoted $\text{cone}_2(\tau) \subset H_2(M, \partial M; \mathbb{R})$, to be the cone of classes carried by the branched surface $\tau^{(2)}$.

The first goal of this section is to prove the following theorem relating the above cones.

Theorem 5.1 (Cones). *For any veering triangulation τ of M with dual graph Γ and flow graph Φ , we have $\text{cone}_1(\Gamma) = \text{cone}_1(\Phi)$. Moreover,*

$$\text{cone}_2(\tau) = \text{cone}_1^\vee(\Gamma) = \text{cone}_1^\vee(\Phi),$$

where we have identified $H^1(M; \mathbb{R}) = H_2(M, \partial M; \mathbb{R})$ via Lefschetz duality.

We now turn to the combinatorial definitions and observations we need to prove Theorem 5.1.

5.1.1. Relating Γ , Φ , and B^s . An oriented path, ray, or cycle in the dual graph Γ is called a *dual path, ray, or cycle*, respectively. Similarly, we call an oriented path, ray, or cycle in the flow graph Φ a *flow path, ray, or cycle*.

As a first step of the proof of Theorem 5.1, we will prove Proposition 5.7 which asserts that the dual cycles and flow cycles generate the same cone in $H_1(M; \mathbb{R})$, i.e. $\text{cone}_1(\Gamma) = \text{cone}_1(\Phi)$. While it is clear from the transversality of Φ and $\tau^{(2)}$ that each flow cycle is homotopic to a dual cycle, the converse is not necessarily true; it may be necessary to square a dual cycle before it is homotopic to a flow cycle. This is not apparent from the picture we have developed so far of the relationship between Γ and Φ (Lemmas 4.6 and 4.7) which has been purely local in nature. Hence we continue our discussion from Section 4.3 of the interplay between B^s , Γ , and Φ , broadening our scope beyond a single sector. It will be convenient to work in the universal cover \tilde{M} of M . Let $\tilde{\Phi}$, $\tilde{\Gamma}$, and \tilde{B}^s denote the lifts to \tilde{M} of Φ , Γ , and B^s respectively.

Since B^s is topologically the dual 2-complex to τ , each component U of $M \setminus B^s$ is homeomorphic to $T_U \times [0, 1]$, where T_U is the component of ∂M contained in U . We call U a *tube* of B^s . Let \tilde{U} denote a single lift of U to \tilde{M} , which we also call a *tube* of \tilde{B}^s . There is a collection of branch lines which are entirely contained in $\partial\tilde{U}$ which we call *the branch lines of \tilde{U}* . Each connected component of $(\partial\tilde{U} \cap \tilde{B}^s) \setminus \{\text{branch lines of } \tilde{U}\}$ is called a *band* of \tilde{U} . The image of a band of \tilde{U} under the covering projection is an immersed annulus in M which we call a *band* of U .

Lemma 5.2 ($\tau^{(2)}$ inside a tube). *Let \tilde{U} be a tube of \tilde{B}^s covering a tube U of B^s . Let T be the boundary component of M contained in U , and let \tilde{T} be the lift of T contained in \tilde{U} . The intersection $\tilde{\tau}^{(2)} \cap \text{cl}(\tilde{U})$ can be identified with $(\tilde{\tau}^{(2)} \cap \tilde{T}) \times [0, 1]$ under the homeomorphism $\text{cl}(\tilde{U}) \cong \tilde{T} \times [0, 1]$. This gives an identification of the tessellation $\tau^{(2)} \cap \partial\tilde{U}$ with the boundary tessellation of \tilde{T} .*

Proof. This is a consequence of the duality of $\tau^{(2)}$ and B^s . The identification of the two tessellations can be visualized as in Figure 8. ■

We will now use the structure of this tessellation on $\partial\tilde{U}$ to reconstruct the intersections of $\tilde{\Gamma}$ and $\tilde{\Phi}$ with $\partial\tilde{U}$. In summary, this can be done as follows: $\tilde{\Gamma} \cap \partial\tilde{U}$ is the dual graph to the tessellation $\tilde{\tau}^{(2)} \cap \partial\tilde{U}$, and each complementary component of $\tilde{\Gamma} \cap \partial\tilde{U}$ is a sector of \tilde{B}^s . Thus the data of which Γ turns are branching/anti-branching is determined




Fig. 8. Each tip of a tetrahedron is a flat triangle that determines a cell of the tessellation of the corresponding tube's boundary. We show the picture for a downward (left) and upward (right) flat triangles.

by Lemma 4.6, and the combinatorics of $\tilde{\Phi} \cap \partial\tilde{U}$ are determined by Lemma 4.7. The complete picture is shown in Figure 9; we emphasize that each individual sector in the picture is exactly as in Figure 7. We justify this picture and add more detail in the lemmas to follow.




Fig. 9. Part of the boundary of a tube in \tilde{M} and its intersection with $\partial\tilde{\tau}^{(2)}$ (gray), $\tilde{\Gamma}$ (red), and $\tilde{\Phi}$ (green). One band has been shaded light blue.

Taking advantage of the identification between the two tessellations as in Lemma 5.2, we will speak of *ladders*, *ladderpoles*, and *flat triangles* on the boundaries of tubes in \tilde{M} .

Lemma 5.3 (Γ on boundary of tube). *Let $\tilde{U}, U, \tilde{T}, T$ be as above. Then*

- (i) $\tilde{\Gamma} \cap \partial\tilde{U}$ is topologically the dual graph to $\tilde{\tau}^{(2)} \cap \partial\tilde{U}$, and the edge orientations of the former agree with the coorientations of the latter;
- (ii) each component of $\partial\tilde{U} \setminus \tilde{\Gamma}$ is a sector of \tilde{B}^s ,
- (iii) the top of each of these sectors lies in a downward triangle, the bottom of each sector lies in an upward triangle,
- (iv) each branch line of \tilde{U} bisects an upward ladder, and each upward ladder is bisected by a single branch line.

Proof. Statement (i) is a direct consequence of B^s being the 2-complex dual to $\tau^{(2)}$ and Γ being its 1-skeleton. Statement (ii) follows immediately from the fact that $\tilde{\Gamma}$ is the branching locus of \tilde{B}^s . Each sector contains exactly one vertex of the tessellation (corresponding to the edge dual to that sector) and the top of each sector lies in the flat triangle lying atop that vertex. This flat triangle must be downward, giving (iii).

Let λ be an upward ladder in $\partial\tilde{U}$. Since λ comprises a biinfinite sequence of flat triangles such that each flat triangle shares a unique edge with its successor, (i) implies

that there is a unique biinfinite $\tilde{\Gamma}$ -path contained in λ . Because any two consecutive flat triangles visited by γ share a vertex in the tessellation of $\partial\tilde{U}$, any two consecutive edges e, f of γ lie in the boundary of some \tilde{B}^s -sector s . By (iii), neither e nor f is topmost in a side of s , so they define a branching turn by Lemma 4.6. Therefore each upward ladder contains a branch line of \tilde{U} .

It remains to show that no other biinfinite $\tilde{\Gamma}$ -path is a branch line. Since any two branch lines must be disjoint, it suffices to show that a downward ladder λ' in $\partial\tilde{U}$ does not define a branch line. Let t be a flat triangle in λ' and let v be the cusp of t such that the two sides of t meeting v are rungs of λ' . Further suppose that t is the last triangle in the fan of v that is contained in λ' . Then the triangle in λ' directly above t is a downward triangle having v at its bottom. Hence, t is the topmost flat triangle in its side of the fan for v . If γ' is the $\tilde{\Gamma}$ -path determined by λ' , then we note that γ' crosses every rung of λ' . By Lemma 4.6, the turn of γ' at the vertex in t is anti-branching. This completes the proof. ■

Thus each band b of \tilde{U} contains two halves of upward ladders meeting ∂b , and one downward ladder interior to b .

Since Φ is in dual position, the vertex sets of $\tilde{\Phi}$ and $\tilde{\Gamma}$ are the same. The intersection of $\tilde{\Phi}$ with $\partial\tilde{U}$ is therefore a directed graph whose vertices correspond to the flat triangles of the tessellation. We describe its edges in the following lemma.

Lemma 5.4 (Φ on boundary of tube). *Fix $U, \tilde{U}, T, \tilde{T}$ as above. Identifying the vertices of $\tilde{\Phi} \cap \partial\tilde{U}$ with the flat triangles in the tessellation of $\partial\tilde{U}$, the edges of $\tilde{\Phi} \cap \partial\tilde{U}$ have the following properties:*

- (i) *Each downward triangle t has a single outgoing edge to a triangle t' which is downward, lies in the same ladder as t , and is the endpoint of a $\tilde{\Gamma}$ -path starting at t .*
- (ii) *Each upward triangle has two outgoing edges, which have endpoints in different downward ladders.*

Proof. Let t be a downward flat triangle in $\partial\tilde{U}$. Then t is topmost in the fan corresponding to exactly one of its 0-vertices (see Fact 3.3 or consider e.g. Figure 8). Let v be the 0-vertex of t such that t is *not* topmost in the fan corresponding to v , and let t' be the downward flat triangle having v as a 0-vertex. Lemma 4.7 and the picture we have developed in Lemmas 5.2 and 5.3 gives that t has a single outgoing edge, and its endpoint is in t' .

Now let t be an upward flat triangle in $\partial\tilde{U}$. It again follows from Lemmas 4.7, 5.2, and 5.3 that t has two outgoing edges and that they connect to the downward triangle sharing π -vertex with t , and to the downward triangle whose π -vertex is the 0-vertex of t for which t is not topmost in the corresponding fan. One sees that these two downward flat triangles lie in separate upward ladders. ■

For a visual summary of Lemmas 5.3 and 5.4, see Figure 9. As previously remarked, one can recover the picture from only the tessellation on the boundary of a tube by drawing the directed graph dual to the tessellation and applying Lemmas 4.6 and 4.7.

If ℓ is a branch line containing points p and q , then we say q lies below p if there is an oriented branch segment in ℓ from q to p , where we include the empty segment (so p lies below itself).

Lemma 5.5 (Φ in a band). *The restriction of $\tilde{\Phi}$ to a band b in \tilde{B}^s has the following properties:*

- (i) *For every $\tilde{\Phi}$ -vertex q lying in b , there is a unique $\tilde{\Phi}$ -ray $\rho_b(q)$ in $\tilde{\Phi}|_b$ with q as its initial vertex.*
- (ii) *For each $\tilde{\Gamma}$ -vertex p interior to b , the union of the portions of the two branch lines through p lying below p and in b divide b into two components, b_p^+ and b_p^- .*
- (iii) *If q lies in b_p^- on a branch line of ∂b , the ray $\rho_b(q)$ intersects a branch segment in b containing p at a point below or equal to p .*

A diagram of the situation in Lemma 5.5 is shown in Figure 10.




Fig. 10. An illustration of Lemma 5.5.

Proof. Statement (i) follows from Lemma 5.4. Statement (ii) follows from Lemma 5.3 and our understanding of which turns in the boundary of a sector are branching from Lemma 4.6. Statement (iii) follows from (i) and (ii): the ray $\rho_b(q)$ eventually must cross from b_p^- to b_p^+ , and it must do so at a $\tilde{\Gamma}$ -vertex lying below p . ■

5.1.2. Dual cycles, flow cycles, and carried classes. Any dual cycle c in Γ has a decomposition into branch segments meeting at anti-branching (AB -) turns. The following lemma shows that the parity of the number of anti-branching segments composing the cycle is an invariant of its homotopy class.

Lemma 5.6. *If dual cycles c_1 and c_2 in Γ are homotopic in M , then the number of AB -turns in c_1 has the same parity as the number of AB -turns in c_2 .*

Proof. Since M deformation retracts to B^s and Γ is the 1-skeleton of B^s , it suffices to build a homomorphism $\mathbb{Z}: \pi_1(B^s) \rightarrow \mathbb{Z}/2$ such that if c is a dual cycle, then $\mathbb{Z}(c)$ is the number of AB turns mod 2.

The branched surface B^s is the base space of a vector bundle $E \rightarrow B^s$ defined by taking the tangent plane at each point of B^s . For any loop γ in B^s , this pulls back to a plane bundle over a circle. If the pullback bundle is orientable, set $\mathcal{Z}(c) = 0$, and set $\mathcal{Z}(c) = 1$ otherwise. It is a standard fact that \mathcal{Z} defines a homomorphism. In fact, \mathcal{Z} is equal to the map $\pi_1(B^s) \rightarrow \mathbb{Z}/2$ induced by the first Stiefel–Whitney class of $E \rightarrow B^s$; see e.g. [24, Chapter 3].

Now let c be a dual cycle. We will show that

$$\mathcal{Z}(c) \equiv \#\{AB\text{-turns of } c\} \pmod{2}.$$




Fig. 11. When a dual cycle makes an anti-branching turn, we can homotope it to lie in a strip like the green one shown above. The blue strip corresponds to a branching turn.

For this, observe that each time c passes through a tetrahedron it makes either a branching or an anti-branching turn, and this portion of c can be homotoped to lie in one of the two strips shown in Figure 11. From the picture we see that these strips glue together to give an annulus if and only if the number of anti-branching turns is even. Indeed, at any point along the path in the branch locus, a local orientation of the plane field is obtained from the direction of the path and the rule that the branching is “on the left.” Now the AB path in Figure 11 enters the tetrahedron where the branching is on one side and exits the tetrahedron where the branching is on the other side. Hence the orientation of the plane flips, which is to say it disagrees with the orientation carried continuously from the entry point. Since the opposite is true for the branching turn, this gives the desired parity property for the loop. ■

The next proposition establishes half of Theorem 5.1. Although our emphasis here is on the equality of the cones, the stronger fact that for any positive transversal of $\tau^{(2)}$ either it or its square is homotopic to a flow cycle will be central in [30].

Proposition 5.7. *For each dual cycle c of Γ , there is a flow cycle f of Φ such that f is homotopic to c^i for $i = 1, 2$. Moreover, if c is composed of an even number of branch segments, then f is homotopic to c .*

Hence the flow graph and the dual graph determine the same cones in $H_1(M; \mathbb{R})$:

$$\text{cone}_1(\Gamma) = \text{cone}_1(\Phi).$$

Proof. We first explain that every branch curve is homotopic to a flow cycle. Let γ be a branch curve in B^s . Then γ determines a tube T_γ of $M \setminus B^s$. Let b be a band of T_γ . The

core curve of b is homotopic to γ . By Lemma 5.5, there is a Φ -ray lying exclusively in b , and it must be eventually periodic. Hence there is a Φ -cycle c in b , which must then be homotopic to γ^i for $i \geq 1$. The fact that $i = 1$ is a consequence of the uniqueness part of Lemma 5.5 (i).

Next, let c be a dual cycle which is not a branch curve, and let \tilde{c} be a lift to \tilde{M} . Let $g \in \pi_1(M)$ be the deck transformation that generates the cyclic subgroup stabilizing \tilde{c} . Further suppose that g translates \tilde{c} in its positive direction. Express \tilde{c} as a concatenation of $\tilde{\Gamma}$ -edges $(\dots, e_{-2}, e_{-1}, e_0, e_1, e_2, \dots)$ and let p_i be the terminal point of e_i .

We build a $\tilde{\Phi}$ -ray inductively. There are exactly two branch lines through p_0 . Let q_0 be any point on either of these branch lines below p_0 . Suppose q_i is on a branch line below p_i . If p_{i+1} lies on the same branch line containing q_i and p_i (including the case where $q_i = p_i$), let $q_{i+1} = q_i$. If p_{i+1} does not lie in the same branch line as p_i and q_i , let b be the band containing the branch line through p_i and q_i in its boundary and with p_{i+1} in its interior. Then $q_i \in b_{p_{i+1}}^-$, so as in Lemma 5.5 the (unique) ray $\rho_b(q_i)$ intersects a branch line containing p_{i+1} at a point q_{i+1} below or equal to p_{i+1} . Let f_i be the $\tilde{\Phi}$ -segment from q_i to q_{i+1} traversed by $\rho_b(q_i)$. (These two cases are depicted in the first two images in Figure 12, where the points q'_i, q'_{i+1} can be ignored for now.) We will show that the concatenation $\tilde{f} = (f_1, f_2, f_3, \dots)$ of these segments form a preperiodic ray in $\tilde{\Phi}$.




Fig. 12. Cases (i)–(iii) (ordered left to right) proving that \tilde{f} is to the left of \tilde{f}' . In the first case, it is possible that $q'_i = p_i$. In the second and third cases, it is possible that $q_{i+1} = p_{i+1}$. In the third case, it is possible that $q_i = p_i$. The Φ -segments f_i and f'_i are drawn in green.

Let ℓ_i be the branch line containing e_i , and let b_i be the band of \tilde{B}^s determined by the property that it contains ℓ_i in its boundary and contains the first edge of \tilde{c} after e_i not lying in ℓ_i . As a comprehension check, we note that if e_i and e_j are contained in the same branch line, then $b_i = b_j$.

Fix some $j \geq 0$ such that $\ell_j \neq \ell_{j+1}$. Consider p_j , which as a reminder is the terminal point of e_j and the initial point of e_{j+1} . By construction, q_j lies below p_j on ℓ_j or ℓ_{j+1} .

- First, suppose that q_j lies below p_j on ℓ_{j+1} (this includes the case where $q_j = p_j$). Let e_k be the last edge of \tilde{c} contained in ℓ_j . Then $q_j = q_{j+1} = \dots = q_k$, and q_{k+1} is equal to either p_{k+1} or one of finitely many points below p_{k+1} on a branch segment in b_k .
- Next, suppose that q_j lies below p_j on ℓ_j . Then q_{j+1} lies in b_j and is equal to either p_{j+1} or one of finitely many points below p_{j+1} on a branch segment in b_j .

Recall that the deck transformation g stabilizes \tilde{c} . The discussion above shows that there are finitely many vertices r_1, \dots, r_m such that, for each $i > 0$, \tilde{f} must pass through one of $g^i(r_1), \dots, g^i(r_m)$. Moreover, since the construction shows that each subray of \tilde{f} is determined by any initial vertex, it follows that the projection f of \tilde{f} to M is eventually periodic, giving a Φ -cycle that is homotopic to a positive integer multiple of c .

By truncating an initial segment of \tilde{f} we can assume \tilde{f} is periodic under some power of g . More precisely, there exist k, ℓ such that $g^k(q_0) = q_\ell$. By construction, \tilde{f} is determined by any of its vertices and so for any $i \in \mathbb{Z}$, either \tilde{f} has no common vertices with $g^i(\tilde{f})$, or the two are equal. To bound the period of \tilde{f} we first need to discuss some extra structure associated to the branch lines through \tilde{c} .

First, label each ℓ_i with either an L or R such that if ℓ_j and ℓ_{j+1} are adjacent (i.e. intersecting and not equal) branch lines along c , as in the above construction, then they receive opposite labels. Such a labeling is uniquely determined by giving ℓ_1 , the branch line containing the edge e_1 , the label L . Note that this labeling is well-defined: if \tilde{c} leaves a branch line and then returns to it, it does so only after an even number of anti-branching turns by Lemma 5.6.

If p is a vertex of \tilde{c} and ℓ^L and ℓ^R are the branch lines through p with indicated labels, we define an order on the vertices lying below p on ℓ^R and ℓ^L . If q_1, q_2 are vertices on these lines which lie below p , then we say that q_1 is *to the left of* q_2 (at p) if any of the following hold:

- q_1 lies in ℓ^L and q_2 lies in ℓ^R ,
- q_1 and q_2 lie in ℓ^L and q_1 lies below q_2 , or
- q_1 and q_2 lie in ℓ^R and q_2 lies below q_1 .

Now if c is composed of an even number of branch segments (i.e. c has an even number of anti-branching turns), then g preserves the L/R labeling on the branch lines through \tilde{c} , and therefore preserves the order relation defined above. For the remainder of the argument, we assume this to be the case. Otherwise, after replacing c (and g) by its square, the same argument applies.

Let p be any vertex of \tilde{c} and let q be a vertex of \tilde{f} such that q lies below p in a branch line ℓ , as in the construction of \tilde{f} . (If $q = p$, then take ℓ to be the branch line containing the edge of \tilde{c} with initial vertex p . Suppose that ℓ is labeled by L . Set $\tilde{f}' = g^{-1}(\tilde{f})$. If $\tilde{f} = \tilde{f}'$, then the proof is complete. So assume that $\tilde{f} \neq \tilde{f}'$ and so these flow lines have no vertices in common.

By construction, there is also a vertex q' of \tilde{f}' such that q' lies below p on some branch line. Since q and q' are not equal, one lies strictly to the left of the other. For the sake of argument, suppose q lies strictly to the left of q' . We claim that this implies that for any other vertex x of \tilde{c} , the vertex of \tilde{f} below x on some branch line lies strictly to the left of the vertex of \tilde{f} below x on some branch line. (As shorthand, we say that \tilde{f} lies to the left of \tilde{f}' at x .) It suffices to prove the claim for vertices after p along \tilde{c} and we do so by induction.

Suppose that \tilde{f} is to the left of \tilde{f}' at the vertex p_i along \tilde{c} . Let ℓ_i^L and ℓ_i^R be the left and right branch lines through p_i , respectively, and let q_i, q'_i be the vertices of \tilde{f}, \tilde{f}' ,

respectively, that lie below p_i on either ℓ_i^L or ℓ_i^R . Let us suppose that p_{i+1} (the next vertex along \tilde{c}) also lies on ℓ_i^L ; the proof in the opposite case is symmetric. We have the following cases (see Figure 12):

- (i) If q_i and q'_i are on ℓ_i^L , then $q_{i+1} = q_i$ and $q'_{i+1} = q'_i$. Hence, q_{i+1} is to the left of q'_{i+1} at p_{i+1} .
- (ii) If q_i and q'_i are on ℓ_i^R (and $q_i \neq p_i$), then both f_i and f'_i (the segments of \tilde{f} , \tilde{f}' joining q_i, q'_i to q_{i+1}, q'_{i+1} , respectively) are contained in the same band b . Since $\ell_i^L = \ell_{i+1}^L$, we see that q_{i+1}, q'_{i+1} are both contained in ℓ_{i+1}^R . As \tilde{f} and \tilde{f}' are assumed to be distinct, the segments f_i, f'_i do not intersect in b and so their terminal vertices are ordered along ℓ_{i+1}^R in the same manner as they were ordered along ℓ_i^R . Hence, q_{i+1} is to the left of q'_{i+1} at p_{i+1} .
- (iii) If q_i and q'_i are on different branch lines, then it must be that q_i is on ℓ_i^L and q'_i is on ℓ_i^R . In this case, $q_{i+1} = q_i$ and q'_{i+1} lies in ℓ_{i+1}^R (where $q'_{i+1} = p_{i+1}$ is possible). Hence, q_{i+1} is to the left of q'_{i+1} at p_{i+1} .

We have shown that \tilde{f} lies to the left of $\tilde{f}' = g^{-1}(\tilde{f})$ at $g^i(p)$ for each $i \in \mathbb{Z}$. By $\langle g \rangle$ -equivariance, this implies that $g^{i+1}(\tilde{f})$ lies strictly to the left of $g^i(\tilde{f})$ at p for every $i > 0$. This, however, contradicts that \tilde{f} is periodic under some power of g . Hence, we must have $\tilde{f} = g(\tilde{f})$ and so f is homotopic in M to c . ■

For the other half of Theorem 5.1, we need two more lemmas.

Recall that a directed graph is *strongly connected* if for any two of its vertices u and v , there is a directed path from u to v .

Lemma 5.8. *The dual graph Γ is strongly connected. Moreover, any cycle in Γ , not necessarily directed, is homologous to a linear combination of dual cycles.*

Hence, the dual cycles generate $H_1(M; \mathbb{Z})$.

Proof. We first show that Γ is strongly connected. Consider a tetrahedron T , and let $S(T)$ be the union of all tetrahedra accessible from T via a directed path in Γ . It is immediate that $S(T)$ has no top faces on its boundary. Since the tetrahedra composing $S(T)$ have an equal number of top and bottom faces, and each top face is glued to a bottom face, $S(T)$ has no bottom faces in its boundary. Hence $S(T) = M$, proving strong connectivity.

Now, let x be a cycle in Γ . We can break x into a concatenation of segments which are maximal with respect to the property that they either completely agree or disagree with the edge orientations of Γ . We write $x = (f_1, b_1, \dots, f_n, b_n)$, where f and b are chosen to evoke “forward” and “backward”. Denote the initial vertex of f_i by p_i , the terminal vertex by q_i , and observe that this dictates that b_i travels from q_i to p_{i+1} , where indices are taken modulo n .

Let α_i be a directed path from q_i to p_i , which exists by the strong connectivity of Γ . Let A_i be the directed cycle (f_i, α_i) , and let B be the directed cycle $(b_n^{-1}, \alpha_n, \dots, b_1^{-1}, \alpha_1)$. Then x is homologous to $A_1 + \dots + A_n - B$. Since $\pi_1(\Gamma) \rightarrow \pi_1(M)$ is surjective, this completes the proof. ■

Recall that if A is a convex set in \mathbb{R}^d , a *supporting hyperplane* for A is a hyperplane H such that $H \cap \partial A \neq \emptyset$ and A is contained in one of the two closed half-spaces determined by H . We will use the following finite-dimensional version of the Hahn–Banach Theorem:

Lemma 5.9 (Hahn–Banach). *Let C be a strongly convex polyhedral cone in a finite-dimensional vector space V . If K is a subspace of V with $K \cap \text{int}(C) = \emptyset$, then K can be extended to a supporting hyperplane H for C . If $K \cap C = \{0\}$, we can choose H such that $H \cap C = \{0\}$.*

Proof. Apply the separating hyperplane theorem (see, for example, [9, Section 2.5]), to K and $\text{int}(C)$ for the first statement. For the second, slightly enlarge C to C' so that C' is still strongly convex with $K \cap \text{int}(C') = \emptyset$ and apply the same theorem to K and $\text{int}(C')$. ■

The following is a strengthening of [36, Proposition 2.12]. The appeal to Hahn–Banach comes from [44, Theorem I.7] and is made explicit in [34, Theorem 5.1], where the second statement of the lemma is proven.

Lemma 5.10. *Suppose that D is a finite directed graph and $\eta \in H^1(D; \mathbb{R})$ is nonnegative on all directed cycles. Then there is a nonnegative cocycle $m: E(D) \rightarrow \mathbb{R}_{\geq 0}$ representing η .*

If η is positive on all directed cycles, then m can be taken to be positive on all directed edges of D .

Proof. We assume η is not identically 0, since in that case the statement is clearly true. First suppose that D is strongly connected. Let $E = E(D)$, let $Z_1 \subset \mathbb{R}^E$ be the subspace of cycles in D , and let \mathbb{R}_+^E and $\mathbb{R}_{\geq 0}^E$ be the open positive orthant and its closure, respectively.

Note that $Z_1 \cap \mathbb{R}_+^E$ is nonempty by strong connectivity, and that η is strictly positive on this set since $\eta \neq 0$. By Lemma 5.9, $\ker(\eta)$ can be extended to a supporting hyperplane H for $\mathbb{R}_{\geq 0}^E$. Let V be the span of a vector in $Z_1 \cap \mathbb{R}_+^E$, and choose a decomposition $H = \ker(\eta) \oplus W$ for some subspace W . Then $\mathbb{R}^E = V \oplus \ker(\eta) \oplus W$, because neither $\ker(\eta)$ nor W intersect \mathbb{R}_+^E . Thus we can define a linear functional $m: \mathbb{R}^E \rightarrow \mathbb{R}$ which extends η by requiring $m|_W = 0$. Note that m is a nonnegative cocycle representing η . If η is strictly positive on cycles in D , then we can choose H so that $H \cap \mathbb{R}_{\geq 0}^E = \{0\}$, guaranteeing that m is a positive cocycle. The same argument shows that the result holds when D is a union of strongly connected components.

Now suppose that D is arbitrary. Let R be the union of recurrent components of D , i.e. the union of all maximal strongly connected subgraphs. Let A be the complementary subgraph of R in D . Let c be a cocycle representing η , and write $c = c_R + c_A$, where c_R and c_A are supported on R and A . By the case above, c_R can be written as $c_+ + \delta h$ for some 0-cochain h , where c_+ is either positive or nonnegative depending on whether η is positive or nonnegative on directed cycles.

Consider the quotient D_R obtained by collapsing each component of R to a vertex. Note that D_R contains no directed cycles, so the edge orientations induce a partial order

on its vertices and there is a function $f: V(D_R) \rightarrow \mathbb{R}$ compatible with this partial order. We may lift f to a function F on D which is constant on each component of R . Then δF is a coboundary which is 0 on $E(R)$ and positive on $E(A)$. By replacing F with a sufficiently large multiple of itself, we can guarantee that for each edge a of A , $\delta F(a) > |c_A(a)|$. Then $c_+ + c_A + \delta F$ is the desired cocycle. ■

Proposition 5.11. *There is an equality of cones*

$$\text{cone}_1^\vee(\Gamma) = \text{cone}_2(\tau).$$

That is, the classes carried by $\tau^{(2)}$ are precisely those whose algebraic intersection with any closed positive transversal to $\tau^{(2)}$ is nonnegative.

Moreover, any integral class $\alpha \in \text{cone}_1^\vee(\Gamma) \cap H^1(M; \mathbb{Z})$ is represented by a surface S carried by $\tau^{(2)}$.

Proof. Since elements of $\text{cone}_1(\Gamma)$ are nonnegative combinations of directed cycles which intersect the faces of τ positively, it is clear that $\text{cone}_2(\tau) \subset \text{cone}_1^\vee(\Gamma)$.

Now let $0 \neq \eta \in \text{cone}_1^\vee(\Gamma)$ and let $i^*\eta \in H^1(\Gamma)$ be its pullback to Γ . By Lemma 5.10 we can represent $i^*\eta$ by a nonnegative cocycle $m: E(\Gamma) \rightarrow \mathbb{R}_{\geq 0}$ on the edges of Γ . Using the identification between directed edges of Γ and cooriented faces of τ , the cocycle m gives a 2-chain $m_2: \{\tau\text{-faces}\} \rightarrow \mathbb{R}$.

From the duality between Γ and τ , it follows that m_2 is a cycle rel boundary and thus defines a class in $H_2(M, \partial M)$. More explicitly, for a τ -edge e , let σ_e denote the B^s -sector pierced by e . Let ℓ and r be the two sides of σ_e . To check that m_2 is indeed a relative cycle, we need only check that $m(\ell) = m(r)$. Since the loop $i(l \cdot \bar{r})$ is trivial in M ,

$$m(l \cdot \bar{r}) = i^*\eta(\ell \cdot \bar{r}) = \eta(i(\ell \cdot \bar{r})) = 0,$$

and $m(\ell) = m(r)$ as required. Hence m_2 determines a carried class $h \in \text{cone}_2(\tau)$ whose algebraic intersection with any dual cycle c is equal to $\eta(c)$. Since the dual cycles generate $H_1(M)$, h is the Lefschetz dual of η .

The “moreover” statement now follows from Lemma 2.2. ■

5.2. Faces of the Thurston norm ball

Here we show that the veering triangulation always determines an *entire* face of the Thurston norm ball, generalizing what was known for the layered case. We use this to give a criterion (Theorem 5.15) to detect fiberedness, which can be thought of as a combinatorial version of Fried’s condition for a flow to be circular [17, Theorem D].

From its definition, we see that the dual graph $\Gamma = \Gamma_\tau$ is a cycle in M and so represents a homology class $[\Gamma] \in H_1(M)$. Further, for any surface Σ carried by $\tau^{(2)}$, we see that $\chi(\Sigma) = -\frac{1}{2}\langle [\Gamma], [\Sigma] \rangle$, where $\langle \cdot, \cdot \rangle$ is algebraic intersection. This is because the intersection of the carried surface with $[\Gamma]$ equals the number of ideal triangles in an ideal triangulation of the surface.

With this in mind, we define the combinatorial *Euler class* of τ to be

$$e_\tau = -\frac{1}{2} \langle [\Gamma], \cdot \rangle.$$

We have $e_\tau \in H^2(M, \partial M)$. Recall from Section 2.2 that $\tau^{(2)}$ is a taut branched surface. Hence, if Σ is carried by $\tau^{(2)}$ then $-e_\tau([\Sigma]) = x([\Sigma])$, where $x(\cdot)$ is the Thurston norm.

Theorem 5.12 (The whole face and nothing but the face). *The subset of $H_2(M, \partial M)$ on which the Thurston norm x is equal to $-e_\tau$ is exactly $\text{cone}_2(\tau)$. Hence, $\text{cone}_2(\tau)$ is equal to the cone over a face \mathbf{F}_τ of $B_x(M)$.*

We say that τ *determines* the face \mathbf{F}_τ . Here we note that, as in work of Mosher [37], the empty face is allowed. The proof of Theorem 5.12 will require additional terminology that we now turn to explain.

5.2.1. The unstable branched surface B^u . Recall that in Section 4.2 we introduced the stable branched surface B^s as a particular smoothing of the 2-complex of M dual to τ . According to Lemma 4.3, this smoothing is characterized by the condition that for each face f of τ , $f \cap B^s$ is the train track whose large branch meets the bottom edge of f .

In an analogous way, we define the *unstable branched surface B^u* to be the branched surface with the same underlying 2-complex such that for each face f of τ , $f \cap B^u$ is the train track whose large branch meets the top edge of f . Here, the top edge of a face f is the unique top edge of the tetrahedron t that contains f as a top face. The intersection of B^u with a tetrahedron is shown in Figure 13. Just as with the stable branched surface, the branch locus of B^u can be naturally identified with the dual graph Γ . We fix the position of both B^s and B^u in M , but make no assumption about how they sit with respect to each other.




Fig. 13. The unstable branched surface B^u in a tetrahedron t depending on whether the bottom edge of t is left or right veering.

Exactly as in Section 4.3, we can define *unstable branch loops* as smoothly immersed closed curves in the branch locus of B^u .

We note that if S is an embedded surface in M transverse to B^u and B^s , then the intersections $S \cap B^u$ and $S \cap B^s$ are train tracks on the surface S , and we refer to the complementary regions as *patches*. (Here, we make no assumption on the Euler characteristic of a patch.) If S is oriented, then each intersection with Γ has a sign, and since these intersections are in bijection with the cusps of $S \cap B^s$ (and $S \cap B^u$), each cusp of these tracks is either *positive* or *negative* accordingly.

5.2.2. Partial branched surfaces and a result of Landry. To apply results from [29] we replace M with the corresponding compact manifold as follows: First, let \mathring{M} be the manifold obtained by truncating the cusps of the tetrahedra of τ , as in Section 2.1. We continue to use $\tau^{(2)}$ to denote the corresponding branched surface with boundary in \mathring{M} . Now let M be the homeomorphic manifold obtained by attaching a thickened torus $T^2 \times [0, 1]$ to each boundary component of \mathring{M} .

In the terminology of [29], $\tau^{(2)}$ is a *partial branched surface* of M with respect to $U = M \setminus \text{int}(\mathring{M})$. A properly embedded surface $S \subset M$ is *carried* by the partial branched surface $\tau^{(2)} \subset \mathring{M}$ if S has no components completely contained in U , $S \setminus \text{int}(U) \subset \mathring{M}$ is carried by $\tau^{(2)}$, and each component of $S \cap U$ is a properly embedded π_1 -injective annulus in U with either one or both boundary components on $\partial \mathring{M}$.

The following theorem [29, Theorem 8.1] is a key ingredient for Theorem 5.12:

Theorem 5.13 (Landry). *Let S be an incompressible, boundary incompressible surface in M . Further suppose that S has the property that for any surface S' isotopic to S that is transverse to B^u and B^s , either*

- *one of $S' \cap B^u$ or $S' \cap B^s$ has a nullgon or monogon patch, or*
- *for both tracks $S' \cap B^u$ and $S' \cap B^s$, each negative cusp belongs to a bigon patch.*

Then S is isotopic to a surface carried by the partial branched surface $\tau^{(2)} \cap \mathring{M}$.

We remark that it is observed in [29] that the tracks $S \cap B^{u/s}$ never have monogon patches.

5.2.3. The proof of Theorem 5.12. Before beginning the proof, we recall a formula for computing the Euler characteristic of surfaces using train tracks (see e.g. [11]). For a surface C with finitely many punctures and cusped boundary, we can define

$$\text{index}(C) = 2\chi(C) - (\text{no. of boundary cusps}).$$

If t is a train track on a surface S of finite genus with finitely many punctures, then

$$2\chi(S) = \sum_C \text{index}(C),$$

where the sum is taken over all patches of S with respect to t . Note that if $\chi(S) \leq 0$, then the only patches with positive index are nullgons and monogons.

Proof of Theorem 5.12. Let $W \subset H_2(M, \partial M)$ be the cone on which $-e_\tau = x$. Any class contained in $\text{cone}_2(\tau)$ is represented by a surface carried by $\tau^{(2)}$ by Lemma 2.2. As discussed at the beginning of this section, such a surface Σ is taut and has $-e_\tau([\Sigma]) = x([\Sigma])$. This yields the containment $\text{cone}_2(\tau) \subset W$.

To show the reverse containment it suffices to produce, for any integral class $\alpha \in W$, a surface representing α and carried by $\tau^{(2)}$.

Let S_α be any taut surface representing α . We will use Theorem 5.13 to prove that S_α is isotopic to a surface carried by the partial branched surface $\tau^{(2)} \cap \mathring{M}$. For this, suppose

that S is a surface transverse to the branched surfaces B^u and B^s and isotopic to S_α . Let $t^s = S \cap B^s$ and $t^u = S \cap B^u$ be the corresponding train tracks on S . By Theorem 5.13 it suffices to assume that neither of these tracks has a nullgon or monogon patch and to show that every negative cusp belongs to a bigon patch. The arguments for t^s and t^u are identical, so we work with $t = t^u$.

Recall that there is a bijection between intersections $S \cap \Gamma$ and cusps of t , and that each cusp of t is the cusp of exactly one patch $D \subset S$. (Since S misses the vertices of Γ , the track t is generic.) Also recall that the sign of a cusp is the sign of the corresponding intersection point of Γ and S .

It is clear that the smooth annular patches (i.e. topological annuli without cusps) do not contribute to the signed number of cusps, and we claim the same is true for the patches that are topological disks (i.e. unpunctured n -gons). Indeed, suppose that D is a patch which is a topological disk contained in a tube V of B^u , defined as in Section 5.1.1. Since D is a disk in V which is unpunctured, ∂D is a trivial curve in ∂V . Hence, ∂D meets each unstable branch curve of ∂V (i.e. the intersection of the cusps of B^u with V) with algebraic intersection number 0. Since this intersection number is also the signed number of cusps in D , we see that the signed number of cusps for any disk is also 0.

Partition the set of patches on S into the set of topological disks and smooth annuli \mathcal{DA} and the rest $n\mathcal{DA}$. Note that each patch in $n\mathcal{DA}$ has both negative index and non-positive Euler characteristic.

The quantity $\langle [S], [\Gamma] \rangle$ can be computed as the signed number of intersections between S and Γ , or alternatively as the signed count of cusps over all patches of S with respect to $t = S \cap B^u$. Since the signed number of cusps of annuli and disks is 0, we have

$$2\chi(S) = -\langle [S], [\Gamma] \rangle = \sum_{n\mathcal{DA}} -\text{algcusps}(D),$$

where $\text{algcusps}(D)$ denotes the signed number of cusps in ∂D and the sum is over patches that are not topological disks or smooth annuli. Subtracting this expression for $2\chi(S)$ from the expression in terms of the indices of patches provided above gives

$$0 = \sum_{\mathcal{DA}} \text{index}(C) + \sum_{n\mathcal{DA}} (\text{index}(D) + \text{algcusps}(D)).$$

We claim that each term of the above sum is nonpositive. First, each term of the first sum is nonpositive, since each patch has nonpositive index following our assumption that there are no patches which are nullgons or monogons. For the second sum, the term for D in $n\mathcal{DA}$ is equal to $2\chi(D) + (\text{algcusps}(D) - \text{cusps}(D))$, which is nonpositive since $\chi(D) \leq 0$ and $\text{algcusps}(D) \leq \text{cusps}(D)$. Hence, each term of the sums is nonpositive and so it must be that each term is actually 0. We conclude that each patch in \mathcal{DA} is an annulus or a bigon with cusps whose signs cancel, and that each patch in $n\mathcal{DA}$ is a once-punctured disk with only positive cusps.

The same argument applies to $t^s = S \cap B^s$, and so by Theorem 5.13 we conclude that, after an isotopy, S is carried by the partial branched surface $\tau^{(2)} \cap \overset{\circ}{M}$. In partic-

ular, $S \cap \mathring{M}$ is carried by $\tau^{(2)}$ and under the canonical isomorphism $H_2(\mathring{M}, \partial\mathring{M}) = H_2(M, \partial M)$, it represents the class α . Hence, $\alpha \in \text{cone}_2(\tau)$ and the proof is complete. ■

5.3. Dimension and fiberedness of cones via the veering polynomial

We begin with the following lemma showing that P_Φ determines the cone of homology directions of τ .

Lemma 5.14. *For a veering triangulation τ , the cone of homology directions $\text{cone}_1(\Gamma)$ is generated by $i_*(\text{supp}(P_\Phi))$, where $i_*: H_1(\Phi; \mathbb{R}) \rightarrow H_1(M; \mathbb{R})$ is induced by the inclusion of the flow graph into M .*

Proof. By Theorem 5.1, $\text{cone}_1(\Gamma) = \text{cone}_1(\Phi)$. Since an arbitrary directed cycle in Φ is a sum of simple directed cycles, it suffices to know that the homology class of every simple directed cycles of Φ appears in the support of the Perron polynomial P_Φ . But this follows immediately from equation (6). ■

The following theorem summarizes the main results of this section and gives a characterization of fibered faces in the spirit of Fried's criterion.

Theorem 5.15. *Let τ be a veering triangulation. Then $\text{cone}_2(\tau)$ is equal to the subset of $H_2(M, \partial M)$ on which $x = -e_\tau$, which is the cone over a face \mathbf{F}_τ of $B_x(M)$. The codimension of $\mathbb{R}_+\mathbf{F}_\tau$ is equal to the dimension of the largest linear subspace contained in $\text{cone}_1(\Gamma)$.*

Furthermore, the following are equivalent:

- (i) $i_*(\text{supp}(P_{\Phi_\tau}))$ lies in an open half-space of $H_1(M; \mathbb{R})$,
- (ii) there exists $\eta \in H^1(M)$ with $\eta([\gamma]) > 0$ for each closed τ -transversal γ ,
- (iii) τ is layered,
- (iv) \mathbf{F}_τ is a fibered face.

Proof. Theorem 5.12 gives that $\text{cone}_2(\tau)$ equals the cone over a face \mathbf{F}_τ of $B_x(M)$. The statement about its codimension follows from the discussion at the beginning of Section 5.1 (or see [18, Section 1.2, Fact 10]) after we recall that the dual cone of $\text{cone}_2(\tau)$ equals $\text{cone}_1(\Gamma)$ by Proposition 5.11.

It remains to show the equivalence of the conditions.

For (i) \Leftrightarrow (ii), since $i_*(\text{supp}(P_{\Phi_\tau}))$ lies in an open half-space of $H_1(M; \mathbb{R})$, there is an $\eta \in H^1(M)$ that is positive on all directed cycles in $i_*(\text{supp}(P_{\Phi_\tau}))$. In particular, no such directed cycle is 0 in $H_1(M; \mathbb{R})$. Hence, if γ is a closed τ -transversal, then $[\gamma] \in H_1(M; \mathbb{R})$ is in the positive span of directed cycles in $i_*(\text{supp}(P_{\Phi_\tau}))$ by Lemma 5.14 and so $\eta([\gamma]) > 0$. For the converse, if η is positive on closed transversals, then it is positive on $i_*(\text{supp}(P_{\Phi_\tau}))$. This implies that $i_*(\text{supp}(P_{\Phi_\tau}))$ lies in an open half-space.

The implication (iii) \Rightarrow (ii) is immediate because a fiber of a layered veering triangulation positively intersects every closed transversal.

Next we turn to (ii) \Rightarrow (iii). Let $\eta \in H^1(M)$ be positive on closed transversals of τ . Then its pullback to the dual graph Γ is positive on directed cycles and so by Lemma 5.10 it is represented by a positive cocycle m on the edges of Γ . After perturbing slightly, we can assume that $m: E(\Gamma) \rightarrow \mathbb{Q}_+$ and so for some $n \geq 0$, $n \cdot m$ assigns a positive integer to each edge of Γ . As in the proof of Proposition 5.11, this implies that $n \cdot \eta$ is presented by a surface S *fully carried* by $\tau^{(2)}$, meaning that S traverses each face of τ . From this, it follows that the components of $M \setminus S$ are I -bundles and that τ is a layered triangulation on any component of S .

Finally, (iii) \Rightarrow (iv) is clear because a “layer” of a layered triangulation is a fiber of M . The reverse implication is more difficult and reserved for Proposition 5.16. ■

The following proposition completes the proof of Theorem 5.15. In short, it states that the only veering triangulations that carry fibers are layered. It can be thought of as a combinatorial version of [14, Proposition 4.5].

Proposition 5.16. *Let τ be a veering triangulation of M whose associated face $\mathbf{F}_\tau \subset B_x(M)$ is fibered. Then τ is layered.*

The proof is inspired by Agol’s proof of virtual fibering [1]. Recall that $G = \pi_1(M)$ is virtually special by Wise [46, Theorem 17.14] (see also Cooper–Futer [12, Theorem 1.4] and Groves–Manning [22, Theorem A]) and so it is virtually RFRS by Agol [1, Corollary 2.3]. This means that there exists a finite index subgroup $G_0 \leq G$ and subgroups $G_0 \supset G_1 \supset \cdots$ such that

- (1) $\bigcap_i G_i = 1$,
- (2) G_i is a normal, finite index subgroup of G_0 ,
- (3) for each i , the map $G_i \rightarrow G_i/G_{i+1}$ factors through $G_i \rightarrow H_1(G_i; \mathbb{Z})/\text{torsion}$.

Such a chain of subgroups is called an *RFRS tower*.

Proof of Proposition 5.16. Let $G = \pi_1(M)$ and let $G \geq G_0 \supset G_1 \supset \cdots$ be an RFRS tower for the finite index subgroup G_0 . For each i , let \tilde{M}_i be the corresponding cover of M .

Since fibers are the unique taut surfaces in their homology class, τ carries all the fibers in the cone $\mathbb{R}_+ \mathbf{F}_\tau$ by Theorem 5.12. Let $\tilde{\tau}_i$ be the lifted veering triangulation of \tilde{M}_i . Since $\tilde{\tau}_i$ carries the lifted fibers from τ , we see that each cone $\mathbb{R}_+ \mathbf{F}_i := \mathbb{R}_+ \mathbf{F}_{\tilde{\tau}_i}$ is fibered and hence top-dimensional.

Suppose that τ is not layered; hence, no surface is fully carried by $\tau^{(2)}$. We choose a (multiple of a) fiber S that is maximal with respect to the weights placed on faces of τ . That is, S is carried by τ and traverses every face that is traversed by any carried surface. Such a surface can be constructed by summing weights of finitely many fibers that traverse the faces traversed by some carried surface. Of course, the complete preimage \tilde{S}_i of S in \tilde{M}_i is a (multiple of a) fiber carried by $\tilde{\tau}_i$ whose weights on faces are lifted from those of S . We claim that \tilde{S}_i also has the weight maximality property for $\tilde{\tau}_i$:

Claim 1. *The carried surface \tilde{S}_i traverses every face of $\tilde{\tau}_i$ that is traversed by any carried surface. In particular, each $\tilde{\tau}_i$ is also not layered for $i \geq 0$.*

Proof of Claim 1. Suppose that Z is a carried surface that traverses a face not traversed by \tilde{S}_i . Then the system of weights on faces of $\tilde{\tau}_i$ corresponding to Z pushes down to τ to give a system of weights (satisfying the matching conditions) and determines a surface carried by τ that traverses faces not traversed by S , a contradiction. ■

Since $\tilde{\tau}_i$ is not layered, it has dual cycles that do not intersect \tilde{S}_i ; see (ii) \Rightarrow (iii) in the proof of Theorem 5.15. (Here, we set $\tilde{M}_{-1} = M$, $\tilde{\tau}_{-1} = \tau$, and $\tilde{S}_{-1} = S$.) These dual cycles lie in the “guts” of \tilde{S}_i defined as follows: Let $\mathcal{G}(\tilde{S}_i)$ be the open region obtained by taking the open tetrahedra of $\tilde{\tau}_i$ together with the open faces of $\tilde{\tau}_i^{(2)}$ that are *not* traversed by \tilde{S}_i . Any dual cycle not crossing \tilde{S}_i lives in $\mathcal{G}(\tilde{S}_i)$ by construction.

Claim 2. *The image of $H^1(\tilde{M}_i) \rightarrow H^1(\mathcal{G}(\tilde{S}_i))$ is trivial for $i \geq -1$.*

Proof of Claim 2. Suppose that some nontrivial α is in the image of the map. Since $\mathbb{R}_+ \mathbf{F}_i$ is top-dimensional in $H^1(\tilde{M}_i) = H_2(\tilde{M}_i, \partial\tilde{M}_i)$, i.e. it has nonempty interior, there is a $\beta \in \mathbb{R}_+ \mathbf{F}_i$ such that $\alpha + \beta \in \mathbb{R}_+ \mathbf{F}_i$. But then $\alpha + \beta$ is represented by a surface Z carried by $\tilde{\tau}_i$ that has nonzero intersection number with an oriented loop in $\mathcal{G}(\tilde{S}_i)$. This contradicts Claim 1, which implies that Z is disjoint from $\mathcal{G}(\tilde{S}_i)$. ■

Now by Claim 2, each component of $\mathcal{G}(\tilde{S}_0)$ lifts homeomorphically to \tilde{M}_1 . Indeed, this amounts to the claim that any loop in $\mathcal{G}(\tilde{S}_0)$ is trivial under the homomorphism $G_0 \rightarrow H_1(\tilde{M}_0)/\text{torsion}$. By definition of \tilde{S}_i , all the lifts of $\mathcal{G}(\tilde{S}_0)$ to \tilde{M}_1 live in $G(\tilde{S}_1)$. Then again applying Claim 2, we see that $\mathcal{G}(\tilde{S}_0)$ lifts to \tilde{M}_2 . Continuing in this way, we find that $\mathcal{G}(\tilde{S}_0)$ lifts to each \tilde{M}_i and hence for each component C of $\mathcal{G}(\tilde{S}_0)$, $\pi_1(C) \subset \bigcap_i G_i$. Hence, each component of $\mathcal{G}(\tilde{S}_0)$ has trivial fundamental group. This however contradicts the fact that if τ (and hence $\tilde{\tau}_0$) is not layered, then $\mathcal{G}(\tilde{S}_0)$ contains dual cycles, which are homotopically essential by a combinatorial version of a theorem of Novikov; see [43, Theorem 3.2] or the references found therein. The proof of Proposition 5.16 is complete. ■

6. Relating the polynomials V_τ and Θ_τ

In this section, we establish a precise version of the identity

$$V_\tau = \Theta_\tau \cdot \prod (1 \pm g_i),$$

where the $g_i \in G$ are represented by certain directed cycles in the dual graph $\Gamma = \Gamma_\tau$. We will see in Section 7 that this is a generalization of McMullen’s Determinant Formula from the fibered setting.

After some definitions we will state Theorem 6.1, which is the main result of this section.

6.1. The AB-polynomial

As before, we let \tilde{M}^{ab} denote the universal free abelian cover of M .

Recall the definitions of $L: \mathbb{Z}[G]^E \rightarrow \mathbb{Z}[G]^E$ and $L^\Delta: \mathbb{Z}[G]^F \rightarrow \mathbb{Z}[G]^E$ and their respective cokernels, $\mathcal{E}(\tilde{\tau})$ and $\mathcal{E}^\Delta(\tilde{\tau})$ (equations (2) and (4)). To relate the two we recall from Lemma 3.2 the fact that, if f is a bottom face of a tetrahedron t and e its bottom edge, then there is a unique top face f' of t such that $L^\Delta(f) + L^\Delta(f') = L(e)$. This face f' is characterized by the fact that the turn in the dual graph Γ at t given by the pair (f, f') is *anti-branching* (see Lemma 4.5).

The correspondence $f \mapsto f'$ defines a map $A: F \rightarrow F$, which after lifting to \tilde{M}^{ab} induces a $\mathbb{Z}[G]$ -module homomorphism $A: \mathbb{Z}[G]^F \rightarrow \mathbb{Z}[G]^F$. The correspondence $f \mapsto e$ sending a face to its bottom edge gives a 2-1 map $\varepsilon: F \rightarrow E$, which also extends to $\varepsilon: \mathbb{Z}[G]^F \rightarrow \mathbb{Z}[G]^E$. Defining $L^{\mathcal{AB}} = I + A$, we can express the above identity as

$$L^\Delta \circ L^{\mathcal{AB}} = L \circ \varepsilon. \quad (8)$$

Now define the AB -module $\mathcal{AB}(\tilde{\tau})$ as the cokernel of $L^{\mathcal{AB}}$, that is, via

$$\mathbb{Z}[G]^F \xrightarrow{L^{\mathcal{AB}}} \mathbb{Z}[G]^F \rightarrow \mathcal{AB}(\tilde{\tau}) \rightarrow 0. \quad (9)$$

As $L^{\mathcal{AB}}$ is a square matrix, its fitting ideal is principally generated by $\det(L^{\mathcal{AB}})$ and we define the AB -polynomial of τ by

$$V^{\mathcal{AB}} = V_\tau^{\mathcal{AB}} = \det(L^{\mathcal{AB}}) \in \mathbb{Z}[G].$$

We will see in Lemma 6.2 that $V^{\mathcal{AB}}$ has the form $\prod (1 \pm g_i)$ over certain $g_i \in G$. Our main result will then be:

Theorem 6.1. *Suppose that $\text{rank}(H_1(M)) > 1$. Then*

$$V_\tau = V^{\mathcal{AB}} \cdot \Theta_\tau$$

up to multiplication by a unit in $\mathbb{Z}[G]$.

When the rank of homology is 1, the conclusion holds up to multiplying by $1 \pm t$, where t generates G . See Remark 6.18.

Let us give a short sketch of the proof.

We first introduce AB -cycles, which are just the cycles of the permutation A , and give in Lemma 6.2 a factorization formula for $V^{\mathcal{AB}}$.

Proposition 6.6 gives the “easy” direction of the theorem, namely $V \mid (V^{\mathcal{AB}} \cdot \Theta_\tau)$, which follows from standard facts about Fitting ideals (but we supply an explicit proof).

In Lemma 6.8 we give an alternative presentation of \mathcal{E}^Δ that allows us to express Θ_τ as the gcd of minors of a matrix involving L and selected columns of L^Δ .

Lemma 6.9, the AB -cycle equation, is an identity combining L , L^Δ and the AB -cycles. In Proposition 6.12 we use this equation to enable column operations that express V_τ as one of the minors of the matrix in Lemma 6.8. This gives us, in Corollary 6.13, a result of the form

$$(\Theta_\tau \cdot \prod (1 \pm g_i)) \mid V_\tau,$$

where the product is over a restricted collection of cycles called a reducing family.

We then introduce *AB-chains* which are cycles formed of segments of *AB*-cycles in a restricted manner. We find in Lemma 6.14 and Corollary 6.15 that

$$(V^{\mathcal{AB}} \cdot \Theta_\tau) \mid (1 \pm z)V_\tau,$$

where z is the image in G of such a cycle.

In Lemma 6.16 we show that the homology classes of *AB*-chains and *AB*-cycles suffice to generate G and use this to show that, by applying Corollaries 6.13 and 6.15 over all *AB*-cycles and *AB*-chains, we can establish

$$(V^{\mathcal{AB}} \cdot \Theta_\tau) \mid V_\tau,$$

which completes the proof.

6.2. Cycles and the factorization of $V^{\mathcal{AB}}$

The map $A: F \rightarrow F$ is a permutation and we call its cycles the *AB*-cycles of τ . Recall that *dual cycles* and *dual paths* are directed cycles and paths respectively in the dual graph Γ . Each *AB*-cycle c determines a unique dual path whose Γ -edges correspond to the τ -faces in c . We will often speak of these two objects interchangeably. In this language, an *AB*-cycle is a directed cycle in Γ that makes only *anti-branching* turns in the sense of Section 4.3.

A τ -face $f \in F$ is dual to a Γ -edge pointing into the tetrahedron for which f is a bottom face, and we can label this tetrahedron (which is dual to a Γ -vertex) by its bottom τ -edge which is exactly $\varepsilon(f)$. We say that an *AB*-cycle *passes through the faces* f, Af, A^2f, \dots and, with a slight abuse of notation, *encounters the τ -edges* $\varepsilon(A^i f)$.

Note that a τ -face f belongs to a unique *AB*-cycle, while each τ -edge is encountered twice by *AB*-cycles (possibly the same one). This corresponds to the fact that the dual graph Γ is 4-valent.

Each *AB*-cycle c determines a homology class in M and hence an image $g = [c] \in G$. Let $k = k(c)$ denote the length of the cycle, or the number of $f \in F$ which it passes through. This structure allows us to compute $V^{\mathcal{AB}}$:

Lemma 6.2. *Let c_1, \dots, c_n denote the *AB*-cycles of τ , k_i their lengths, and $g_i = [c_i]$ their images in G . Then we may identify*

$$\mathcal{AB}(\tilde{\tau}) \cong \bigoplus_{i=1}^n \frac{\mathbb{Z}[G]}{(1 + (-1)^{k_i+1} g_i)},$$

where the map from the i th copy of $\mathbb{Z}[G]$ into $\mathcal{AB}(\tilde{\tau})$ maps 1 to any (fixed) face in c_i . In particular,

$$V^{\mathcal{AB}} = \prod_{i=1}^n (1 + (-1)^{k_i+1} g_i)$$

up to a unit factor in $\mathbb{Z}[G]$.

Remark 6.3. We will often write the expression $1 + (-1)^{k_i+1} g_i$ as $1 \pm g_i$, for brevity.

Proof of Lemma 6.2. In this proof we abbreviate $\mathcal{AB} := \mathcal{AB}(\tilde{\tau})$. Fix a face f_i for each cycle c_i and define a map

$$\gamma: \mathbb{Z}[G]^n \rightarrow \mathbb{Z}[G]^F$$

as in the statement, taking the generator of the i -th copy of $\mathbb{Z}[G]$ to f_i . Let $\hat{\gamma}$ denote γ composed with the surjection $\mathbb{Z}[G]^F \rightarrow \mathcal{AB}$. The map $\hat{\gamma}$ is surjective because, in \mathcal{AB} , the identity $f = -Af$ holds (since $L^{\mathcal{AB}}(f) = f + Af$ is a relation), so one face from each cycle suffices to generate the rest.

Lifting the permutation A to the faces in the universal cover, it acts as a translation on the lift of each cycle. Thus if we denote the distinguished lift of f_i by f_i again, we have the A -orbit $f_i \rightarrow Af_i \rightarrow \cdots \rightarrow A^{k_i} f_i$, where $A^{k_i} f_i$ is just $g_i \cdot f_i$, the translate of the lift f_i by the element of the deck group associated to the cycle.

Since in \mathcal{AB} we have $A^j f_i = -A^{j+1} f_i$, we obtain the relation $g_i \cdot f_i = (-1)^{k_i} f_i$. Hence $\hat{\gamma}$ factors through the quotient $\bigoplus_{i=1}^n \mathbb{Z}[G]/(1 \pm g_i)$. An inverse map is easily constructed by appealing to the same relations, and so we obtain the desired description of \mathcal{AB} .

This gives us a new presentation of \mathcal{AB} , namely $\mathbb{Z}[G]^n \rightarrow \mathbb{Z}[G]^n \rightarrow \mathcal{AB}$, where the matrix is diagonal with $1 \pm g_i$ along the diagonals. Since the Fitting ideal of \mathcal{AB} is independent of presentation, the expression for $V^{\mathcal{AB}}$ follows. ■

6.3. Diagrams and presentations

It is helpful to organize our maps in the following commutative diagram:

Proposition 6.4. *We have the following commutative diagram with exact rows and columns:*

$$\begin{array}{ccccccc}
 \mathbb{Z}[G]^F & \xrightarrow{\varepsilon} & \mathbb{Z}[G]^E & \longrightarrow & 0 & & \\
 L^{\mathcal{AB}} \downarrow & & L \downarrow & & & & \\
 \mathbb{Z}[G]^F & \xrightarrow{L^\Delta} & \mathbb{Z}[G]^E & \longrightarrow & \mathcal{E}^\Delta(\tilde{\tau}) & \longrightarrow & 0 \\
 \downarrow & & \downarrow & & \parallel & & \\
 \mathcal{AB}(\tilde{\tau}) & \xrightarrow{\overline{L^\Delta}} & \mathcal{E}(\tilde{\tau}) & \longrightarrow & \mathcal{E}^\Delta(\tilde{\tau}) & \longrightarrow & 0 \\
 \downarrow & & \downarrow & & & & \\
 0 & & 0 & & & &
 \end{array}$$

Proof. Surjectivity of ε is immediate from the definition, and commutativity of the top square is equation (8).

Exactness of the middle row is the definition of \mathcal{E}^Δ , and exactness of the first two columns is the definition of \mathcal{AB} and \mathcal{E} .

Commutativity of the bottom left square is the definition of $\overline{L^\Delta}$.

The definition of the arrow $\mathcal{E}(\tilde{\tau}) \rightarrow \mathcal{E}^\Delta(\tilde{\tau})$, and exactness of the bottom row, follow from a short diagram chase. ■

6.3.1. *Presenting \mathcal{E} .* We can now consider an alternative presentation of $\mathcal{E} = \mathcal{E}(\tilde{\tau})$:

Proposition 6.5. *Consider the module homomorphism*

$$\mathcal{L}: \mathbb{Z}[G]^F \oplus \mathbb{Z}[G]^F \rightarrow \mathbb{Z}[G]^F \oplus \mathbb{Z}[G]^E$$

given in block form by the $(F + E) \times (2F)$ matrix

$$\mathcal{L} = \left[\begin{array}{c|c} L^{\mathcal{AB}} & I \\ \hline 0 & L^\Delta \end{array} \right].$$

Then $\mathcal{E}(\tilde{\tau})$ is isomorphic to the cokernel of \mathcal{L} .

Proof. Denote the cokernel of \mathcal{L} by \mathcal{E}' and note \mathcal{E}' is the free module $\mathbb{Z}[G]^F \oplus \mathbb{Z}[G]^E$ modulo two relations for each face f :

$$L^{\mathcal{AB}}(f) = f + f' = 0, \quad f + L^\Delta(f) = 0.$$

Hence, if $e = \varepsilon(f)$, then

$$\begin{aligned} L(e) &= L^\Delta(f) + L^\Delta(f') \\ &= (f + L^\Delta(f)) + (f' + L^\Delta(f')) - L^{\mathcal{AB}}(f) \in \text{Im}(\mathcal{L}), \end{aligned}$$

and so the factor inclusion $\mathbb{Z}[G]^E \rightarrow \mathbb{Z}[G]^F \oplus \mathbb{Z}[G]^E$ descends to a module homomorphism $\mathcal{E} \rightarrow \mathcal{E}'$.

An inverse homomorphism is induced by first defining a homomorphism $\mathbb{Z}[G]^F \oplus \mathbb{Z}[G]^E \rightarrow \mathbb{Z}[G]^E$ by

$$(f, e) \mapsto e - L^\Delta(f).$$

Since this maps the first relation on \mathcal{E}' to a tetrahedron relation of \mathcal{E} and the second relation on \mathcal{E}' to 0, it induces a map $\mathcal{E}' \rightarrow \mathcal{E}$. This is easily seen to be the inverse of the one defined above. ■

From this we obtain one direction of our main theorem:

Proposition 6.6.

$$V_\tau \mid (V^{\mathcal{AB}} \cdot \Theta_\tau).$$

Proof. The Fitting ideal (V_τ) of \mathcal{E} is independent of the presentation of \mathcal{E} . Hence, we know that (V_τ) is also the ideal generated by the minors of \mathcal{L} of size $|F| + |E|$. In particular, V_τ divides the product of $\det(L^{\mathcal{AB}})$ with each minor of L^Δ of size $|E|$. Hence $V_\tau \mid (V^{\mathcal{AB}} \cdot \Theta_\tau)$. ■

Remark 6.7. This is actually a general fact: If $P \rightarrow Q \rightarrow R \rightarrow 0$ is an exact sequence of finitely generated modules then $\text{Fitt}(P) \cdot \text{Fitt}(R) \subset \text{Fitt}(Q)$. For example, this can be derived from [38, Chapter 3, Exercise 2]. We can apply this to the sequence $\mathcal{AB} \rightarrow \mathcal{E} \rightarrow \mathcal{E}^\Delta \rightarrow 0$ from Proposition 6.4. We chose to give an explicit proof for completeness.

6.3.2. *Presenting \mathcal{E}^Δ .* The following presentation of $\mathcal{E}^\Delta(\tilde{\tau})$ gives us a new description of Θ_τ .

Lemma 6.8. *Let c_1, \dots, c_n be the collection of all AB -cycles and let f_i be a face of c_i . Then $\mathcal{E}^\Delta(\tilde{\tau})$ is presented by the $|E| \times (n + |E|)$ matrix*

$$[L^\Delta(f_1) \mid \dots \mid L^\Delta(f_n) \mid L].$$

Hence, Θ_τ is the gcd of the $|E| \times |E|$ minors of this matrix.

Here we have used the following notation: If A and B are matrices with the same number of rows, then $[A|B]$ denotes the matrix whose columns are the columns of A followed by the columns of B . If A or B is a vector, then we interpret it as a column vector.

Proof of Lemma 6.8. This matrix defines a map $M: \mathbb{Z}[G]^{|E|+n} \rightarrow \mathbb{Z}[G]^E$ and it suffices to show that its image equals the image of L^Δ .

Returning to the diagram of Proposition 6.4, since ε is surjective, the image of L must be contained in the image of L^Δ . Thus $\text{Im}(M) \subset \text{Im}(L^\Delta)$.

Now for any $f \in F$ we have $L^\Delta(f + Af) = L(\varepsilon(f))$. Hence $L^\Delta(Af)$ is contained in the span of $\text{Im}(L)$ and $L^\Delta(f)$. Repeating inductively, the entire cycle of f_1 has L^Δ -image in the span of $\text{Im}(L)$ and $L^\Delta(f_1)$. Applying this to all the cycles we see that the image of L^Δ is contained in the image of M .

The statement about Θ_τ follows again from the fact that the Fitting ideal is independent of presentation. ■

6.4. Factoring V_τ

We next turn to describing our primary means of factoring the veering polynomial.

6.4.1. *The AB -cycle equation.* The cycle structure of the permutation A gives rise to a useful identity relating L and L^Δ . This uses a combination of the discussion in the proof of Lemma 6.2 and the identity (8).

Lemma 6.9 (*AB -cycle equation*). *Let c be an AB -cycle, $e(c)$ the set of τ -edges encountered by c , f one of the τ -faces of c , and g the class $[c]$ in G . Then the following identity holds:*

$$\sum_{e \in e(c)} \alpha_e \cdot L(e) = (1 \pm g) \cdot L^\Delta(f), \quad (10)$$

where α_e is a unit in $\mathbb{Z}[G]$ if e appears once along c , and a sum of two units if e appears twice.

Proof. We identify f with its selected lift in \tilde{M}^{ab} , and enumerate a lift of the cycle as $f_i = A^i f$, $i = 0, \dots, k$, where k is the period of c , so that $f_k = A^k f = gf$. Then

each $\tilde{\tau}$ -edge $\varepsilon(f_i)$ can be written as $a_i e_i$, where e_i is one of the elements of $e(c)$ (again identified with its selected lift in \tilde{M}^{ab}) and $a_i \in G$. The identity $L \circ \varepsilon = L^\Delta \circ (I + A)$ gives us

$$a_i L(e_i) = L^\Delta(f_i) + L^\Delta(f_{i+1}).$$

If we combine these in an alternating sum, we therefore obtain, after cancellations,

$$\sum_{i=0}^{k-1} (-1)^i a_i L(e_i) = L^\Delta(f_0) + (-1)^{k-1} L^\Delta(f_k).$$

Now since $f_k = gf$, the right-hand side is $(1 \pm g)L^\Delta(f)$. On the left side, if e appears once as e_i then we have a term of the form $\pm a_i L(e)$, and if it appears twice as e_i and e_j we combine two terms to get $(\pm a_i \pm a_j)L(e)$. This gives the desired statement. ■

As an immediate corollary we have:

Corollary 6.10. *Let c be an AB -cycle and e an edge appearing exactly once in c . Then*

$$L(e) = u(1 \pm g) \cdot L^\Delta(f) + \sum_{e' \in e(c) \setminus \{e\}} \beta_{e'} \cdot L(e'), \quad (11)$$

where $g = [c] \in G$, f is a face of c , u is a unit in $\mathbb{Z}[G]$, and each $\beta_{e'}$ is a sum of at most two units in $\mathbb{Z}[G]$.

6.4.2. Factoring with reducing families of cycles. A collection \mathcal{C} of AB -cycles forms a *reducing family of cycles* if it can be ordered $\mathcal{C} = \{c_1, \dots, c_k\}$ so that for each $1 \leq i \leq k$, there is a τ -edge e_i through which c_i passes exactly once and which is not visited by any c_j for $j > i$. Such an ordering is not necessarily unique. However, given such an ordering, we call it the *preferred order*. We call the τ -edges e_i the *distinguished edges* of the cycles in the family (these too may not be unique, but we can make an arbitrary choice). In the results and proofs to follow, given a reducing family $\{c_1, \dots, c_k\}$ we will often write $g_i = [c_i] \in G$.

We first show that reducing families are not hard to find:

Lemma 6.11 (Finding families). *Any proper subset of AB -cycles is a reducing family.*

Proof. For the argument, we think of AB -cycles as directed cycles in the dual graph Γ , so that τ -edges of a cycle correspond to Γ -vertices.

A collection of directed cycles in Γ that crosses every vertex either zero or two times must in fact cross every vertex, since Γ is strongly connected by Lemma 5.8. Thus, a proper subset \mathcal{C} of AB -cycles must cross some vertex v_1 of Γ exactly once. Let c_1 be the cycle that contains v_1 . Now apply the same argument to $\mathcal{C} \setminus \{c_1\}$ and continue inductively. ■

Reducing families provide factorizations of V_τ of the following type:

Proposition 6.12. *For any reducing family of cycles \mathcal{C} , we have the factorization (up to a unit $\pm h \in \mathbb{Z}[G]$)*

$$V_\tau = \left(\prod_{c_i \in \mathcal{C}} (1 \pm g_i) \right) \cdot \det[L^\Delta(f_1) \mid \cdots \mid L^\Delta(f_r) \mid L'],$$

where f_i is the (unique) face in the cycle $c_i \in \mathcal{C}$ such that $\varepsilon(f_i)$ is the distinguished edge e_i and L' is the matrix obtained from L by removing the columns corresponding to the distinguished edges of \mathcal{C} .

Proof. Let $\mathcal{C} = \{c_1, \dots, c_r\}$ be written in preferred order, and reorder the columns of L so that the i th column (for $i \leq i \leq r$) corresponds to the distinguished edge e_i of c_i . Since e_1 appears exactly once in c_1 , Corollary 6.10 of the AB -cycle equation implies that we can perform a column replacement on the first column of L to obtain (up to units)

$$\begin{aligned} V_\tau &= \det([L(e_1) \mid L'_1]) = \det([(1 \pm g_1)L^\Delta(f_1) \mid L'_1]) \\ &= (1 \pm g_1) \cdot \det([L^\Delta(f_1) \mid L'_1]), \end{aligned}$$

where L'_1 denotes L with its first column removed.

Now since \mathcal{C} is a reducing family, the cycle c_2 does not encounter e_1 . This is to say that equation (10) for the cycle c_2 does not involve $L(e_1)$. Hence, again using that e_2 appears exactly once in c_2 , we can apply a column replacement on the second column of $[L^\Delta(f_1) \mid L'_1]$ to obtain

$$V_\tau = (1 \pm g_1)(1 \pm g_2) \cdot \det([L^\Delta(f_1) \mid L^\Delta(f_1) \mid L'_{1,2}],$$

where $L'_{1,2}$ is obtained by removing the first two columns from L . Continuing in this manner proves the proposition. \blacksquare

The proposition along with Lemma 6.8 gives the following:

Corollary 6.13. *For each reducing family of cycles \mathcal{C} ,*

$$\left(\Theta_\tau \cdot \prod_{c_i \in \mathcal{C}} (1 \pm g_i) \right) \mid V_\tau.$$

Proof. By Proposition 6.12, $V' = V_\tau / \prod_{c_i \in \mathcal{C}} (1 \pm g_i)$ defines an element of $\mathbb{Z}[G]$ up to a unit, and it suffices to show $\Theta_\tau \mid V'$. For this, note that the expression for V' given by Proposition 6.12 is a minor appearing in the definition of Θ_τ from Lemma 6.8. \blacksquare

6.4.3. Factoring with AB -chains. We now consider a class of dual cycles which will be useful for factoring V_τ .

First, define the AB -length of a dual cycle c to be its number of branching turns. For any dual cycle c of positive AB -length there is a decomposition of c as a concatenation of dual paths $c = (p_1, \dots, p_k)$ such that each anti-branching turn of c is interior to some p_i , and where k is the AB -length of c . This decomposition is unique up to cyclic permutation and we call it the AB -decomposition of c . We say such a c is an AB -chain if it has the following two properties with respect to its AB -decomposition $c = (p_1, \dots, p_k)$:

- (1) (*AB-simple*) Each p_i is a proper subpath of some AB -cycle c_i such that $c_i \neq c_j$ for $i \neq j$.
- (2) (*Endpoint simple*) Each endpoint of each p_i is visited exactly once by c .

Note that if c is a simple (i.e. embedded) dual cycle with AB -length $k > 0$ then c is automatically endpoint simple, although we allow nonsimple AB -chains. The endpoints of the paths p_1, \dots, p_k in the AB -decomposition correspond to the branching turns of c , and we call the corresponding edges of τ the *branching τ -edges* of c . If c_i is the AB -cycle containing p_i , we say c uses the AB -cycles c_1, \dots, c_k .

Note that an AB -cycle has AB -length 0 and so the above definitions do not apply.

Lemma 6.16 below will furnish many useful AB -chains, but for now let us explain how these objects can be used to factor (a multiple of) V_τ .

Lemma 6.14. *Suppose that c is an AB -chain which has branching τ -edges e_1, \dots, e_k and which uses distinct AB -cycles c_1, \dots, c_k . Let f_i be the face in c with $\varepsilon(f_i) = e_i$. Then, up to a unit in $\mathbb{Z}[G]$,*

$$(1 \pm z) \cdot V_\tau = \left(\prod_{1 \leq i \leq k} (1 \pm g_i) \right) \cdot \det[L^\Delta(f_1) \mid \dots \mid L^\Delta(f_k) \mid L'_c],$$

where $z = [c]$, $g_i = [c_i]$, and L'_c is the matrix obtained from L by removing the columns corresponding to e_1, \dots, e_k .

Proof. Let (p_1, \dots, p_k) be the AB -decomposition of c . We label τ -edges so that p_i is an oriented path from e_i to e_{i+1} with indices taken mod k . (Here, we are using the correspondence between vertices of Γ and edges of τ .)

Lift c to \tilde{M}^{ab} so that it begins with a fixed lift of e_1 which we also denote e_1 . Then each p_i lifts to a path from $a_i e_i$ to $a_{i+1} e_{i+1}$, where a_i denotes the group element translating a fixed lift of e_i to the one encountered by our lift of c . Thus $a_1 = 1$ and the last τ -edge is $a_{k+1} e_1$, where a_{k+1} is exactly $z = [c]$.

First suppose that $k \geq 2$. (The case of $k = 1$ is easier and handled later.) Note that because c is AB -simple and $k \geq 2$, each e_i is encountered by two distinct AB -cycles c_i, c_{i+1} and therefore no single AB -cycle encounters e_i twice. Hence applying the AB -cycle equation (Lemma 6.9) to c_i we obtain

$$\pm a_i L(e_i) \pm a_{i+1} L(e_{i+1}) + \sum_{e'} \alpha_{e'} L(e') = h_i (1 \pm g_i) L^\Delta(f_i), \quad (12)$$

where $h_i \in G$ corrects for the difference between our lift of p_i here and the lift of c_i from Lemma 6.9. In equation (12) above, the sum on the left-hand side is over all e' in $e(c_i) \setminus \{e_i, e_{i+1}\}$, where $e(c_i)$ is the set of τ -edges encountered by c_i and each $\alpha_{e'}$ is a sum of at most two units in $\mathbb{Z}[G]$. For the last segment p_k the left-hand side becomes $\pm a_k L(e_k) \pm z L(e_1) + \sum_{e'} \alpha_{e'} L(e')$.

We claim that no edge encountered by c_i other than $\{e_i, e_{i+1}\}$ is a branching edge of c . In other words, for each $1 \leq i \leq k$ and for each $e' \in e(c_i) \setminus \{e_i, e_{i+1}\}$, $L(e')$ is a column of L'_c . Otherwise we would have $e' = e_j$ for some $j \in \{1, \dots, k\} \setminus \{i, i+1\}$.

This would force e_j to be part of distinct AB -cycles c_i, c_{j-1}, c_j , which is impossible as each τ -edge is encountered by at most two AB -cycles.

This enables us to make the following computation, in which the equalities are taken up to multiplication by units $\pm h$:

$$\begin{aligned} \left(\prod_{1 \leq i \leq k} (1 \pm g_i) \right) \cdot \det[L^\Delta(f_1) \mid \cdots \mid L^\Delta(f_k) \mid L'_c] \\ = \det[(1 \pm g_1)L^\Delta(f_1) \mid \cdots \mid (1 \pm g_k)L^\Delta(f_k) \mid L'_c] \\ = \det[\pm L(e_1) \pm a_2 L(e_2) \mid \cdots \mid \pm a_k L(e_k) \pm z L(e_1) \mid L'_c]. \end{aligned}$$

In the last equality we use the substitution of equation (12) in each column and then use the columns of L'_c to remove all the terms that do involve the branching τ -edges e_i .

Now adding (or subtracting) the i th column to (from) the $(i + 1)$ st column, starting at $i = 1$, the columns are replaced by $\pm L(e_1) \pm a_{i+1} L(e_{i+1})$ for $i = 1, \dots, k - 1$, and $\pm L(e_1) \pm z L(e_1)$ for the k th column. Thus we have

$$\begin{aligned} \det[\pm L(e_1) \pm a_2 L(e_2) \mid \cdots \mid \pm L(e_1) \pm z L(e_1) \mid L'_c] \\ = \det[\pm L(e_1) \pm a_2 L(e_2) \mid \cdots \mid \pm(1 \pm z) \cdot L(e_1) \mid L'_c] \\ = (1 \pm z) \cdot \det[\pm L(e_1) \pm a_2 L(e_2) \mid \cdots \mid L(e_1) \mid L'_c] \\ = (1 \pm z) \cdot \det[\pm a_2 L(e_2) \mid \cdots \mid L(e_1) \mid L'_c] \\ = (1 \pm z) \cdot V_\tau. \end{aligned}$$

In the penultimate line we used the k th column to remove the $L(e_1)$ terms from the previous columns. In the last line we note that the determinant is, up to unit multiples, equal to the determinant of L , or V_τ .

Finally, when $k = 1$ we have that p_1 is a proper subpath of c_1 which starts and ends at e_1 and forms the directed cycle c . In particular, c_1 encounters e_1 twice and so in place of equation (12) we have

$$\pm L(e_1) \pm z L(e_1) + \sum_{e'} \alpha_{e'} L(e') = h_1(1 \pm g_1)L^\Delta(f_1), \quad (13)$$

where notation is as before. Then proceeding exactly as above, we have

$$\begin{aligned} (1 \pm g_1) \cdot \det[L^\Delta(f_1) \mid L'_c] &= \det[(1 \pm g_1)L^\Delta(f_1) \mid L'_c] = \det[\pm L(e_1) \pm z L(e_1) \mid L'_c] \\ &= (1 \pm z) \cdot V_\tau, \end{aligned}$$

and the proof is complete. ■

As a corollary we obtain the following factorization:

Corollary 6.15. *For each AB -chain c with $z = [c]$,*

$$(V^{AB} \cdot \Theta_\tau) \mid (1 \pm z) \cdot V_\tau.$$

Proof. Let \mathcal{C}_c be the collection of AB -cycles not used by c . Since AB -chains are non-empty by definition, \mathcal{C}_c is a proper collection and hence a reducing family by Lemma 6.11. Note that \mathcal{C}_c may be empty if c uses all of the AB -cycles in Γ .

Each branching τ -edge of c is contained in two AB -cycles used by c . Since a τ -edge is encountered by at most two AB -cycles, the branching τ -edges are not encountered by any of the AB -cycles of \mathcal{C}_c .

Enumerate the branching τ -edges of c by e_1, \dots, e_k as above, and let e_{k+1}, \dots, e_n denote the distinguished τ -edges of the cycles of \mathcal{C}_c . Let f_i be the associated τ -faces as before.

Let L'' be the matrix obtained from L by removing the columns corresponding to e_1, \dots, e_n . We see that

$$\begin{aligned} (1 \pm z) \cdot V_\tau &= \left(\prod_{1 \leq i \leq k} (1 \pm g_i) \right) \cdot \det[L^\Delta(f_1) \mid \cdots \mid L^\Delta(f_k) \mid L'_c] \quad (\text{by Lemma 6.14}) \\ &= \left(\prod_{1 \leq i \leq k} (1 \pm g_i) \right) \cdot \left(\prod_{k+1 \leq j \leq n} (1 \pm g_j) \right) \cdot \det[L^\Delta(f_1) \mid \cdots \mid L^\Delta(f_n) \mid L''] \\ &= V^{AB} \cdot \det[L^\Delta(f_1) \mid \cdots \mid L^\Delta(f_n) \mid L''], \end{aligned}$$

The second equality above follows from techniques in the proof of Proposition 6.12. More specifically, we repeatedly apply Corollary 6.10 using the facts that \mathcal{C}_c is a reducing family and that no AB -cycle in \mathcal{C}_c encounters any of e_1, \dots, e_k .

As in the proof of Corollary 6.13,

$$[L^\Delta(f_1) \mid \cdots \mid L^\Delta(f_n) \mid L'']$$

is a minor of the presentation matrix for $\mathcal{E}^\Delta(\tilde{\tau})$ from Lemma 6.8. It follows that Θ_τ divides its determinant. \blacksquare

6.5. AB -chains and homology

Until now we have not proved the existence of any AB -chains, but in fact there are enough of them to generate the homology of Γ :

Lemma 6.16. *The AB -chains and AB -cycles together generate $H_1(\Gamma)$.*

Proof. We first show that any dual cycle c is homologous to an integer linear combination of AB -cycles and AB -simple dual cycles. For this, we induct on the AB -length of c . If c has AB -length 0 then c is an AB -cycle; also, if c has AB -length 1 then c is easily seen to be an AB -chain, so is AB -simple.

Now suppose c has AB -length $k > 1$ with AB -decomposition $c = (p_1, \dots, p_k)$, and suppose p_1 and p_j are part of the same AB -cycle g . Now:

- Let g_1 be the directed subpath of g from the initial point of p_1 to the terminal point of p_j . Note that this includes the edges p_1 and p_j .
- Let g_2 be the directed subpath of g from the initial point of p_j to the terminal point of p_1 .
- Let $G_1 = (g_1, p_{j+1}, \dots, p_k)$.
- Let $G_2 = (g_2, p_2, \dots, p_{j-1})$.




Fig. 14. Decomposing the dual cycle c as $c = G_1 + G_2 - g$ as in the proof of Lemma 6.16, where we have drawn the cycles so that orientations are clockwise in the page.

Then $c = G_1 + G_2 - g$ in $H_1(\Gamma)$ (see Figure 14). Note that G_1, G_2 both have smaller AB -length than c , as does g whose AB -length is 0. By induction, G_1 and G_2 are each homologous to a \mathbb{Z} -linear combination of AB -cycles and AB -simple dual cycles, whereby c is as well.

We next claim that any AB -simple dual cycle is also endpoint simple unless it is the concatenation of two AB -cycles. Combined with the discussion above, this will show that each directed cycle is an integer linear combination of AB -chains and AB -cycles, proving the lemma by Lemma 5.8.

For this, suppose that c is an AB -simple dual cycle. We may assume that the AB -length of c is greater than 1 since otherwise c is an AB -chain or AB -cycle. As before, let $c = (p_1, \dots, p_k)$ be the AB -decomposition of c . Let c_i be the AB -cycle containing p_i . By AB -simplicity, each c_i is distinct.

By symmetry it suffices to show that the terminal vertex v of p_1 appears exactly once in c unless c is the concatenation of two dual cycles. Suppose v appears again in the directed cycle c . We split the possibilities into four cases.

Case 1: v belongs to p_j for $j \neq 1, 2$. In this case v is contained in the AB -cycle c_j and by AB -simplicity c_1, c_2, c_j are all distinct. This is impossible since any Γ -vertex is contained in at most two distinct AB -cycles.

Case 2: v is an interior vertex of either p_1 or p_2 . In this case either c_1 or c_2 must meet the vertex v twice, which is impossible since v meets c_1 and c_2 once each and so meets neither twice.

Case 3: v is the initial vertex of p_1 . This implies v is the terminal vertex of p_k , so v belongs to the AB -cycles c_k, c_1, c_2 . We conclude that $k = 2$ so $c = (p_1, p_2)$ and each of p_1 and p_2 is an AB -cycle as claimed.

Case 4: v is the terminal vertex of p_2 . A symmetric argument to Case 3 shows that in this case c is also a concatenation of two AB -cycles.

This proves our claim that every AB -simple dual cycle is either endpoint simple or the concatenation of two AB -cycles, completing the proof of the lemma. ■

In the proof of Theorem 6.1 we will also need this short lemma about relatively prime elements in $\mathbb{Z}[G]$:

Lemma 6.17. *In the group ring $\mathbb{Z}[G]$, $1 \pm g$ and $1 \pm h$ are relatively prime if $\langle g, h \rangle \leq G$ has rank 2.*

Proof. Since $1 - g^2 = (1 + g)(1 - g)$, it suffices to prove that $1 - g$ and $1 - h$ are relatively prime. Let n be the rank of G , and write $G = \langle u \rangle \oplus \langle t \rangle \oplus \mathbb{Z}^{n-2}$, where $g = u^a$ and $h = u^b t^c$ for $a, c \geq 1$ and $b \geq 0$. Begin by factoring

$$1 - g = 1 - u^a = \prod_{d|a} \Phi_d(u),$$

where $\Phi_d(u)$ is the d th cyclotomic polynomial in the variable u . Note that each Φ_d is irreducible in $\mathbb{Z}[G]$ since $R_d = \mathbb{Z}[G]/(\Phi_d) \cong \mathbb{Z}[G/\langle u \rangle][\omega_d]$ is a domain. Here ω_d is a primitive d th root of unity.

Since $\mathbb{Z}[G]$ is a UFD, it suffices to show that $\Phi_d(u)$ does not divide $1 - h$ for any d . If it did, then $1 - h$ would be zero in the ring R_d . But by further quotienting $R_d \rightarrow \mathbb{Z}[G/\langle u \rangle]$, we have

$$1 - h = 1 - t^c = \prod_{d|c} \Phi_d(t) \quad \text{in } \mathbb{Z}[G/\langle u \rangle],$$

which is clearly nonzero since t is primitive in $G/\langle u \rangle$. ■

6.6. Completing the proof

We can now assemble the proof of this section's main theorem on factoring V_τ .

Proof of Theorem 6.1. For any AB -cycle c with homology class g , Let \mathcal{C}_c be the set of cycles other than c , which by Lemma 6.11 is a reducing family. Corollary 6.13 gives us

$$\left(\Theta_\tau \cdot \prod_{c_i \in \mathcal{C}_c} (1 \pm g_i) \right) \mid V_\tau,$$

and multiplying by $1 \pm g$ gives

$$(V^{AB} \cdot \Theta_\tau) \mid (1 \pm g) \cdot V_\tau.$$

This, together with Corollary 6.15, implies that

$$(V^{AB} \cdot \Theta_\tau) \mid p \cdot V_\tau, \tag{14}$$

where p is the gcd of all polynomials of the form $1 \pm g$ for g the homology class of an AB -cycle and $1 \pm z$ for z the homology class of an AB -chain. Since $\text{rank}(H_1(M)) > 1$, Lemma 6.16 implies there are such homology classes that are independent in $H_1(M)$ (i.e. do not generate a rank 1 subgroup). By Lemma 6.17, the corresponding polynomials are relatively prime and so p is a unit in $\mathbb{Z}[G]$. Combining this with Proposition 6.6 completes the proof. ■

Remark 6.18. If $\text{rank}(H_1(M)) = 1$, so that $G = \langle t \rangle$, the conclusion of Theorem 6.1 still holds after possibly multiplying V_τ by $1 \pm t$. To see this, first note that Proposition 6.6 is unaffected by the rank assumption. Second, as in the proof of Theorem 6.1, equation (14) holds where p is the gcd of polynomials of the form $1 \pm t^{k_i}$, with each t^{k_i} corresponding to an AB -chain or AB -cycle. The homology classes of AB -chains and AB -cycles still generate $H_1(M)$ by Lemma 6.16, so the gcd of the k_i is 1. Factoring $1 - t^{k_i}$ and $1 + t^{k_i}$ into cyclotomics, we see that the only possible divisors of all the $1 \pm t^{k_i}$ are $\Phi_1 = 1 - t$ and $\Phi_2 = 1 + t$. However, at most one of Φ_1, Φ_2 is a common divisor of the $1 \pm t^{k_i}$ since $1 - t \nmid 1 + t^{k_i}$ and $1 + t \nmid 1 - t^{k_i}$ for k_i odd.

7. The fibered case and the Teichmüller polynomial

In this section, we examine the veering polynomial V_τ in the case of a layered triangulation τ , when τ is canonically associated to a *fibered face* \mathbf{F}_τ of the Thurston norm ball (as in Theorem 5.15). Ultimately, we demonstrate the connection between the veering polynomial and McMullen’s Teichmüller polynomial. We first do this for the manifolds admitting veering triangulations (Theorem 7.1) and then derive the general case (Proposition 7.2) using the fact that any hyperbolic fibered manifold admits a veering triangulation after puncturing along the singular orbits of its associated suspension flow. We conclude by showing how these results combine with those of the previous section to give a “determinant formula” for computing the Teichmüller polynomial using only the veering data (Corollary 7.3).

Throughout this section, all veering triangulations are layered.

7.1. Teichmüller polynomials from veering triangulations

Fix a layered veering triangulation τ of a manifold M and denote its associated fibered face by $\mathbf{F} = \mathbf{F}_\tau$. Since our goal is to introduce a veering interpretation of the Teichmüller polynomial, we closely follow McMullen’s paper [33] and refer the reader there for the complete construction.

Let \mathcal{L} be the two-dimensional expanding lamination of M associated to \mathbf{F} , and $\tilde{\mathcal{L}}$ its lift to \tilde{M}^{ab} , the free abelian cover of M with deck group G . We recall that up to isotopy \mathcal{L} is obtained by suspending the expanding lamination of the monodromy for any fiber in $\mathbb{R}_+\mathbf{F}$ [33, Corollary 3.2]. To orient the reader, in this article we mainly use the terms unstable/stable rather than expanding/contracting, respectively, although their meanings are the same.

Associated to $\tilde{\mathcal{L}}$ and its G action is a *module of transversals* $T(\tilde{\mathcal{L}})$, which like the modules of Section 3 is a $\mathbb{Z}[G]$ -module. We will not require its precise definition here but instead rely on several of its properties that will be recalled below. Let $\Theta_{\mathbf{F}}$ be the Teichmüller polynomial associated to \mathbf{F} . By definition, $\Theta_{\mathbf{F}}$ is the gcd of the elements of the Fitting ideal for $T(\tilde{\mathcal{L}})$ and is hence well-defined up to a unit $\pm g \in \mathbb{Z}[G]$.

We will show in Theorem 7.1 that $\Theta_{\mathbf{F}} = \Theta_\tau$ up to a unit, where Θ_τ is the taut polynomial for τ defined in Section 3.1.

Before stating Theorem 7.1, we address two technical issues arising from the different conventions used to define Teichmüller and taut polynomials. First, our definition of the face module $\mathcal{E}^\Delta(\tilde{\tau})$ from Section 3.1 is most naturally related to the module of transversals for the *contracting lamination* dual to \mathcal{L} . To remedy this, we define a $\mathbb{Z}[G]$ -module $\mathcal{E}^\nabla(\tilde{\tau})$ that is a variation on the definition of the face module. For this, we note that each face f lies at the top of a unique tetrahedron and we refer to the edge \mathbf{t} at the top of this tetrahedron as the *top* edge of f . The $\mathbb{Z}[G]$ -module $\mathcal{E}^\nabla(\tilde{\tau})$ is defined exactly as $\mathcal{E}^\Delta(\tilde{\tau})$ but with the following modification: for each face f with top edge \mathbf{t} , instead of the relation from equation (3) we use the relation

$$\mathbf{t} = \mathbf{w} + \mathbf{z}, \quad (15)$$

where \mathbf{w}, \mathbf{z} are the other two edges of f . As shown by Parlak [40], the $\mathbb{Z}[G]$ -modules $\mathcal{E}^\Delta(\tilde{\tau})$ and $\mathcal{E}^\nabla(\tilde{\tau})$ are isomorphic. Indeed, an isomorphism is induced by the map $\mathbb{Z}[G]^E \rightarrow \mathbb{Z}[G]^E$ which sends an edge \mathbf{e} to itself if it is right veering and to $-\mathbf{e}$ if it is left veering. For each face f , this map exchanges the relation from equation (3) with the relation from (15). This follows easily from two facts: (1) every face f of τ has edges of each veer, and (2) the bottom and top edges of f have the same veer. The first fact is immediate from the veering definition and the second follows from Fact 3.3. Thus we get a $\mathbb{Z}[G]$ -module homomorphism which is easily seen to be invertible.

Second, to define his module $T(\tilde{\mathcal{L}})$, McMullen uses the action of G on transversals by taking preimages under deck transformations, whereas our action of G on the edges of $\tilde{\tau}$ (which form a special class of transversals) is by taking images. To deal with this discrepancy, we introduce the group isomorphism $\text{inv}: G \rightarrow G$ defined by $\text{inv}(g) = g^{-1}$. We extend this to a ring isomorphism $\text{inv}: \mathbb{Z}[G] \rightarrow \mathbb{Z}[G]$ and use it to define a $\mathbb{Z}[G]$ -module $\mathcal{E}_*^\Delta(\tilde{\tau})$ which as a \mathbb{Z} -module is equal to the face module $\mathcal{E}^\Delta(\tau)$ from Section 3.1 and whose $\mathbb{Z}[G]$ -module structure is determined by

$$g \cdot x = \text{inv}(g)x,$$

where $g \in G$ and $x \in \mathcal{E}_*^\Delta(\tau)$. In words, we are simply replacing the action of an element by its inverse. We define $\mathcal{E}_*^\nabla(\tilde{\tau})$ similarly.

From the definitions, it is clear that the gcd of the elements of the Fitting ideal of $\mathcal{E}_*^\Delta(\tilde{\tau})$ (and of $\mathcal{E}_*^\nabla(\tilde{\tau})$ by the discussion above) is equal to $\text{inv}(\Theta_\tau)$. In what follows, we will use the symmetry of the Teichmüller polynomial [33, Corollary 4.3] which states that $\Theta_F = \text{inv}(\Theta_F)$, up to a unit $\pm g \in \mathbb{Z}[G]$.

Theorem 7.1. *Let τ be a layered veering triangulation of M representing a fibered face \mathbf{F} . The module of transversals $T(\tilde{\mathcal{L}})$ is isomorphic as a $\mathbb{Z}[G]$ -module to $\mathcal{E}_*^\nabla(\tilde{\tau})$ and hence to $\mathcal{E}_*^\Delta(\tilde{\tau})$. In particular,*

$$\Theta_\tau = \Theta_F \quad \text{up to a unit in } \mathbb{Z}[G].$$

Before the proof we need some additional preliminaries. Let S be a fiber of M representing a class in the cone over \mathbf{F} . Let $\psi: S \rightarrow S$ be the corresponding monodromy so

that M can be recovered as the mapping torus of ψ :

$$M = \frac{S \times [0, 1]}{(s, 1) \sim (\psi(s), 0)}.$$

We denote ψ 's expanding lamination by λ so that the suspension of λ is isotopic to \mathcal{L} . Either directly from Agol's construction [2], from [35, Lemma 3.2], or from Theorem 5.15 above, the fiber S is carried by τ up to isotopy. We fix any carrying map of S into $\tau^{(2)}$. Then if we pull back edges and faces we obtain an ideal triangulation \mathcal{T} of S and a simplicial map $S \rightarrow M$. Such a map is called a *section* in the terminology of [35]. The ideal triangulation \mathcal{T} of S is dual to a train track \mathcal{V} on S which is obtained by pulling back the intersection of S with the *unstable* branched surface B^u . We observe that with our setup, the track \mathcal{V} carries λ and is an invariant track for ψ in the sense that $\psi(\mathcal{V})$ is carried by \mathcal{V} . This can be seen by noting that the sequence of upward diagonal exchanges from \mathcal{T} to itself in τ corresponds to a sequence of diagonal exchanges from $\psi(T)$ to T on S , and this sequence is dual to a *folding* sequence from $\psi(\mathcal{V})$ to \mathcal{V} (see e.g. Figure 13.)

Next we follow a modified version of the discussion in [33, Section 3]. An elevation \tilde{S} of S to \tilde{M}^{ab} is the cover of S corresponding to the kernel of the homomorphism $\pi_1(S) \rightarrow H_1(S) \rightarrow G$, whose image we denote by H . Fix a lift $\tilde{\psi}: \tilde{S} \rightarrow \tilde{S}$ and, using that M is the mapping torus of ψ , split $G = H \oplus \mathbb{Z}u$. Here, u acts on $\tilde{M}^{\text{ab}} = \tilde{S} \times \mathbb{R}$ as the deck transformation $\tilde{\Psi}$ mapping $(s, t) \mapsto (\tilde{\psi}(s), a(s, t))$, where $a(\cdot, \cdot)$ is real-valued function such that for each $s \in \tilde{S}$ and $t \in \mathbb{R}$, $a(\psi(s), t) \leq a(s, t)$. (Since the fixed section $S \rightarrow M$ may not be an embedding, we cannot necessarily make the conventional choice $a(s, t) = t - 1$.)

Similarly, McMullen shows [33, Theorem 3.5] that if $\tilde{\lambda}$ denotes the lift of λ to \tilde{S} and $T(\tilde{\lambda})$ is the $\mathbb{Z}[H]$ -module of transversals for $\tilde{\lambda}$, then $\tilde{\mathcal{L}} = \tilde{\lambda} \times \mathbb{R}$ and $T(\tilde{\mathcal{L}}) = T(\tilde{\lambda})$ as $\mathbb{Z}[H]$ -modules. Then by considering the action of u on $T(\tilde{\lambda})$, McMullen establishes that

$$T(\tilde{\mathcal{L}}) \cong \frac{T(\tilde{\lambda}) \otimes \mathbb{Z}[u]}{\text{im}(uI - \tilde{\psi}^*)},$$

as $\mathbb{Z}[H][u] = \mathbb{Z}[G]$ -modules. Here $\tilde{\psi}^*$ is the action of $\tilde{\psi}$ on $T(\tilde{\lambda})$ defined by taking preimages of transversals.

To connect the discussion back to the veering triangulation τ , we recall a consequence of Guéritaud's construction [23] of τ in the layered setting. We refer the reader to [35] for additional details. Fix a ψ -invariant quadratic differential q on S and let \tilde{q} be its lift to \tilde{S} . Then there is a projection $\Pi: \tilde{M}^{\text{ab}} \rightarrow \tilde{S}$ that maps edges of $\tilde{\tau}$ to saddle connections of \tilde{q} . More precisely, Π induces a bijection between edges of $\tilde{\tau}$ and saddle connections of \tilde{q} that span singularity-free, immersed euclidean rectangles whose vertical/horizontal sides are segments of the vertical/horizontal foliations of \tilde{q} . The projection Π is also equivariant in the sense that $\Pi(\tilde{\Psi}(\tilde{e})) = \tilde{\psi}(\Pi(\tilde{e}))$. In particular, $\Pi(\tilde{e})$ can be naturally regarded as a transversal of $\tilde{\lambda}$, which is the lamination associated to the horizontal foliation of \tilde{q} . See Figure 15.

We now turn to the proof of Theorem 7.1.




Fig. 15. A face f with top edge t . Each edge of f is dual to a branch of the intersection $f \cap B^u$ (left) and is a transversal to the expanding lamination (right).

Proof of Theorem 7.1. Let $\mathcal{E}_H^\nabla(\tilde{\tau})$ denote the $\mathbb{Z}[H]$ -module obtained from the $\mathbb{Z}[G]$ -module $\mathcal{E}_*^\nabla(\tilde{\tau})$ by restricting scalars. For each edge \tilde{e} of $\tilde{\tau}$, $\Pi(\tilde{e})$ is a transversal of $\tilde{\lambda}$ and this induces a homomorphism from the free module on edges of $\tilde{\tau}$ to the free module on transversals. This is a $\mathbb{Z}[H]$ -module homomorphism because the action on each is given by pullback under deck transformations. Moreover, this homomorphism descends to a $\mathbb{Z}[H]$ -module homomorphism $\Pi_*: \mathcal{E}_H^\nabla(\tilde{\tau}) \rightarrow T(\tilde{\lambda})$. Indeed, it is easily checked that relations among edges of $\tilde{\tau}$ from (15) are satisfied in $T(\tilde{\lambda})$; see Figure 15. In fact, we have:

Claim 3. $\Pi_*: \mathcal{E}_H^\nabla(\tilde{\tau}) \rightarrow T(\tilde{\lambda})$ is an isomorphism of $\mathbb{Z}[H]$ -modules.

Proof of claim. The preimage of \mathcal{V} in \tilde{S} is a $\tilde{\psi}$ -invariant train track $\tilde{\mathcal{V}}$ which is dual to the ideal triangulation $\tilde{\mathcal{T}}$ of \tilde{S} induced by the simplicial map $\tilde{S} \rightarrow \tilde{M}^{\text{ab}}$. Mapping each branch of $\tilde{\mathcal{V}}$ to its dual edge of $\tilde{\mathcal{T}}$, regarded as an edge of $\tilde{\tau}$, induces a $\mathbb{Z}[H]$ -module homomorphism $\mathcal{D}: T(\tilde{\mathcal{V}}) \rightarrow \mathcal{E}_H^\nabla(\tilde{\tau})$. Here $T(\tilde{\mathcal{V}})$ is the $\mathbb{Z}[H]$ -module generated by the edges of $\tilde{\mathcal{V}}$ modulo relations imposed by the switch conditions—these relations are mapped by \mathcal{D} to the relations from (15) used to define $\mathcal{E}_H^\nabla(\tilde{\tau})$. See the left-hand side of Figure 15.

The composition $\Pi_* \circ \mathcal{D}: T(\tilde{\mathcal{V}}) \rightarrow T(\tilde{\lambda})$ can alternatively be described as follows: If $\tilde{f}: \tilde{\lambda} \rightarrow \tilde{\mathcal{V}}$ is the carrying map, then the branch b is mapped to $f^{-1}(x)$, where x is any point in the interior of b . Hence, [33, Theorem 2.5] proves that the composition $\Pi_* \circ \mathcal{D}: T(\tilde{\mathcal{V}}) \rightarrow T(\tilde{\lambda})$ is an isomorphism of $\mathbb{Z}[H]$ -modules. This immediately gives that Π_* is surjective. Injectivity will follow from the fact that \mathcal{D} is surjective.

To prove that \mathcal{D} is surjective, we show the more general statement that for any ideal triangulation $\tilde{\mathcal{T}}$ of \tilde{S} corresponding to a simplicial map $\tilde{S} \rightarrow \tilde{M}^{\text{ab}}$ lifting a section $S \rightarrow M$, the edges of $\tilde{\mathcal{T}}$ generate $\mathcal{E}_H^\nabla(\tilde{\tau})$. This follows from two facts. First, if $\tilde{\mathcal{T}} \rightarrow \tilde{\mathcal{T}}'$ is a diagonal exchange corresponding to pushing through a single tetrahedron t of $\tilde{\tau}$, then the new edge of $\tilde{\mathcal{T}}'$ is a sum or difference of two edges of $\tilde{\mathcal{T}}$ that belong to t (depending on whether the tetrahedron lies above or below $\tilde{\mathcal{T}}$). Second, for any edge \tilde{e} of $\tilde{\tau}$ there is triangulation $\tilde{\mathcal{T}}'$ of \tilde{S} containing \tilde{e} and a finite sequence $\tilde{\mathcal{T}} = \tilde{\mathcal{T}}_1, \dots, \tilde{\mathcal{T}}_n = \tilde{\mathcal{T}}'$ so that each $\tilde{\mathcal{T}}_i \rightarrow \tilde{\mathcal{T}}_{i+1}$ corresponds to simultaneously pushing through an H -orbit of tetrahedra of $\tilde{\tau}$ [35, Lemma 3.2 and Proposition 3.3]. These two facts show that every edge of $\tilde{\tau}$ is a linear combination of edges of $\tilde{\mathcal{T}}$, so \mathcal{D} is surjective as desired. ■

Returning to the proof of the theorem, we note that since $\Pi(\tilde{\Psi}^{-1}(\tilde{e})) = \tilde{\psi}^*(\Pi(\tilde{e}))$, the isomorphism Π_* takes the action of $\tilde{\Psi}^{-1}$ on $\mathcal{E}_H^\nabla(\tilde{\tau})$ to the action of $\tilde{\psi}^*$ on $T(\lambda)$. Hence, it extends to an isomorphism of $\mathbb{Z}[H] \otimes \mathbb{Z}[u] = \mathbb{Z}[G]$ -modules:

$$\frac{\mathcal{E}_H^\nabla(\tilde{\tau}) \otimes \mathbb{Z}[u]}{\text{im}(uI - \tilde{\Psi}^{-1})} \rightarrow \frac{T(\lambda) \otimes \mathbb{Z}[u]}{\text{im}(uI - \tilde{\psi}^*)},$$

where the second module is $T(\tilde{\mathcal{L}})$, as explained above. Since the first module simply reintroduces the action of u , it is isomorphic to $\mathcal{E}_*^\nabla(\tilde{\tau})$. Hence, $\mathcal{E}_*^\nabla(\tilde{\tau})$ is isomorphic to $T(\tilde{\mathcal{L}})$ as a $\mathbb{Z}[G]$ -module.

Using the remarks preceding the theorem, we also find that $\mathcal{E}_*^\Delta(\tilde{\tau})$ is isomorphic to $T(\tilde{\mathcal{L}})$ as a $\mathbb{Z}[G]$ -module and we conclude that $\text{inv}(\Theta_\tau) = \Theta_F = \text{inv}(\Theta_F)$ and so $\Theta_\tau = \Theta_F$, up to a unit. This completes the proof of Theorem 7.1. \blacksquare

7.2. The general case via puncturing

On its face Theorem 7.1 applies only when starting with a fibered face associated to a veering triangulation. However, the following proposition shows that the constructions are compatible with puncturing along orbits of the associated suspension flow.

Let N be any hyperbolic 3-manifold with fibered face \mathbf{F}_N . There is a flow φ such that each fibration associated to $\mathbb{R}_+ \mathbf{F}_N$ may be isotoped so that the first return map to a fiber is pseudo-Anosov and that, up to reparametrization, φ is the associated suspension flow [16, Theorem 14.11]. The orbits of the pseudo-Anosov's singularities are called the *singular orbits* of φ .

Now let M be the result of puncturing N along the singular orbits of φ and let τ be the associated veering triangulation of M (see Section 2.2). Let $i: M \hookrightarrow N$ be the inclusion map. Then $i^*: H^1(N) \rightarrow H^1(M)$ maps the cone over \mathbf{F}_N into the cone over a fibered face $\mathbf{F}_M := \mathbf{F}_\tau$ of M .

The first statement of the next proposition is observed by McMullen to prove [33, equation 6.1]. We provide some details using results from [33, Section 4].

Proposition 7.2. *With the setup as above, $\Theta_{\mathbf{F}_N} = i_*(\Theta_{\mathbf{F}_M})$ up to a unit. Hence, if τ is the layered veering triangulation on M associated to the fibered face \mathbf{F}_M , then*

$$\Theta_{\mathbf{F}_N} = i_*(\Theta_\tau) \quad \text{up to a unit } \pm g \in \mathbb{Z}[G].$$

Proof. The second claim follows immediately from the first and Theorem 7.1.

We recall that if Θ is the Teichmüller polynomial associated to a fibered face \mathbf{F} , then there is a unique $g \in \text{supp}(\Theta)$ such that for all α in the interior of $\mathbb{R}_+ \mathbf{F}$,

$$\alpha(g) > \alpha(h) \quad \text{for all } h \in \text{supp}(\Theta) \setminus \{g\}.$$

This follows from the proof of [33, Theorem 6.1]; see also item 5 in the subsection “Information packaged in Θ_F ” of [33, Section 1]. Hence, we can normalize Θ by multiplying by the unit $\pm g$ so that for any α as above, $\Theta(u^\alpha)$ has a positive constant term and all other

terms have negative exponent (here we are thinking of α as a first cohomology class). We perform this normalization on both $\Theta_{\mathbf{F}_N}$ and $\Theta_{\mathbf{F}_M}$.

As in the proof of [33, Corollary 4.3], to show that $\Theta_{\mathbf{F}_N} = i_*(\Theta_{\mathbf{F}_M})$ it suffices to show that for each integral $\alpha = [S] \in \mathbb{R}_+ \mathbf{F}_N$, there is the equality of specializations $\Theta_{\mathbf{F}_N}(u^\alpha) = i_*(\Theta_{\mathbf{F}_M})(u^\alpha)$. This is because $\mathbb{R}_+ \mathbf{F}_N$ is open and so one can find such an α so that the values $\alpha(g)$ are distinct for all $g \in \text{supp}(\Theta_{\mathbf{F}_N}) \cup \text{supp}(i_*(\Theta_{\mathbf{F}_M}))$.

Next we note that $i_*(\Theta_{\mathbf{F}_M})(u^\alpha) = \Theta_{\mathbf{F}_M}(u^{i^*\alpha})$. The pullback $i^*\alpha$ is dual to the class $[\mathring{S}]$, where \mathring{S} is obtained from S by puncturing at the singularities of its monodromy ψ . Hence, by [33, Theorem 4.2], up to units $\pm u^k$, $\Theta_{\mathbf{F}_N}(u^\alpha)$ is equal to the characteristic polynomial of ψ and $\Theta_{\mathbf{F}_M}(u^{i^*\alpha})$ is equal to the characteristic polynomial of $\mathring{\psi}$, the induced monodromy of \mathring{S} (unless the two-dimensional lamination \mathcal{L} is orientable, in which case we multiply first by $u - 1$ in both cases). However, the characteristic polynomial from [33] is defined solely in terms of ψ 's action on its expanding lamination λ , and since we have punctured S at an ψ -invariant set in the complement of λ , λ is also the expanding lamination of $\mathring{\psi}$. Hence, the maps $\psi: \lambda \rightarrow \lambda$ and $\mathring{\psi}: \lambda \rightarrow \lambda$ are equal. (This can also be seen by considering the intersection of \mathcal{L} with S and \mathring{S} in N .) We conclude that their characteristic polynomials are equal and therefore $\Theta_{\mathbf{F}_N}(u^\alpha) = \Theta_{\mathbf{F}_M}(u^{i^*\alpha})$ up to a unit $\pm u^k$. However, it is clear from our normalizations that this unit must be the identity, and the proof is complete. ■

7.3. Computing $\Theta_{\mathbf{F}}$ via V_τ

Here we state an analog of McMullen's determinant formula which follows immediately from Theorems 7.1 and 6.1. By further applying Proposition 7.2, we obtain a general method to compute the Teichmüller polynomial from the veering triangulation on the associated fully punctured manifold.

Corollary 7.3. *Let (M, τ) be a layered veering triangulation associated to a fibered face \mathbf{F} and suppose that $\text{rank}(H_1(M)) > 1$. Then*

$$\Theta_{\mathbf{F}} = \frac{V_\tau}{V_\tau^{AB}} = \frac{\det L}{\det L^{AB}},$$

where L is as in Section 3 and L^{AB} is as in Section 6.

Note that in this case (i.e. when τ is layered), $V_\tau^{AB} \neq 0$ since no AB -cycle (nor any dual cycle) can be trivial in H_1 by Theorem 5.15.

Remark 7.4. Some care is required to compare Corollary 7.3 with McMullen's determinant formula [33, Theorem 3.6]. For this we first define, given a veering triangulation τ , its negative $-\tau$ to be the veering triangulation obtained by reversing the coorientation on faces. With this definition, we have $\mathcal{E}^\nabla(\tilde{\tau}) = \mathcal{E}^\Delta(\neg\tilde{\tau})$.

If the train track \mathcal{V} from above is used in McMullen's construction, then his determinant formula states that

$$\Theta_{\mathbf{F}} = \frac{\det(uI - P_E)}{\det(uI - P_V)},$$

where P_E and P_V are the matrices with entries in $\mathbb{Z}[G]$ that represent the action of the lifted monodromy on the branches and switches of the lifted track \tilde{V} , respectively, as defined in [33, Section 3]. One can directly show that $V_{-\tau} = \det(uI - P_E)$ up to a unit in $\mathbb{Z}[G]$. One would then expect that $V_{-\tau}^{AB} = \det(uI - P_V)$ up to a unit. However, this is the case only if the implied definition of P_V in McMullen’s paper is interpreted with signs chosen according to whether the train track carrying map associated to the monodromy preserves or reverses the implicit orientation at switches. With these signs taken appropriately, McMullen’s equation 3.4 commutes as claimed.

Acknowledgments. We thank Anna Parlak for interesting conversations on the subject of this paper and for pointing out an algebraic mistake in an earlier draft. We also thank Spencer Dowdall, Curtis McMullen, and Henry Segerman for helpful feedback as well as the referee for useful corrections.

Funding. This work was partially supported by NSF grants DMS-1610827, DMS-1744551, and the Sloan Foundation.

References

- [1] Agol, I.: [Criteria for virtual fibering](#). J. Topol. **1**, 269–284 (2008) Zbl [1148.57023](#) MR [2399130](#)
- [2] Agol, I.: [Ideal triangulations of pseudo-Anosov mapping tori](#). In: Topology and geometry in dimension three, Contemporary Mathematics 560, American Mathematical Society, Providence, RI, 1–17 (2011) Zbl [1335.57026](#) MR [2866919](#)
- [3] Agol, I.: Comparing layered triangulations of 3-manifolds which fiber over the circle. MathOverflow discussion, <http://mathoverflow.net/questions/106426> (2012)
- [4] Algom-Kfir, Y., Hironaka, E., Rafi, K.: [Digraphs and cycle polynomials for free-by-cyclic groups](#). Geom. Topol. **19**, 1111–1154 (2015) Zbl [1354.57007](#) MR [3336279](#)
- [5] Baik, H., Wu, C., Kim, K., Jo, T.: [An algorithm to compute the Teichmüller polynomial from matrices](#). Geom. Dedicata **204**, 175–189 (2020) Zbl [1436.57029](#) MR [4056698](#)
- [6] Bell, M.: flipper (computer software). pypi.python.org/pypi/flipper (2013–2018)
- [7] Bell, M. C.: The pseudo-Anosov and conjugacy problems are in $\text{NP} \cap \text{co-NP}$. arXiv:[1410.1358](#) (2014)
- [8] Billet, R., Liechti, L.: Teichmüller polynomials of fibered alternating links. Osaka J. Math. **56**, 787–806 (2019) Zbl [1437.57025](#) MR [4020637](#)
- [9] Boyd, S., Vandenberghe, L.: [Convex optimization](#). Cambridge University Press, Cambridge (2004) Zbl [1058.90049](#) MR [2061575](#)
- [10] Calegari, D.: [Problems in foliations and laminations of 3-manifolds](#). In: Topology and geometry of manifolds (Athens, GA, 2001), Proceedings of Symposia in Pure Mathematics 71, American Mathematical Society, Providence, RI, 297–335 (2003) Zbl [1042.57501](#) MR [2024640](#)
- [11] Casson, A. J., Bleiler, S. A.: [Automorphisms of surfaces after Nielsen and Thurston](#). London Mathematical Society Student Texts 9, Cambridge University Press, Cambridge (1988) Zbl [0649.57008](#) MR [964685](#)
- [12] Cooper, D., Futer, D.: [Ubiquitous quasi-Fuchsian surfaces in cusped hyperbolic 3-manifolds](#). Geom. Topol. **23**, 241–298 (2019) Zbl [1444.57013](#) MR [3921320](#)
- [13] Fenley, S. R.: [Foliations with good geometry](#). J. Amer. Math. Soc. **12**, 619–676 (1999) Zbl [0930.53024](#) MR [1674739](#)
- [14] Fenley, S. R.: [Surfaces transverse to pseudo-Anosov flows and virtual fibers in 3-manifolds](#). Topology **38**, 823–859 (1999) Zbl [0926.57009](#) MR [1679801](#)

- [15] Floyd, W., Oertel, U.: [Incompressible surfaces via branched surfaces](#). *Topology* **23**, 117–125 (1984) Zbl [0524.57008](#) MR [721458](#)
- [16] Fried, D.: Fibrations over S^1 with pseudo-Anosov monodromy. In: *Travaux de Thurston sur les surfaces*, Astérisque **66-67**, 251–266 (1979) Zbl [0446.57023](#)
- [17] Fried, D.: [The geometry of cross sections to flows](#). *Topology* **21**, 353–371 (1982) Zbl [0594.58041](#) MR [670741](#)
- [18] Fulton, W.: [Introduction to toric varieties](#). *Annals of Mathematics Studies* 131, Princeton University Press, Princeton, NJ (1993) Zbl [0813.14039](#) MR [1234037](#)
- [19] Futer, D., Guéritaud, F.: [Explicit angle structures for veering triangulations](#). *Algebr. Geom. Topol.* **13**, 205–235 (2013) Zbl [1270.57054](#) MR [3031641](#)
- [20] Futer, D., Taylor, S. J., Worden, W.: [Random veering triangulations are not geometric](#). *Groups Geom. Dynam.* **14**, 1077–1126 (2020) Zbl [1467.57011](#) MR [4167033](#)
- [21] Giannopoulos, A., Schleimer, S., Segerman, H.: A census of veering structures. <https://math.okstate.edu/people/segerman/veering.html> (2020)
- [22] Groves, D., Manning, J. F.: [Quasiconvexity and Dehn filling](#). *Amer. J. Math.* **143**, 95–124 (2021) Zbl [1460.57020](#) MR [4201780](#)
- [23] Guéritaud, F.: [Veering triangulations and Cannon–Thurston maps](#). *J. Topol.* **9**, 957–983 (2016) Zbl [1354.57025](#) MR [3551845](#)
- [24] Hatcher, A.: *Vector bundles and K-theory*. Online book, <http://pi.math.cornell.edu/~hatcher/VBKT/VB.pdf> (2017)
- [25] Hodgson, C. D., Rubinstein, J. H., Segerman, H., Tillmann, S.: [Veering triangulations admit strict angle structures](#). *Geom. Topol.* **15**, 2073–2089 (2011) Zbl [1246.57034](#) MR [2860987](#)
- [26] Lackenby, M.: [Taut ideal triangulations of 3-manifolds](#). *Geom. Topol.* **4**, 369–395 (2000) Zbl [0958.57019](#) MR [1790190](#)
- [27] Landry, M.: [Taut branched surfaces from veering triangulations](#). *Algebr. Geom. Topol.* **18**, 1089–1114 (2018) Zbl [1396.57033](#) MR [3773749](#)
- [28] Landry, M.: [Stable loops and almost transverse surfaces](#). *Groups Geom. Dynam.* **17**, 35–75 (2023) Zbl [07672076](#) MR [4563328](#)
- [29] Landry, M. P.: [Veering triangulations and the Thurston norm: homology to isotopy](#). *Adv. Math.* **396**, art. 108102, 53 pp. (2022) Zbl [1483.57016](#) MR [4370465](#)
- [30] Landry, M. P., Minsky, Y. N., Taylor, S. J.: [Flows, growth rates, and the veering polynomial](#). *Ergodic Theory Dynam. Systems* (online, 2022)
- [31] Lanneau, E., Valdez, F.: [Computing the Teichmüller polynomial](#). *J. Eur. Math. Soc.* **19**, 3867–3910 (2017) Zbl [1380.37070](#) MR [3730517](#)
- [32] Mazur, B., Wiles, A.: [Class fields of abelian extensions of \$\mathbb{Q}\$](#) . *Invent. Math.* **76**, 179–330 (1984) Zbl [0545.12005](#) MR [742853](#)
- [33] McMullen, C. T.: [Polynomial invariants for fibered 3-manifolds and Teichmüller geodesics for foliations](#). *Ann. Sci. École Norm. Sup. (4)* **33**, 519–560 (2000) Zbl [1013.57010](#) MR [1832823](#)
- [34] McMullen, C. T.: [Entropy and the clique polynomial](#). *J. Topol.* **8**, 184–212 (2015) Zbl [1353.37033](#) MR [3335252](#)
- [35] Minsky, Y. N., Taylor, S. J.: [Fibered faces, veering triangulations, and the arc complex](#). *Geom. Funct. Anal.* **27**, 1450–1496 (2017) Zbl [1385.57025](#) MR [3737367](#)
- [36] Mosher, L.: [Surfaces and branched surfaces transverse to pseudo-Anosov flows on 3-manifolds](#). *J. Differential Geom.* **34**, 1–36 (1991) Zbl [0754.58031](#) MR [1114450](#)
- [37] Mosher, L.: [Dynamical systems and the homology norm of a 3-manifold. II](#). *Invent. Math.* **107**, 243–281 (1992) Zbl [0758.58027](#) MR [1144424](#)
- [38] Northcott, D. G.: [Finite free resolutions](#). *Cambridge Tracts in Mathematics* 71, Cambridge University Press, Cambridge (1976) Zbl [0328.13010](#) MR [0460383](#)

- [39] Oertel, U.: [Incompressible branched surfaces](#). *Invent. Math.* **76**, 385–410 (1984) Zbl [0539.57006](#) MR [746535](#)
- [40] Parlak, A.: [Computation of the taut, the veering and the Teichmüller polynomials](#). *Experimental Math.* (online, 2021)
- [41] Parlak, A.: [The taut polynomial and the Alexander polynomial](#). *J. Topol.* **16**, 720–756 (2023)
- [42] Schleimer, S., Segerman, H.: From veering triangulations to link spaces and back again. [arXiv:1911.00006](#) (2019)
- [43] Schleimer, S., Segerman, H.: [Essential loops in taut ideal triangulations](#). *Algebr. Geom. Topol.* **20**, 487–501 (2020) Zbl [1437.57023](#) MR [4071381](#)
- [44] Sullivan, D.: [Cycles for the dynamical study of foliated manifolds and complex manifolds](#). *Invent. Math.* **36**, 225–255 (1976) Zbl [0335.57015](#) MR [433464](#)
- [45] Thurston, W. P.: A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.* **59**, no. 339, i–vi and 99–130 (1986) Zbl [0585.57006](#) MR [823443](#)
- [46] Wise, D. T.: [The structure of groups with a quasiconvex hierarchy](#). *Annals of Mathematics Studies* 209, Princeton University Press, Princeton, NJ (2021) Zbl [07286269](#) MR [4298722](#)